

WID DATA QUALITY

MANUEL ARIAS-OSORIO
IGNACIO FLORES
ROWAIDA MOSHRIF
ANA VAN DER REE

TECHNICAL NOTE N°2026/05

OCTOBER 2026

WORLD
INEQUALITY
..... LAB

WID Data Quality

Technical Note

Manuel Arias-Osorio¹, Ignacio Flores², Rowaida Moshrif³, Ana Van Der Ree⁴
September 2026

¹ World Inequality Lab (WIL); manuel-esteban.arias-osorio@psemail.eu

² Paris School of Economics, Università degli studi Roma Tre; ignacio.flores@psemail.eu

³ WIL; rowaida.moshrif@psemail.eu

⁴ WIL; ana.vanderree@psemail.eu

Table of Contents

1. Introduction	3
Info box: Where to find the data quality on WID.world	4
2. Grades Descriptions	5
Short descriptions	5
Table 1. Distributed concepts	5
Table 2. Aggregate concepts	6
3. How to turn data quality grades in a binary indicator	6
4. Special grading rules for post-tax income series	8
5. Special grading rules for fiscal income series	8
6. Special grading rules for aggregate macroeconomic series	9
6.1. Aggregate income	9
6.1.1. In-house calculations	10
6.1.2. Macroeconomic identities enforcement	11
6.1.3. Extensions to onelets a, m, and w	11
6.2. Aggregate wealth	12
6.3. Indices and exchange rates	12
6.3.1. MER exchange rates	12
6.3.2. PPP exchange rates	14
6.3.3. Real exchange rates	14
7. Special grading rules when a variable has both distributed and aggregate observations	15
8.1. Distributional regional series	16
8.2. Macroeconomic regional aggregates	16
9. Summarizing data quality as a weighted average	17
9.1. Time-decay weighting of data quality series	17
9.2. Simulation of logistic weighting functions	18
9.3. Parameter values used for average quality score calculation	19
References	20
Appendix	21
Long descriptions	21
Table 3. Distributed concepts	21
Table 4. Aggregate concepts	24

1. Introduction

The accurate measurement of income and wealth inequality in WID.world relies on the release of detailed microdata, aggregated tabulations, and survey results by statistical agencies and governments. To make sense of the diverse information that is available, and which we use in our estimates, we assign data quality grades from 0 to 5 at the observational level for series. The data quality grades vary by year within a country-series combination. This time dimension enables users to track how data quality evolves over time, spot improvements, identify gaps in coverage, and compare trends across countries.

Observational data quality has been developed for the following distributional series: pre-tax income, household wealth, post-tax national income, post-tax disposable income, female labour income share, and fiscal income. It has equally been developed for the following aggregate series: all aggregate income subcomponents, aggregate wealth subcomponents, population groups, indices and exchange rates. This project is ongoing for the remainder of the series but in total this covers 99.5% of the observations in the World Inequality Database.

In this technical note, we first lay out our upgraded grading framework. We then discuss the construction and use of the weighted average quality score, a tool that provides an overall picture of data quality for each country's series.

Info box: Where to find the data quality on WID.world

How to download the data quality?

The yearly data quality grades are downloaded by default as an additional column when downloading WID data, either via our user-friendly [direct download](#) on the website or via the statistical software packages in [R](#) or [STATA](#). There is no need to request the metadata variables, it downloads by default in WID.data.⁵

How do we summarize data quality over time?

To provide an overall picture of data quality for each country and series, we combine the yearly grades into a single score as a weighted average. The average data quality score is a weighted average of the data quality across years for a given country-series. More recent years receive greater weight, while the influence of older years gradually decreases. This ensures that the score reflects current data quality without disregarding valuable historical information. As such, it reflects both the richness of the underlying data sources and the consistency of coverage over time.

The average score can take non-integer values between 0 and 5, rounded to one decimal place. For a technical explanation of the score construction and the exact weights attributed in the time-decay weighting, please see [Section 8](#).

⁵ If you cannot see the data quality column when downloading via the R or STATA command, you may be using an older version of the WID package. Please uninstall and reinstall the package to access the latest version. The column is included automatically in new downloads and does not need to be requested separately.

2. Grades Descriptions

The data quality framework assigns grades from 0 to 5, with higher grades generally reflecting greater reliance on directly observed and higher-quality source data, and lower grades reflecting a greater degree of estimation or imputation. Because the underlying sources and estimation methods differ across WID concepts, the precise meaning of each grade is concept-specific. The tables below provide a summary of the grading scale for each concept. See the [Appendix](#) for the longer grade descriptions.

