

# Health Index methods and development: 2015 to 2021

The Health Index measures health outcomes and risk factors over time, for different geographic areas. This methodology explains how we have constructed this tool.

Contact:  
Greg Ceely  
Health.Data@ons.gov.uk  
+44 207 5928692

Release date:  
16 June 2023

Next release:  
To be announced

## Table of contents

1. [Overview](#)
2. [Handling data changes](#)
3. [Aims of the Health Index](#)
4. [Future developments](#)
5. [Process for constructing the Health Index](#)
6. [Theoretical framework \(COIN Step 1\)](#)
7. [Data selection \(COIN Step 2\)](#)
8. [Methods overview](#)
9. [Geographical aggregation](#)
10. [Imputation of missing data \(COIN Step 3\)](#)
11. [Multivariate analysis \(COIN Step 4\)](#)
12. [Homogenising the data \(Normalisation, COIN Step 5\)](#)
13. [Weighting \(COIN Step 6\)](#)
14. [Aggregating indicators \(COIN Step 7\)](#)
15. [Sensitivity analysis \(COIN Step 8\)](#)
16. [The scale of the Health Index](#)
17. [Expert Advisory Group members](#)

18. [Related links](#)

19. [Cite this methodology](#)

# 1 . Overview

- This methodology article accompanies two releases presenting the Health Index results at local and national levels, and an article detailing the indicators contained within the Health Index.
- The Health Index has been designed with the support of health experts to present a single number measuring the health of an area, with a clear breakdown of how different measures of health are combined to produce this value.
- The development of the Health Index has followed guidance by the Competence Centre on Composite Indicators and Scoreboards (COIN).
- Data have been selected from a wide variety of sources to allow comparisons across time and by geography, down to lower-tier local authority (LTLA) level.
- Data selection has been based on agreed principles, such as the aim to measure health and its drivers rather than direct measures of health services.
- Factor analysis has been used to group individual indicators of health into subdomains, guided by expert advice; factor analysis results have informed each indicator's weight towards the total Index value.

## 2 . Handling data changes

### Handling missing data for 2020 and 2021

The coronavirus (COVID-19) pandemic has naturally affected public health. It also forced changes to how some health variables were measured and whether certain indicators could be measured at all.

While COVID-19 and related restrictions have had the most obvious effect on 2020 and 2021 results and data collection, it is important to note that there may be other events or longer-term trends, which have led to differences between results.

Of the 56 indicators that form the Health Index, 21 were unavailable or unusable in 2021, compared with 12 in 2020 and two in 2019. In 2021 results, this includes four indicators for which the underlying data are not yet published but are expected to be later in the year. Prioritising timeliness for this release makes the Health Index more useful.

At the indicator level, all missing values were handled following our existing imputation methodology, which in these cases means 2019 or 2020 scores were often held constant. There is more detail and information about how we impute missing values in [Section 10: Imputation of missing data \(COIN Step 3\)](#).

The children and young people subdomain was especially affected by the pandemic, with only one and two out of five indicators available for 2021 and 2020, respectively. Our conversations with public health experts suggested the available indicators present a different trend to that expected for the indicators held constant.

Measures, such as pupil absenteeism, were unavailable but are expected to have increased, which would have contributed to a lower Health Index score. For this reason, we have decided to hold the children and young people subdomain itself constant for 2020 and 2021. This avoids misleading users about the impact of the pandemic on this subdomain. The available indicators are included in the Health Index for 2020 and 2021, but they do not affect their subdomain score for this year.

In 2021 we replicated the process above for the access to services subdomain. In this subdomain only one indicator had usable data, and we would expect to see variation in the other indicators if they were not imputed. To avoid presenting misleading results the subdomain is held constant in 2021, even though the indicators are included in the Health Index for 2021.

## Updating data

The existing data were reviewed to ensure we use the most up-to-date and accurate data. For 11 indicators, the back series was updated where data producers had published new versions of their data. This was because of updates to methodology, improvements or revisions made by the producer since we last collected the data or because of re-weighting to more recent population estimates. These indicators are:

- all four personal well-being indicators
- cancer screening (bowel screening component)
- child poverty
- disability
- healthy eating
- job-related training
- overweight and obesity in adults
- suicides

For a further four indicators, the back series was updated as part of our calculations of the Health Index, including how we handled the data to produce our indicators. These indicators are:

- personal crime - to remove bike theft and shoplifting, because these are already counted in the low-level crime indicator
- rough sleeping - to keep in true zeros, where these were imputed previously
- sedentary behaviour - using the correct year convention as applied to other indicators
- sexually transmitted infections (diagnosis and test components) - new data provided through Fingertips

The size and extent of any of these changes varies by indicator. Detailed information about the changes to each indicator are available on request.