Short descriptions

Table 1. Distributed concepts

Grade	Distributed income	Distributed wealth	Female labour income share
0	Imputed regional average	Income inequality-based weighting	Imputed regional average
1	Carry forward or backward	Survey data	Predicted values (OLS) fully imputed
2	Interpolation	Wealth/estate tabulated tax data, probate microdata, simplified multiplier method	Predicted values (OLS) level-adjusted
3	Only one type of source: survey data or tax data.	Survey data scaled to national accounts wealth aggregates	Interpolation
4	Tax and survey data	Mixed model of income capitalization or tabulated wealth/estate tax data and survey data	Observed microdata – supplementary source
5	Tax and survey microdata	Wealth tax microdata or administrative wealth registers	Observed microdata – benchmark source

Table 2. Aggregate concepts⁶

Grade	Aggregate income	Aggregate wealth*	Indices & Exchange Rates	Population
0	Imputed regional average	Imputed regional average	Assumed static value	N/A
1	Carry forward/backward or simple extrapolation	Predicted and extended with regional trend	Imputed from related countries	Carry forward or backward
2	Predicted values (mainly OLS)	Observed and extended with regional trend	Imputed from domestic indicators	Imputed using domestic growth rates
3	Interpolation - Enforce	Predicted using macroeconomic covariates	Interpolation	Interpolation
4	Observed (historical SNA or academic sources)	Observed and extended with OLS prediction	Observed (historical SNA, academic sources, or forecasts)	Observed historical compilations
5	Observed (latest SNA)	Observed	Observed (latest SNA)	Observed (UN WPP or latest SNA)

*Aggregate wealth grade definitions are currently being revised and may be subject to change.

3. How to turn data quality grades in a binary indicator

For visualisation purposes, the 0–5 data quality grades are grouped into two categories: estimates primarily based on observed sources, and estimates for which such sources are only partially available or entirely unavailable. The first category is visualised on WID.world using solid lines, the second is shown using dashed lines.

There is an important distinction between aggregate and distributed series, reflected in the differences of the underlying data and the methodological challenges involved. Most major macroeconomic aggregates are reported by national statistical offices and can generally be treated as observed, even when limited in-house transformations are required to ensure consistency across variables and compliance with macroeconomic identities. Distributed data sources, by contrast, are typically only partially available -

⁶ The values calculated in-house are derived using accounting identities or assumptions based on already calculated variables. Their grade reflects the quality of the underlying components (see [section 6](#)).

either in terms of time or quality - and require more substantial adjustments, including combining with other supporting sources and recalibrating with macroeconomic aggregates. Their construction therefore involves a considerably greater degree of in-house intervention, which is detailed in the [DINA Guidelines](#) (Chancel et al., 2025). Although both aggregate and distributed series are visualised using solid/dashed segments on WID.world, users should consult the concept-specific grade definitions to understand the methodological choices associated with each estimate.

For distributed concepts, the grade depends on whether the estimate uses observed data for the year in question or observed data from nearby years. Pre-tax income series receive grades 3–5 when survey and/or tax data are available for the year in question from national statistical agencies. Grades 1–2 are assigned when no such data is available for the year itself, but there is observed data from closeby years that can be used to estimate the missing value. This situation is common because national statistical agencies often release micro-level surveys and tax data biennially or at longer intervals. Grades 1–2 therefore could still indicate that the series is grounded in neighboring observed country-level sources, although additional estimation is required to fill a temporal gap. Such estimates may remain entirely appropriate for many forms of analysis and should not automatically be disregarded. The lower grade reflects the need for temporal interpolation or extrapolation, rather than a complete absence of relevant data. Grade 0 is assigned when no observed country-level data are used for any year in the series and the entire series must be imputed as a last resort to provide coverage.

The cut-off grade used in WID.world for the solid/dashed line visualisation varies by concept, as shown below.

Concept	Grades classified as solid lines on WID.world
Distributed income	Grade ≥ 3
Distributed wealth	Grade ≥ 1
Female labour income share	Grade ≥ 4
Aggregate income	Grade ≥ 3
Indices & Exchange Rates	Grade ≥ 3
Population	Grade ≥ 3
Aggregate wealth	Grade ≥ 3

4. Special grading rules for post-tax income series

For the post-tax income series - both the post-tax national income series (*di inc*) and the post-tax disposable income series (*ca inc*)⁷ - we assign yearly data quality grades by inheriting the grades from the country's pre-tax income series (*pt inc*). So, for a given year, a country will have the same data quality grade in post-tax and pre-tax income series.