## Handling data and population estimates following Census 2021

The Health Index 2021 includes the time-period in which Census 2021 data is available. As a result of the census, we have handled our data in varying ways because of the individual effect this has had on our indicators and associated mid-year population estimates.

Household overcrowding data for 2021 is now available and the 2011 and 2021 data are comparable, allowing us to build a more accurate reflection of household overcrowding in each year. There was previously one datapoint for 2011, which was being used as the 2015 baseline and every other year's score, in line with our imputation methods. This meant household overcrowding scores were stable at 100 for England, and lower-tier local authorities (LTLAs) retained the same score every year.

We used linear interpolation with the 2011 and 2021 data to estimate the level of household overcrowding in 2015. The estimated value for 2015 was used as the baseline in 2015 (a score of 100 for England). This means that the household overcrowding indicator has a time series of data from 2015 in the Health Index 2021, where scores are not held constant across the years, at any geography, in the Health Index.

Population estimates had the largest impact on our Health Index 2021. The Office for National Statistics (ONS) population estimates have shown differences between Census 2021 data and estimates rolled forward since the 2011 Census. Population estimates are important in our methods, and we use them in relation to our geographies and regional medians to help us create the Health Index scores from indicator to country level.

To prevent the Health Index scores showing artificial increases and decreases in scores because of the change from 2020 mid-year estimated population to Census population in 2021, we have re-used 2020 population estimates for 2021. Both the Integrated Care Systems (ICS) scores and our Health Index scores have applied this same method. The release of Health Index 2022 will include a revised back-series using ONS' latest population estimates.

The impact of the census population is also reflected in a few of the indicators we use in the Health Index. Where data providers use ONS population estimates to produce their data, the difference in population for census compared with the population estimates up to 2020 has led to an inconsistent time series. In these cases, we are using the previously published back series up to 2020 and using our imputation method for 2021. The indicators that are affected are alcohol misuse, frailty and self-harm, all sourced from Hospital Episode Statistics (HES) produced by NHS England.

Once new population estimates are available, the data provider will update the back series and the data will be usable from 2021 onwards.

## 3 . Aims of the Health Index

### The Index's origin

The proposal for a Health Index was made in the 2018 annual report of the government's then Chief Medical Officer (CMO), Dame Sally Davies, entitled [Health 2040 - better health within reach](#).

The Office for National Statistics (ONS) aims to develop the Health Index into a regular publication allowing differences in health to be tracked over time. The work to develop the Health Index so far has been completed in consultation with an Expert Advisory Group (EAG) consisting of representatives from a range of government, academic and third sector organisations. A list of members is available in [Section 17: Expert Advisory Group members](#). Representatives from the health departments of the four UK nations have been involved in its development, with a view to extending coverage beyond England in the future.

The initial "beta" version release was accompanied by a public consultation to collect views on how the Health Index could be improved, which we have been addressing.

## How the Health Index differs from existing products

The Health Index provides a single number headline health indicator to act as a guide for policy and public focus, incorporating changes in health and the factors that influence it. There is no established example of a health index of the type we are currently developing, in England or elsewhere.

In terms of the existing health indicator landscape, there are multiple frameworks in use.

The ONS publishes both health statistics and wider national measures relevant to health, such as:

- [Health state life expectancies](#)
- [Subnational indicators explorer](#)
- [Sustainable Development Goals \(SDGs\)](#)

Publications by other UK government bodies:

- [Adult Social Care Outcomes Framework \(ASCOF\) by NHS England](#)
- [Index of Multiple Deprivation \(IMD\), by the Department of Levelling Up, Housing and Communities](#), which contains a health domain as one of its components
- [NHS Outcomes Framework \(NHS OF\) by NHS England](#)
- [Public Health Outcomes Framework, by the Office for Health Improvement & Disparities](#)
- [Quality and Outcomes Framework \(QOF\) by NHS England](#)

Frameworks used in other countries:

- [County Health Rankings in the US, by the University of Wisconsin](#)
- [Global Burden of Disease, by the independent Institute for Health Metrics and Evaluation](#)
- [New Zealand well-being statistics, by the state statistics authority Stats NZ](#), looking beyond direct well-being questions to topics such as housing conditions

These frameworks all have different purposes and uses, and most contain elements of all three domains defined for the Health Index, described in [Section 6: Theoretical framework \(COIN Step 1\)](#). What the Health Index offers, which these sources individually do not, is a single headline measure of health that is transparent in its construction and can be compared over time. Indicators within the Health Index are grouped into meaningful subdomains and domains, allowing users to trace the drivers of changes in scores. The Health Index can also be compared at different geographical levels.

## Potential users of the Health Index

We expect there to be three broad groups of people using the Health Index, including:

- the media and general public
- policymakers and analysts, in government and local government
- analysts and decision makers outside of government

The media and general public can present and see the headline measures as an indicator of change in the nation's health, and of inequality in health between different geographical areas. This should help people understand their own local area and promote public engagement with health issues.