This is because the pre-tax income distribution serves as the foundation for constructing the post-tax income series (Fisher-Post & Gethin, 2023). Post-tax incomes are obtained by applying estimates of the distributional incidence of taxes and government transfers to the pre-tax series, while retaining all the underlying data inputs. The quality of the original pre-tax distribution is therefore largely preserved in the resulting post-tax series.

5. Special grading rules for fiscal income series

Fiscal income is a legacy concept from the World Top Income Database (WTID), the precursor of the WID. While it is still included in the WID, it has been progressively superseded by DINA income concepts, such as the pre-tax income series, which are better harmonized. Note that in most cases, fiscal income variables are only available for the top 10% of the distribution. For completeness, we use the distributed income grade definitions in Table 1 to develop data quality for the fiscal income series.

Fiscal income estimates are predominantly based on observed tax data and therefore generally receive grades between 3 and 5. Grade 3 is assigned when the series relies on tax tabulations only, as is the case for much of the historical fiscal income data. Grade 4 applies when tax tabulations are supplemented with survey microdata, while Grade 5 is assigned when sufficiently detailed tax microdata are used. As a result, grades 0–2 are rarely, if ever, applicable to the fiscal income series, since in the WID the fiscal series generally only contains years for which underlying source data is available.

It is important to note that these categories necessarily simplify substantial differences in the underlying sources and methodologies. In particular, in the case of the fiscal series, microdata is not inherently superior to tabulations for every purpose: the quality of a fiscal income estimate also depends on the population and income concepts

⁷ In the WID we only offer *ca inc* series for some countries in Europe.

captured by the source, including the coverage of different forms of capital income. Users requiring a more detailed assessment of a particular country or period are therefore encouraged to consult the original papers referenced in the WID publications, which document the construction of each fiscal income series.

6. Special grading rules for aggregate macroeconomic series

In addition to the rich distributional data, the WID includes comprehensive macroeconomic aggregate data that provides context for the distributions presented. These variables can be identified by the percentile p0p100 and the population–age group combination 999i or 992i. The construction of these aggregates differs across the income and wealth components. For income (and the associated indices) , aggregate values are primarily constructed through the selection of observations from different sources and their subsequent processing. For wealth, the estimation process also involves the calibration of models that allow for the statistical estimation of aggregates of assets, liabilities, and wealth.

In most cases, these processes are consistent with the definitions presented in [Table 4](#), and the assigned quality grade therefore directly reflects the quality scale defined there. However, some specific cases require additional rules to determine the data quality. These rules are necessary to preserve the purpose of the grading system: to provide users with an indication of the precision, origin, and processing of each value.

6.1. Aggregate income

While distributional income series result from the simultaneous processing and correction of several inputs, aggregate income data result primarily from the selection of values among the available sources. Consequently, in many cases, the data quality assigned to each observation depends on the quality attributed to the source reporting that value, whether this is a statistical institution, such as a national government or an international organisation, or a dataset constructed as part of a peer-reviewed publication, as described in the grades 4 and 5 of the [Table 4](#).

Nevertheless, WID applies additional calculations to these raw series, including the application of single or multiple accounting identities, as well as different types of imputations (see Chancel, Flores, Moshrif, Nievas, & Piketty (2025) for a description of these data treatments). The quality grade attributed to these imputations can be

classified within the grades 0-3 presented in [Table 4](#). However, some cases have specific characteristics that require additional rules for setting the data quality, which are described below.

6.1.1. In-house calculations

In-house calculations typically involve obtaining the value of an aggregate concept ($K_{c,t}$) for country (c) and year (t) either as a residual from an accounting equation or through an imputation based on assumptions found in the literature, definitions, or indexation to other variables. The different forms that these imputations can take, together with the corresponding data-quality rules, are summarised below. The variables involved in the calculations can be cross-checked in the metadata for each observation-country, in the column “source” of the WID or by clicking in the “Info” button and checking the “TECHNICAL NOTES” in the WID data visualization tool in the wid.world website.

Addition of components: Some variables are constructed through the simple addition of other variables. In this case, the quality grade of the component with the largest value is assigned to the resulting aggregate. The nature of these components can be consulted in the variable-country metadata.

For country c , year t , and aggregate concept $K_{c,t}$ resulting from the sum of aggregate concepts A^i , the grade Q is:

$$Q_{c,t,K} = Q_{c,t,A}, \text{ if } |A^i_{c,t}| > |A^j_{c,t}| \quad \text{for all } j \text{ in } \{0, \dots, n\}, j \neq i$$

Ratio adjustment: For $K_{c,t}$ variables constructed as

$$K_{c,t} = A_{c,t} * (D_{c,t} / E_{c,t})$$

The ratio ($D_{c,t} / E_{c,t}$) is interpreted as an adjustment factor, while $A_{c,t}$ is considered the core value. The quality grade of the resulting variable therefore corresponds to the quality grade of the core value:

$$Q_{c,t,K} = Q_{c,t,A}$$

Addition of an adjusted component: For $K_{c,t}$ variables constructed as

$$K_{c,t} = A_{c,t} + B * (D_{c,t} / E_{c,t})$$

the ratio $D_{c,t} / E_{c,t}$ is again considered an adjustment factor. In order to maintain consistency with the previous rule and to keep the calculation simple, the quality grade is assigned according to the heavier of the terms:

$$Q_{c,t,K} = Q_{c,t,A}, \text{ if } A_{c,t} > B_{c,t}$$

Similarly, the corresponding grade for the adjusted component is retained when $B_{c,t} > A_{c,t}$.

6.1.2. Macroeconomic identities enforcement

Given the diversity of sources and estimation methods used to construct macroeconomic variables, the final stage of the processing involves an enforcement algorithm that solves a system of simultaneous macroeconomic identities. The algorithm modifies existing values by the minimum amount necessary and fills missing values with estimates that allow the system of equations to be satisfied.

When the algorithm modifies an existing reported value, the modification is typically small. Consequently, the data-quality grade attributed to the original reporting source remains unchanged. When a variable is not reported by any of the considered sources and the enforcement algorithm estimates a value that is consistent with the system of equations, the resulting value is assigned a quality grade of 3, as indicated in [Table 4](#).

6.1.3. Extensions to averages, macroeconomic totals, and wealth-to-income ratios

WID estimates of macroeconomic aggregates are internally handled as shares of GDP (WID codes starting with *y*). We also provide versions of these variables in per-capita terms (*a*), as macroeconomic totals in monetary terms (*m*), and as shares of net national income (*w*). These versions are obtained through simple transformations of the *y* versions of the variables: by multiplying by the GDP estimate to obtain monetary aggregates, and later by dividing by the population to obtain per-capita values, or dividing by net national income to obtain shares of net national income. For the resulting series, the assigned data-quality grade corresponds to the grade originally attributed to the corresponding variable in its *y* version. This means that these changes in the unit of a variable do not affect its quality grading, despite the fact that the parameters used to change the unit have their own quality gradings

Although the rules described in [Section 5.1.1](#) could suggest treating GDP as the "heavier" variable and the variable expressed as a share of GDP as an adjustment factor, applying this approach would systematically assign the quality grade of GDP to the resulting aggregate. This would lead to a homogeneous quality grade across the macroeconomic section of the WID and would prevent the quality indicator from signalling the imputations or assumptions underlying the construction of the original variable. In addition, inheriting the data-quality grade from the original y version preserves internal consistency across the different representations of the same underlying variable. Since the a , m , and w versions are simply transformations of the y series, they retain the information on data quality associated with the original estimation and its underlying assumptions.

6.2. Aggregate wealth

Wealth aggregates are constructed using different estimation procedures (see Bauluz, Brassac, Martínez-Toledano, Piketty & Sodano (2026), for details). Consequently, their grading rules differ in some respects from those applied to income aggregates. Nevertheless, there are two important similarities. First, wealth aggregates are originally estimated as shares of GDP (variables starting with y) and later extended to w , m and a . Second, some variables are constructed through in-house calculations. Consequently, the data-quality grading for these variables can follow the same algorithms defined in the *In-house calculations* and *Extensions to a , m , and w* sections above.

6.3. Indices and exchange rates

The indices include the price index (whose WID code is `inyixx`), exchange rates at market exchange rates (MER) (`xlc__x`) and purchasing power parity exchange rates (PPP) (`xlc__p`), real exchange rates (`xrer__`), and time-varying currency units (`inctlu`).

The price indices are constructed from deflators and inflation rates collected directly from international organisations, while the time-varying currency units are constructed from historical information collected for review papers (in this case, Nievas & Piketty (2025), is the source used by WID). The other variables require additional in-house calculations, which in turn require specific rules for assigning data-quality grades.