Policymakers in government and local government can clearly identify which topics related to health are not improving over time, and measure health impacts when assessing policies. The Health Index enables measurement of impacts on health to become more regular and consistent. Local government decision makers can compare health in their area with other places of their choosing, such as areas with similar characteristics, and learn about differences between them.

Analysts outside government, such as academics and those in think tanks and charities, can improve the body of evidence on different aspects of health and the stories this can tell us.

## 4 . Future developments

There are still further refinements that we will make to the Health Index, and we are working towards making the Health Index a National Statistic.

There will be further methods investigations following expert technical review, with some details noted where relevant in [Sections 8 to 16](#).

We will continuously review the data used to ensure we are always including as many relevant concepts as we can and measuring them in the best ways that we can.

We recognise that providing data that are as timely as possible is important to aid use. As such, we are working to improve the timeliness of the index. Data from 2019 were published in March 2022, data from 2020 were published in November 2022, and data from 2021 were published in June 2023. We are looking at ways we can further improve this. The extent to which we can directly improve timeliness for the Health Index overall is limited by when the underlying data used to construct it are available. We are exploring options such as making the index values for individual indicators available between publications of the full Index.

Related work also aims to:

- develop a health projections model to estimate how the Health Index and its components may change in the future, and simulate simple "what if?" scenarios
- expand the Health Index to include analysis of all four nations in the UK
- enable the production of local health indices

## 5 . Process for constructing the Health Index

Our process to construct the Health Index for England largely follows that outlined in the [Organisation for Economic Co-operation and Development \(OECD\) and Joint Research Centre's \(JRC\) Handbook on constructing composite indicators](#) and subsequently, in the [European Commission's 10-Step Guide to Composite Indicators and Scoreboards](#).

The steps included in this guide are:

1. theoretical framework
2. data selection
3. imputation of missing data
4. multivariate analysis
5. normalisation
6. weighting
7. aggregating indicators
8. sensitivity analysis
9. link to other measures
10. visualisation

This article will focus on Steps 1 to 8 and is structured as such. We have renamed Step 5 to "Homogenising the data" to reflect that scale-based transformations are also involved here. Links to other measures are discussed in [Section 3: Aims of the Health Index](#), so are not detailed here.

## 6 . Theoretical framework (COIN Step 1)

The concept of health that the Health Index covers is largely derived from the Chief Medical Officer's (CMO) original recommendation in [Health 2040 - better health within reach](#), which suggested that the Index should be "inclusive of health outcome measures, modifiable risk factors and the social determinants of health".

This encompasses the [World Health Organization's definition of health](#), that "health is a state of complete physical, mental and social well-being, and not merely the absence of disease or infirmity", and adds specificity to the idea of well-being.

The theoretical framework that the CMO alluded to is well known in public health and epidemiology. It divides factors influencing health into three categories.

"Health status or outcomes" is about the occurrence of illness and health events themselves. This includes morbidity measures such as disease prevalence, as well as mortality, life expectancy, and wider well-being measures.

"Modifiable risk factors" (MRFs) refers to things that affect health that can potentially be changed at individual level. This means health-related behaviours (for example, smoking and exercise) and actionable clinical findings (for example, blood pressure). However, it is important to understand that these factors are in the middle of a bigger causal chain.

"Wider social drivers of health" means circumstances that have a major effect on life chances -- including both MRFs and health outcomes - but cannot be addressed at individual level. Examples include:



- unemployment rates
- quality of transport infrastructure
- environmental pollution

Dahlgren and Whitehead's "rainbow" diagram in Policies and strategies to promote social equity in health from the Institute for Future Studies, Stockholm, 1991, is often used to illustrate the relationships between these different factors affecting health.

Considering the definitions mentioned previously, elements of health are divided into three domains for the Health Index, each corresponding to one of these three categories:

- Healthy People - health outcomes
- Healthy Lives - health-related behaviours and personal circumstances
- Healthy Places - wider drivers of health that relate to the places people live

## 7 . Data selection (COIN Step 2)

In our [Health Index contents and definitions](#), it details the indicators that make up the Health Index. Health Index metadata are available in the [Health Index datasets](#).

We conducted a comprehensive review of existing indices and frameworks that had a relation to our chosen definition of health, as discussed in [Section 6: Theoretical framework \(COIN Step 1\)](#). The aim was to understand what content they included, and what of that was relevant to the Health Index.

Following this, we reviewed the wider literature to understand whether there was additional content that the Health Index should include, as its aims, functions and purpose differ from these other products. In conjunction with both steps, a range of data sources that could potentially be used to measure these concepts were identified. We reported these initial proposals for Health Index content to the Expert Advisory Group (EAG) to gain their feedback on the concepts included, how they were measured and whether there were additional concepts we should add.

Using this feedback, a detailed review of the content proposed for inclusion was carried out. This included a critical review of how these should be measured and what data were available to construct the Index presenting those concepts. At all stages of this process, the aim was to maintain a balance between concept and data, ensuring the use of the most appropriate measure without unduly compromising on data quality.

Additional concepts were also added to the Health Index based on consultation feedback. Proposed revisions were shared with the EAG, and their suggestions were also incorporated.

## Data requirements for quality

The data that have been selected to develop the Index have largely come from already-published sources, as this means certain quality standards will already have been met. There are also several requirements for data to be included in the Health Index, focused on ensuring data are consistent and comparable.

Data must be available for enough years to make comparisons over time; currently this means 2015 to 2021. There are some exceptions to this, if it is reasonable to assume that the value being measured would not change much from year to year but will still differ between areas.

There must be reasonable certainty that the data will continue to be produced into the future. This helps the Health Index to remain consistent and comparable in the long term.

To enable local comparisons, data must be available at the level of lower-tier local authority areas (LTLAs). This is the smallest geographical scale available for most health data sources suitable for the Index's needs. The median LTLA in England currently has a population of around 140,000 people, though the range is quite broad.

Ensuring we measure health itself and its drivers, rather than healthcare activity, service performance or policy, was central to designing the Health Index. This has been considered in both the inclusion of concepts and in the ways in which they are measured. Some data sources have been ruled out because they are too directly linked to one of these aspects. For example, if the Index were to include the number of people receiving adult social care as a measure, the overall figure representing health would change if the national thresholds for social care eligibility changed, even if the nation's health did not actually get better or worse.

There are some concepts deemed important enough to be presented in the Index, even where the data could not meet all our requirements. In these instances, careful considerations have been made to understand whether the benefits of their inclusion outweigh the limitations. An example is household overcrowding. Data at the geography levels we need are only available through the census, so they are updated only once every ten years. We do, however, see household overcrowding as sufficiently important to include in the Health Index.

## Handling different timespans

For the Health Index, we want to be able to report results for calendar years for consistency with other Office for National Statistics (ONS) health statistics. However, not all data sources are published on this basis. For example, some statistics are presented in financial years or academic years.

Where data differ from calendar years, we have assigned the data to the year in which most of the source period falls. For example, data in the financial year April 2016 to March 2017 would be used in the Health Index as representing the 2016 calendar year.

For some data sources, three-year aggregates are used to present the data, where counts for individual years would risk being disclosive or volatile. In such cases, to maximise timeliness of Health Index releases, we present three-year averages as the value for the most recent of those three years in the Health Index. If the data producer also uses a three-year aggregate, but to represent the central year of the three, this can lead to inconsistencies between what the Health Index and source data state the score or trend is for an individual year.

## 8 . Methods overview

The methods chosen to construct the Health Index result from extensive research and consultation with experts. Sensitivity analyses were performed to explore how the use of other possible methods could potentially affect Index scores.

The sections that follow detail each step taken to create the Health Index, including:

- geographical aggregation ([Section 9](#))
- imputation of missing data ([Section 10](#))
- multivariate analysis ([Section 11](#))
- homogenising the data ([Section 12](#))
- weighting ([Section 13](#))
- aggregating indicators ([Section 14](#))
- sensitivity analysis ([Section 15](#))
- scaling ([Section 16](#))

## 9 . Geographical aggregation

The Health Index is presented at country, region, upper-tier local authority (UTLA) and lower-tier local authority (LTLA) level using 2021 geographic boundaries, and Integrated Care Systems (ICS) level. There are currently 309 LTLAs, which combine to form 151 UTLAs, which combine to form the nine regions of England.

For the purposes of the Health Index, results for the Isles of Scilly (LTLA code E06000053) and City of London (E09000001) LTLAs are not included for any indicators. This is because exceptionally small populations lead to small sample sizes, which results in unreliable data. For some sources, these LTLAs' data are grouped with nearby LTLAs in the source data. These are Cornwall (E06000052) for the Isles of Scilly, and Hackney (E09000012) for City of London. Where this is the case, we have not made an adjustment to separate them, as their small population size means the impact of their inclusion is minimal.

Not all data sources use these geographies in their original publications. Some sources, such as the GP Patient Survey (GPPS), present data based on health geographies - in this case, GP practices. Some education data are presented using individual schools, and the crime data used in the Health Index are presented using Community Safety Partnerships (CSPs). These data need separate handling to convert them to geographies consistent with the other indicators.