6.3.1. Market exchange rates

The exchange rate to USD (`xlcusx`) is constructed by collecting information from several institutions, mainly the United Nations and international financial institutions for the

period since 1970, and from Nievas & Piketty (2025) for historical years. In this case, the data-quality grade corresponds to the quality assigned to each underlying source.

In contrast to aggregate quantities, where one component may account for most of the level or trend of a composite variable, the construction of indices and exchange rates often combines inputs that are equally important. In these cases, the quality of the resulting variable cannot exceed the quality of its least reliable input. The corresponding rule therefore relies on a minimum function. For example, the exchange rates to EUR and CNY for all countries are constructed by triangulating the exchange rate to USD with the exchange rate of the euro-area countries and the exchange rate of China, respectively.

In the case of the Euro, however, this triangulation requires an additional calculation. The exchange rate of the euro to USD before 1998 is constructed as an average of the exchange rates of France, Germany, Italy, the Netherlands and Spain weighted by its GDP (see Chancel, Flores, Moshrif, Nievas, & Piketty (2025)). Consequently, the quality grade of $Q_{euro,t,k}$ corresponds to a weighted average of the quality grades assigned to the estimated exchange rates k from these countries, also weighted by the GDP for holding consistency, this is:

$$Q_{euro,t,k} = \frac{\sum_c^{C_{euro,t,k}} GDP_{c,t} q_{c,t,k}}{\sum_c GDP_{c,t}} \text{ for all } c \text{ in } \{DE, ES, FR, IT, NL\}$$

These historical estimates are typically assigned a grade of 4, given that series for the main countries are included in Nievas & Piketty (2025), the main source for historical aggregates.

As so, for the exchange rate to EUR (xlceux), the assigned data-quality grade for each country is therefore the minimum of:

- the quality grade of the country's exchange rate to USD; and
- the weighted-average quality grade of the exchange rates of the euro-area countries.

Similarly, for the exchange rate to CNY (xlcyux), the assigned grade is the minimum of:

- the quality grade of the country's exchange rate to USD; and
- the quality grade of China's exchange rate to USD.

6.3.2. Purchasing Power Parity exchange rates

The PPP exchange rate to USD (xlcupsp) is based on an observed benchmark year⁸ and is extended across years by comparing the price index of the country with that of the United States (see Chancel, Flores, Moshrif, Nievas, & Piketty (2025) for a description of these data treatments). For the reference year (normally 2021), the quality grade is directly derived from the quality associated with the underlying data sources or imputation methods. For the years obtained through the extension procedure, the quality grade is assigned as the minimum of the quality grades of the reference year's PPP value, the country's price index and the US price index for the corresponding year.

The exchange rates to EUR and CNY are constructed using a process similar to that described above for MER, through a triangulation. For the Euro, this creates the necessity for the weighted-average PPP exchange rate is estimated for France, Germany, Italy, the Netherlands and Spain.

With these elements is possible PPP exchange rate to EUR (xlceup) obtained through triangulation, with a quality grade equal to the minimum of the quality grade of the exchange rate to USD and the weighted-average quality grade of the euro-area countries. And, symmetrically, the quality grade of the PPP exchange rate to CNY (xlcyup) is the minimum of the quality grade of the PPP exchange rate to USD and the quality grade assigned to China's exchange rate.

6.3.3. Real exchange rates

The real exchange rates (xrerus, xrereu and xreryu) are calculated from the relationship between the MER exchange rate and the PPP exchange rate. Consequently, its data-quality grade is assigned as the minimum of the quality grades of the two underlying exchange rates.

⁸ Worldwide estimates of purchasing power parities (PPPs) are not available on an annual basis. At WID, we rely primarily on data from the World Bank's International Comparison Program, whose latest comparison round was conducted in 2021. This is the reference year for most of the countries.

7. Special grading rules when a variable has both distributed and aggregate observations

While most WID distributions are calculated for equal-split adults (992j) or the equal-split total population (999j) within each country, in some cases we also provide distributional series calculated for the overall individuals population, identified by the 999i population-age group combination, or adults individuals (992i). As mentioned in [Section 6](#), these population-age group combinations also define the standard expression of aggregated income and wealth concepts in the WID. What follows is a specific warning for a minor case of observations.

In specific country-year cases for fiscal income (*fiinc*) and net personal wealth (*hweal*), the quality grade assigned to observations with percentile *p0p100* and population–age group combinations 999i and 992i retains the grade estimated for the corresponding aggregate concept (see [Table 4](#)). For all other cases (that is, for other percentiles of the same variables and population-age group combinations different from *p0p100*), the quality grade corresponds to the distributional definitions provided in [Table 3](#). This special grading rule only applies to 1.2% of the country-year combinations in the WID for which we provide wealth series ending in 992i or 999i.