Data collected for individual GP practices have GP practice codes. Postcodes for GP practices are published by [NHS England](#). This enables GP practices to be grouped by LTLA using the [National Statistics Postcode Lookup \(NSPL\)](#). Similarly, schools can be aggregated to LTLA level using the postcode of the school site. The Health Index therefore assumes the patients and children registered at GP practices and schools reside in the same LTLA as the GP or school they attend. This is broadly but not always true. If any GP or school has missing data for the number of patients or children, respectively, these are excluded from the analysis.

There are 301 CSP police force areas in England. Most of these can be mapped to individual LTLAs using the Office for National statistics (ONS) [Open Geography Portal - Local Authority District to Community Safety Partnerships to Police Force Areas \(December 2022\) Lookup in England and Wales](#). Where one CSP represents multiple LTLAs, the total crime counts are split based on population sizes. This means each of those LTLAs will have the same derived crime rates. Where one LTLA is represented by multiple CSPs, the numbers of crimes in those CSPs are summed to receive a total number for that LTLA. Rates are then derived based on that LTLA's population.

There were no changes to the boundaries or structure of LTLAs from 2015 to 2018. In 2019, 2020 and 2021 there were mergers of LTLAs to form unitary authorities (UAs) or non-metropolitan districts. Some former counties have been abolished and replaced by the resulting UA. There have been no splits of existing LTLAs, which means that in principle all changes can be handled by combining available data. The areas that have changed are:

### 2019:

- Dorset UA (code E06000059) was created from a merger of five non-metropolitan districts (E07000049-53)
- Bournemouth, Christchurch and Poole UA (E06000058) was created from a merger of one non-metropolitan district (E07000048) and two UAs (E06000028 and E06000029)
- the county of Dorset (E10000009), which had comprised the non-metropolitan districts E07000048-53 that were merged into the two UAs above, was abolished
- East Suffolk non-metropolitan district (E07000244) was created from two non-metropolitan districts (E07000205 and E07000206)
- West Suffolk non-metropolitan district (E07000245) was created from a merger of two non-metropolitan districts (E07000201 and E07000204)
- Somerset West and Taunton (E07000246) was created from two non-metropolitan districts (E07000190 and E07000191)

## **2020:**

- Buckinghamshire UA (E06000060) was created from a merger of four non-metropolitan districts (E07000004-7)
- the county of Buckinghamshire (E10000002), which was made of the same four non-metropolitan districts as the new UA, was abolished

## **2021:**

- North Northamptonshire LTLA (E06000061) was created from a merger of four LTLAs (E07000150, E07000152, E07000153, E07000156)
- West Northamptonshire LTLA (E06000062) was created from a merger of three LTLAs (E07000151, E07000154, E07000155)

Aggregation to LTLA geographies used the following process, consistent with previous guidance for the Office for Health Improvement and Disparities (OHID) Fingertips tool.

Method 1 is used if numerators (for example, the number of rough sleepers) and denominators (for example, the total population of the LTLA) are available for all areas that form the new area. These can then be summed and converted into a new combined statistic (for in this case, a percentage).

Method 2 only looks at overall values for each area. The contribution of each area to the combined whole is estimated by multiplying each area's value by the fraction of the combined population present in that area. These adjusted values are then summed.

Method 1 calculates an aggregated value, while Method 2 provides an estimate. Each indicator was aggregated using Method 1 or 2 as appropriate.

If an indicator's source statistic was not a rate or a percentage, Method 2 was automatically used. Method 2 was also used to give an estimate where the statistic was an age-standardised rate.

Some denominators are not based on the whole population of the local authority, or not based on population at all. When using Method 2 in these cases, by making estimations based on population proportions we are assuming that the more specific denominator also follows these proportions. For example, if the denominator is the number of people aged 65 years and over, we are assuming that the areas involved in the merger all have the same proportion of elderly residents.

Considering the boundary changes to LTLAs detailed earlier in this section, these two methods allow results to be calculated in the absence of data with 2021 boundaries. For years prior to and during the transition to these changes, where data are not available for the LTLAs as they exist in 2021, the new areas can be calculated or estimated from the corresponding non-metropolitan districts and UAs. The populations of the previous geographies (from Office for National Statistics mid-year population estimates) are applied to weight the components when aggregating to these new LTLAs. This follows Method 1.

Dorset is a special case, as the new Dorset UA is almost identical to the former county of Dorset, but with Christchurch split off to join Bournemouth and Poole. If data for Dorset are only available in terms of the old county of Dorset, we can still use the population-proportion Method 2 to estimate values for Dorset UA and Bournemouth, Christchurch and Poole UA.

Our future development plans include investigating calculation of age-standardised values for indicators with data, which are not currently age-standardised. In this case, Method 1 would use the population pyramid by age for each of the geographies to calculate the resulting value.

## **10 . Imputation of missing data (COIN Step 3)**

Our approach to imputation is typically simple imputation, in line with how other indices handle missing values. Multiple imputation is being considered to refine our methods in future but is not expected to impact results substantially.

The detail of the approach used, presented in the order that steps were applied, is:

- if we had known values either side of one or more missing values for a lower-tier local authority (LTLA), the missing values were calculated as a linear interpolation of the values either side
- if one or more values were missing with a known value only on one side of the time series, that is, because the missing values were at the start or end of the time series, missing values were replaced with the nearest adjacent value
- if a value was missing for an LTLA for all years, we imputed the median for its region

Once fully constructed, the Health Index consists of 56 indicators, 307 LTLA records and seven time points (2015 to 2021). This amounts to approximately 120,000 data points in total. Just over 22% of all records were affected by imputation in some way (26,593 data points, or 22.1%), either because their value was imputed or one of their components was imputed.

The modal reason for imputation was that the indicator data were not available for an entire year of our time series. This was the case for over 26,000 of the affected records, or 21.7% of the total Health Index data. The regional median imputation method was only required for Rutland LTLA's suicide indicator score. These were the only seven data points to use regional median imputation.

## 11 . Multivariate analysis (COIN Step 4)

Typically for a composite index, we would aim to avoid collinearity (that is, very high correlations) between indicators as that suggests they are measuring the same or very similar concepts, and including both is redundant.

However, by the nature of the Health Index's aims of presenting health at multiple levels, and being transparent in its construction, the Index looks to capture multiple indicators measuring similar principles and cluster these into subdomains for comparison. We also expect many of our indicators to be correlated because the Health Index includes both risk factors and the outcomes that we expect are associated with those risk factors.

While conducting factor analysis to group indicators into the Index's subdomains, as described in [Section 13: Weighting \(COIN Step 6\)](#), we assessed correlation matrices of all indicators within each domain. We used this together with factor analysis when multiple similar data options were available to capture the same concept, to assess which was a better fit for the Index as a whole. We also used the combination of correlations and indicator weights from factor analysis to inform whether some indicators were suitable to include at all.

## 12 . Homogenising the data (Normalisation, COIN Step 5)

It is necessary when constructing an index to transform all indicators to a homogenous scale.

### Scaling

Since the population of different lower-tier local authorities (LTLAs) varies, indicator data are measured in terms of rates and percentages, rather than raw counts. We used Office for National Statistics (ONS) population estimates for this purpose. We have used age-standardised rates where they were applicable and available, but this was only possible for a minority of the data sources used.

To construct an Index from comparable indicators, we needed to ensure the direction of change was consistent. This was not always the case in the underlying data. For example, an increase in healthy eating is better for health, but an increase in smoking is worse. We have adjusted certain indicators so that for all indicators a higher value corresponds with better health. This simply involved multiplying some indicators' values by negative one.

## Normalisation

Factor analysis is used to organise the data into subdomains in a later step, and an underlying assumption of factor analysis is that indicators have a normal distribution. Therefore, when homogenising the data, we needed to address any outliers and then assess the skew and kurtosis of each indicator.

Winsorization was used to bring extreme outliers to within three standard deviations of the mean for that indicator. Then, to assess whether transformations were required, we examined the skew and kurtosis of each indicator, for each year, with threshold values of positive or negative 0.5 and positive or negative 2, respectively. Where values fell outside either threshold, we applied a number of commonly used transformation methods (log, square root, cube root, cube, negative reciprocal) and selected the method which most effectively reduced the skewness and kurtosis of the indicator.

Of the 56 indicators, 21 did not require transformation. Log transformations were applied to 15 indicators. The remaining 20 indicators were transformed via one of the other listed functions.

Data presented in the Health Index itself are not winsorized. This is because, though winsorization increases normality, it also involves altering extreme datapoints. Publishing winsorized values could therefore be misleading, especially for particular LTLA values. For example, if an area were performing exceptionally poorly, its score would be artificially boosted by increasing it to exactly three standard deviations below the mean. If that area then improved or declined while still remaining below the threshold, its value would appear unchanged even though this is not the case.

## Standardisation

The methods available for standardisation are narrowed greatly by the Health Index's need to be comparable across time and geographic area, with additional years of data not affecting the back series values. For this reason, time series standardisation has been used.

Regular standardisation involves subtracting the mean value and dividing by the standard deviation, for each indicator. For a time series like the Health Index, this is not suitable: the mean and standard deviation would change with each additional year of data, which would cause confusion by changing past index scores. If standardisation were applied within individual years, the resultant values would no longer be comparable across years, which is an important attribute of the Health Index.

We follow the [COIN \(2020\) 10-step-guide's](#) recommendation of applying the mean and standard deviation for a base year across the whole time series. This allows for comparisons across time and only causes back series changes if the reference year is updated. This is a common practice used across a number of national statistics. For the Health Index, the base year is 2015.