8. Special grading rules for regional series

Regional series in WID.world are constructed by combining the values of their constituent countries. The combination method depends on the kind of variable: for most macroeconomic aggregates it is a sum, while for distributions it is normally a population-weighted average. To describe the quality of the underlying data, we assign each region a yearly data quality grade calculated as a population-weighted average of the country-level grades for that year. This approach allows regional grades to reflect differences in data quality across countries and changes over time. For region *r*, year *t*, and distributional concept *k*, the grade *Q* is:

$$Q_{r,t,k} = \frac{\sum_c P_{c,t} q_{c,t,k}}{\sum_c P_{c,t}}$$

where $q_{c,t,k}$ is the country-level data quality grade, $P_{c,t}$ is the country's total population in that year, and $C_{r,t,k}$ denotes the countries contributing to the relevant variable. Each country enters the calculation once per year and concept (and, in the case of the distributional series, regardless of the number of percentile observations).

We use total population (npopul999i) as the weighting basis for any variable, including, in the distributional case, both adult-population series (code 992) and total population series (code 999). Using a common population basis ensures that differences in countries' age structures do not themselves generate differences in data quality scores between the two series with the same concept. Where the underlying country population share and grades are identical, the resulting regional scores are therefore identical across age groups.

As it is a weighted average, the regional grades will not necessarily take exact integer values; to ensure consistency with the rest of the yearly grades in WID.world, we round the yearly region grades to the nearest integer within the range 0 to 5. This methodology ensures that both regions expressed in market exchange rates (MER) and purchasing power parity (PPP) have the same data quality.

8.1. Distributional regional series

For each type of distributed concept, the regional series is assigned the population-weighted methodology. As the grades and definitions vary by concept, the regional grades vary too. The post-tax income regional series are assigned the same grades calculated as the pre-tax income regional series, consistent with the country-level treatment described in [Section 4](#). Regional wealth data quality grades are calculated using the country-level wealth grades.

The regions' data quality grades summarize the quality of the underlying country data; it does not measure the precision of individual regional percentile estimates. Changes over time reflect changes in country grades, relative population weights, and, where applicable, country coverage.

8.2. Macroeconomic regional aggregates

In the case of income and wealth aggregates, population weighting allows us to grade the estimated values by capturing the quality of the largest population countries/sub-regions within the region. In the case of population estimates, it reflects

the fact that, by construction, total population is the leading variable, and the other age disaggregations are scaled to this value.

In the case of indices, one additional clarification is needed. Since regions are theoretical constructs, regional indices values are calculated implicitly from the comparison of aggregations of national net national income values expressed in different currencies and adjusted for exchange rates and price differences (see Chancel, Flores, Moshrif, Nievas, & Piketty (2025), for details on these calculations). While the value of the index is calculated implicitly, its quality grade is calculated explicitly, following the same approach as for the aggregates. More specifically, the quality grade is calculated as a population-weighted average of the quality grades of the indexes of the constituent countries.

9. Summarizing data quality as a weighted average

To provide an overall picture of data quality for each country and series, we combine the yearly grades into a single time-invariant score (“average quality score”). The score is a weighted average of the data quality across years. It gives more weight to recent years, reflecting both the richness of the underlying data sources and the consistency of coverage over time.

9.1. Time-decay weighting of data quality series

To account for the relative importance of recent observations, we assign each period t in the series a *time-decay weight* based on its distance from the most recent observation T . The weight is given by a logistic (S-shaped) function:

$$w_t = \frac{1}{1 + \exp\left(-\frac{d_t - c}{k}\right)} \quad \text{where } d_t = T - t$$


where:


- d_t is the *temporal distance* in years between observation t and the most recent year T ;
- c is the *midpoint parameter*, determining the distance at which the weight equals 0.5;
- k is the *slope parameter*, controlling how rapidly weights decay around the midpoint.

These weights are unnormalized, meaning they do not sum to one. Instead, they can be interpreted as importance scores that measure how much each observation contributes to an overall assessment of data quality. Recent data receive weights close to 1, while older data receive weights approaching 0, with the transition between the two regions controlled by c and k . A larger value of c shifts the midpoint farther into the past, assigning higher relative importance to older observations. A larger value of k produces a gentler decay, spreading influence more evenly across time. This logistic specification captures a smooth, flexible transition from high to low importance, rather than an abrupt cutoff. This reflects that information from past years may remain partially relevant but should gradually lose weight as the data become outdated, also avoiding penalizing countries with valuable longer historical series where data quality fades out.

9.2. Simulation of logistic weighting functions

The following simulations illustrate how these weights vary across time for different parameter values. The first example fixes $c = 30$ (midpoint 30 years from the most recent observation) and varies the slope k . The second fixes $k = 8$ (moderate slope) and varies c .





In both cases, the curves illustrate how the logistic specification produces a smooth transition from high to low importance. Smaller values of k yield a sharper decay around the midpoint, while larger values of k flatten the transition. Similarly, increasing c moves the midpoint further into the past, extending the influence of older observations.

9.3. Parameter values used for average quality score calculation

The exact values chosen for the parameters to construct the weighted average score are arbitrarily defined as $c = 20$ and $k = 5$. We also exclude the latest year from the calculation because it is generally extrapolated, since the sources used to calculate the data for the latest year are not publicized immediately.

9.4. Special note for population series

The weighted average quality score is observed at the country-concept level as it is a metadata variable in the WID. This means that, for a given concept, the average score is constant across averages, shares, thresholds, population groups and ages. The population series are a specific exception to the norm where the yearly data quality of the various population series (eg. `npopul999i`, `npopul992i`, `npopul996i`, etc) are not constant at the concept level (`popul`) because the sources and methods for the various age groups are different, and hence the yearly data quality grades may be different too.

For this special case, when it comes to constructing the average quality score, we use the yearly grades of the total population series (`npopul999i`) to arrive at an average score for all series with the concept `popul`. We chose this age group because it

represents the total population, i.e. individuals of all ages, which we believe is the most appropriate choice for providing a general picture of quality for our population series.

References

- Bauluz, L., Brassac, P., Martínez-Toledano, C., Piketty, T., Sodano, A.(2026). *Estimation of Global Wealth Aggregates in WID.world: Methodology*. WIL Technical Note 2026/03
- Chancel, L., Flores, I., Moshrif, R., Nievas, G., & Piketty, T. (2025). *Distributional National Accounts (DINA) Guidelines 2025: Methods and concepts used in the World Inequality Database*
- Fisher-Post, M., & Gethin, A. (2023). *Government redistribution and development, global estimates of tax-and-transfer progressivity, 1980-2019*. WIL Working Paper 2023/17
- Gomez-Carrera R., Moshrif R., Nievas G., Piketty, T., Somanchi A. (2024). *Extending WID Population Series: Projections 2024-2100 & Age/Gender Breakdowns*. WIL Technical Note 2024/12
- Nievas, G., Piketty, T.,(2025) *Unequal Exchange and North-South Relations: Evidence from Global Trade Flows and the World Balance of Payments 1800-2025*. WIL Working Paper 2025

Appendix

Long descriptions

Table 3. Distributed concepts

Grade	Distributed income	Distributed wealth	Female labour income share
0	No data is available (imputed regional average). This is the lowest-quality category and is used only as a last resort to ensure coverage.	Estimates are based on the observed correlation between income inequality and wealth inequality. Wealth distribution is predicted using an income inequality weighted average of other countries' wealth distributions, with adjustments for rich-list and national accounts.	No income and employment data is available (imputed regional average). This is the lowest-quality category and is used only as a last resort to ensure coverage.
1	The value is imputed by carrying the nearest available observation forward or backward, filling short gaps in an otherwise continuous series	Estimates based solely on survey data, statistically adjusted to approximate missing top-end values. They include rich list and national accounts corrections when available (typically from the 1990s-onwards).	Values are estimated using an OLS regression model with female shares of wage and self-employment as the main predictors, employment data are drawn from ILO modelled estimates. Unlike grade 2, no observed income data are available for the country, so the full series is model-generated.

2	Values are estimated by linear interpolation between two known data points, typically for short, internal gaps.	Wealth distribution estimates are imputed either via a simplified mortality-based multiplier method or from micro inheritance/probate data or from wealth tax tabulations. These are supplemented by rich list and / or national accounts adjustments when available (typically from the 1990s-onwards).	Values are estimated using an OLS regression model with female shares of wage and self-employment as the main predictors, employment data are drawn from ILO modelled estimates. Where observed income data exist for the country, the predicted trend is level-adjusted to match those observations, anchoring the series to real data points.
3	Only one type of source is used: either income/consumption surveys, or tax tabulations; the two sources are not jointly available. Tax tabulations are typically the sole source for the historical series (pre 1980) or fiscal income series.	Aggregate wealth figures are derived from national-accounts balancing items, then distributed according to survey shares.	Values are estimated by linear interpolation between two known data points, typically for short, internal gaps.
4	Tax and survey data are both available, but at least one source is in tabulated (not micro) form.	Wealth distribution is modeled using a hybrid of either income-capitalization techniques or wealth/estate tax tabulations and survey information, with rich list/national accounts adjustments.	Female and male wage and self-employment incomes are derived from a supplementary survey source. These observations are used to extend the time series beyond the years covered by the benchmark source, and are therefore of slightly lower priority but remain grounded in observed microdata.