## 13 . Weighting (COIN Step 6)

The Health Index's hierarchical structure means there are multiple levels at which weighting can be relevant, each requiring separate consideration. Indicators must be weighted within their subdomain; subdomains must be weighted within their domain and the domains must be weighted within the overall Index.

## Weighting indicators within subdomains: time series factor analysis

A fundamental assumption of factor analysis is that there is a latent factor that underpins the variables in a group. In terms of the Health Index, that means a single unobserved variable underpinning the indicators within each subdomain.

The indicators within each subdomain will likely be highly correlated, which could lead to double counting in the Index. Factor analysis directly addresses this issue, accounting for the correlation between indicators in their implied weights. It also groups indicators into subdomains based on statistical information, and not just theorised concepts.

As with the normalisation methods, factor analysis cannot be used in its regular form to meet this Index's aims. If the factor analysis were carried out across all observations, the weights would change with each additional year of data. As such, the weights need to be calculated for a set time period, and these weights are held constant until a review date.

The indicators to be included within each domain were decided based on a theoretical approach, with input from the Expert Advisory Group (EAG) and using consultation feedback. We then conducted a separate factor analysis for each domain's bundle of indicators. We also ran factor analysis on all indicators together to see if alternative groupings emerged.

The theoretical position of indicators is not always clear-cut, so conducting factor analysis helped guide which domains indicators fit into best. For example, some indicators are both health outcomes and risk factors - such as children's social, emotional and mental health. As such they could be placed in Healthy People as health outcomes, or Healthy Lives or Places, as risk factors.

We assessed the most suitable results for grouping into subdomains using our correlation matrices and hypothesised indicator groupings. Where groupings were surprising, we re-ran factor analysis using only specific variables to confirm the subdomains we are presenting would not split out into separate factors (subdomains) if allowed to. Each indicator could only be included in one subdomain even if it loaded onto multiple factors, for ease of user interpretation.

Where indicators did not load as expected in our initial hypotheses, we critically considered the sources used for those indicators to check they were measuring the intended information and tested the indicator in different subdomains.

Each indicator's factor loading is the amount of the latent factor (subdomain) variance which that indicator can explain. Weights were constructed for each indicator within each subdomain using the scaled factor loadings within that subdomain. For example, if a subdomain had two indicators with factor loadings of 0.7 and 0.5, respectively, one indicator would receive a weight of 0.7 divided by 1.2 and the other of 0.5 divided by 1.2. The weights for each indicator are presented in our downloadable [Health Index datasets](#).

## Limitations of factor analysis

There are limitations involved with using factor analysis. This method only accounts for the collinearity between indicators and does not derive any measure of the importance of the indicators (COIN, 2020). Furthermore, this method gives lower weight to indicators that are not highly correlated with others. But the low correlation between indicators is often the exact reason why an index is being created, because it suggests the indicator that is not well correlated with others is measuring a different aspect of the whole. There are also subjective choices made within the process that affect the final weights.

## Weighting subdomains within domains: equal weighting

For the purposes of this version of the Health Index, all subdomains have equal weighting within their domain. This means that as the Healthy People domain has five subdomains, each subdomain has a weight of one-fifth of the overall domain. Healthy Lives has four subdomains, so each subdomain has a weight of one-fourth of the overall domain. Healthy Places has five subdomains, so each subdomain has a weight of one-fifth of the overall domain.

This method will be refined in future, as we plan to use a budget allocation process for subdomain weights. We are in the process of consulting our Expert Advisory Group (EAG) and the Faculty of Public Health, who will use their expertise in public health to assign weights based on relative importance of subdomains within each domain.



## Weighting domains to the overall Health Index score: equal weighting

Equal weighting is used to weight the three domains. The Health Index's aim is to offer a broad measure of health and not focus simply on health outcomes. Weighting each of these domains equally satisfies this.

## 14 . Aggregating indicators (COIN Step 7)

The Health Index is aggregated geographically, from lower-tier local authority (LTLA) to upper-tier local authority (UTLA), region, and England as a whole. It is also aggregated to higher index levels, from indicator to subdomain, domain, and overall Health Index score.

The Health Index is aggregated using linear aggregation. Linear aggregation involves taking the (weighted) arithmetic mean of indicators to calculate the Index. This is the simplest aggregation method; however, it introduces compensability into the composite index. This means that poor performance in one area can be offset by good performance elsewhere (COIN, 2020).

Indicators are aggregated into subdomain scores using the weights described in [Section 13: Weighting \(COIN Step 6\)](#). Subdomains are equally weighted within each domain, and each domain is equally weighted to produce an overall Health Index score.

For all levels of the Health Index, when aggregating geographically, lower-tier local authority level areas (LTLAs) are weighted by their mid-year population estimate for that year.