5	Fully integrated micro-level data from tax records and household surveys are available.	High-quality wealth microdata from administrative registers or tax declarations, requiring only minor imputations for unobserved assets.	Female and male wage and self-employment incomes are directly observed from the benchmark survey source for each country. The benchmark source is selected based on breadth of coverage, methodological consistency, and systematic data collection procedures. These observations represent the highest data quality in the series.
---	---	--	--

Table 4. Aggregate concepts

Grade	Aggregate income	Aggregate wealth*	Indices & Exchange Rates	Population
0	No data are available. Values are assigned using a regional average or another assumed value solely to ensure coverage. This is the lowest-quality category and is used only as a last resort.	No data are available. Values are assigned using a regional average to ensure coverage. This is the lowest-quality category and is used only as a last resort.	No data are available; a static value is assumed. This is the lowest-quality category and is used only as a last resort to ensure coverage.	No observation is classified with this score.
1	Values are estimated by extrapolating nearby observations using simple assumptions, such as forward/backward projection, moving averages, regional growth rates, or stable relationships with a related country or former common jurisdiction.	Values are estimated by extrapolating nearby OLS predictions with regional average.	The value is imputed either directly or by projecting an observation using a secondary variable -or the same variable- as a reference for growth observed in one or more neighbouring countries (or countries with historical affiliations) (e.g., imputation of exchange rates based on comparisons of Yugoslav GDP values expressed in different currency terms, with the resulting estimates applied to the former Yugoslav countries).	The value is imputed by projecting the nearest available observation forward or backward, assuming constant relative shares, estimated in a recent reference period, between the country and the aggregate of another country with which it was part of the same jurisdiction at some point in time (e.g. the Czech Republic and Slovakia following the dissolution of Czechoslovakia).

2	Values are estimated using statistical or econometric methods, such as OLS regressions on relevant macroeconomic covariates or signal-extraction techniques (e.g., Hodrick–Prescott filtering).	Values are estimated by extending nearby directly observed data with regional average.	The value is imputed by projecting an observation using a secondary variable observed within the same country as a reference (e.g., KP PPP exchange rate calculated based on the comparison of GDP values expressed in different currency terms), or by carrying the nearest available observation forward or backward.	The value is imputed by projecting the nearest available observation forward or backward, using growth rates or age shares observed in the country's available totals.
3	Values are estimated by linear interpolation between observed data points or derived through the simultaneous adjustment of macroeconomic identities to ensure internal consistency.	Values are estimated using OLS regressions on relevant macroeconomic country-specific covariates rather than past observations.	Values are estimated by linear interpolation between two known data points, filling short internal gaps.	Values are estimated using linear interpolation between adjacent data points to fill short internal gaps.
4	Values are directly sourced from older national accounting frameworks, benchmark academic studies, or historical compilations. Comparability with current standards may be limited.	Values are estimated by extending nearby directly observed data with OLS predictions.	Values are sourced from benchmark academic research, historical compilations, forecasts under the latest SNA (e.g., IMF WEO estimates of deflators for some countries in the past year), or from older SNA frameworks.	Values are directly sourced from historical compilations; comparability with current standards may be limited.

5	Values are directly observed from the most recent national accounts compiled under the latest SNA standards. This is the highest-quality category.	Values are directly observed in official balance sheets (from the most recent national accounts compiled under the latest SNA standards), global asset datasets, or validated academic studies. This is the highest-quality category.	Values are directly observed from the most recent national accounts under the latest SNA, representing the highest-quality data.	Values are directly observed from the UN World Population Prospects or from the most recent national accounts compiled under the latest SNA, representing the highest-quality data.
---	--	---	--	---

Note: The values calculated in-house are derived using accounting identities or assumptions based on already calculated variables. Their grade reflects the quality of the underlying components (see [section 6](#)).

*Aggregate wealth grade definitions are currently being revised and may be subject to change.