Integrated Care System (ICS) results were calculated by aggregating LTLA results. Some LTLAs spanned two ICSs, so the proportion of the LTLA population in each ICS was used in aggregating scores. LTLA to ICS data were provided by the Office for Health Improvement and Disparities (OHID) for years 2018 to 2020. Each Index year's ICS scores were produced with respect to the same year's LTLA to ICS data, but the population data for 2018 were used in calculating ICS index scores for 2015 to 2017. Given the population estimate issues affecting the Health Index ([Section 2: handling data changes](#)), the ICS population estimates for 2021 re-used 2020 population estimates while we await updated population estimates from the Office for National Statistics (ONS) following the Census 2021.

## 15 . Sensitivity analysis (COIN Step 8)

We performed a series of analyses to investigate whether changes to the methods used in [Sections 12 to 14](#) would have substantial effects on the Health Index results, known as sensitivity analysis.

The methods tested in this way, and the alternatives tested, are:

- normalisation using standardisation or min-max scaling
- indicator weighting using factor analysis or equal weights
- indicator, subdomain and domain aggregation: linear, geometric or weighted-adjusted Mazziotta-Pareto indexing

For most steps, although small differences were seen when using different methods, these made no statistically significant difference to the Index scores. Correlations between resultant scores from each pair of methods were above 0.9.

The Office for National Statistics (ONS) has collaborated with The Alan Turing Institute to quality assure the Health Index methods. This includes more detailed sensitivity analysis, such as how ranks between local authorities change as a result of different methods. The Alan Turing Institute aims to publish their assessment of the Health Index methodology, including these findings.

In future, sensitivity analysis will also be undertaken to ensure that the indicator weights produced using factor analysis do not alter greatly when they are derived using different time periods.

## 16 . The scale of the Health Index

The Health Index has been scaled to a base of 100 for England, with a base year of 2015. Values higher than 100 indicate better health than England in 2015, and values below 100 indicate worse health. The scale is such that a score of 110 represents a score that is one 2015 standard deviation higher than England 2015's score for that same indicator, a score of 80 is two standard deviations lower, and so on.

In this way, comparisons both over time and within a single year are simple to understand. The Index scores should also be "future proof" such that a score must be 10 standard deviations lower than England 2015 in order to be zero or become negative.

The weighted arithmetic mean used for geographic aggregation, as described in [Section 14: Aggregating indicators \(COIN Step 7\)](#), leads to the spread of results reducing for higher geographies. The range of scores at regional or national level for any element of the Health Index is smaller than the range at lower-tier local authority (LTLA) level. These higher geographies cannot be re-scaled to have a score difference of 10 representing one standard deviation, because there are too few observations to calculate a robust standard deviation. For example, there will only be nine regional scores per element of the Health Index per year, because there are nine regions of England.

Regional and national-level scores can therefore be interpreted in the same way, in relation to health in England in 2015; more than 100 is better, 100 is the same, and less than 100 is worse. However, more than one standard deviation of the data at these levels has scores between 90 and 110.

## 17 . Expert Advisory Group members

Our Expert Advisory Group (EAG) includes the following UK government departments and arms-length bodies:

- Cabinet Office
- Department for Environment, Food and Rural Affairs (Defra)
- Department for Levelling Up, Housing and Communities (DLUHC)
- Department for Transport (DfT)
- Department of Health and Social Care (DHSC)
- London Health Partnership
- NHS England
- National Institute for Health and Care Excellence (NICE)
- Northern Ireland Health Department
- Office for Health Improvement and Disparities (OHID)
- Office for National Statistics (ONS)
- Public Health Wales
- Scottish Government
- Welsh Government

The group also includes the following organisations, which are not part of the UK government:

- Alan Turing Institute
- Association of Directors of Public Health (ADPH)
- Health Foundation
- Institute for Fiscal Studies (IFS)
- Institute for Social and Economic Research (ISER)
- King's Fund
- Organisation for Economic Co-operation and Development (OECD)
- Royal Society for Public Health (RSPH)
- University College London

We extend our thanks to all members for their valuable input into the Health Index's development.

## 18 . Related links

### [Health in England: 2015 to 2021](#)

Bulletin | Released 16 June 2023

Insights into England's health during the coronavirus (COVID-19) pandemic at national and regional level, using the Health Index for England.

### [How health has changed in your local area: 2015 to 2021](#)

Digital article | Released 16 June 2023

Use our interactive tool to explore how health has changed in each local authority area across England between 2015 and 2021.

### [Health Index contents and definitions](#)

Methodology | Released 9 November 2022

The Health Index for England, which measures the health of the nation; definitions of the indicators used.

## 19 . Cite this methodology

Office for National Statistics (ONS), released 16 June 2023, ONS website, Methodology, [Health Index methods and development: 2015 to 2021](#)