Journal of Income Distribution
Volume 32, number 3-4, 2023, pages 35-39

Notes on the Gini Coefficient

Rajiv Sethi*
Barnard College, Columbia University and the Santa Fe Institute

The Gini coefficient has been introduced to generations of students using some variant of the formula:

$$G = \frac{1}{2\mu n^2} \sum_{i=1}^{n} \sum_{j=1}^{n} |y_i - y_j| \, , \tag{1}$$

where $y_1, \ldots, y_n$ are non-negative income levels in a population of size n and $\mu > 0$ is the mean income within this population (Sen 1973, Ray 1998). This index has an appealing interpretation as the average absolute difference between all pairs of individuals, relative to the mean income in the population. The pairs in this case include individuals paired with themselves, with the corresponding income differences being identically zero. The numerator is unaffected by their inclusion, but the denominator is inflated relative to the case when such pairs are excluded.

This reasoning has led some to favour an alternative version of the index that simply excludes self-matched pairs, and hence involves only $n(n-1)$ instead of n^2 comparisons (Jasso 1979, Deaton 1997, Bowles and Carlin 2020). The resulting measure of inequality is:

$$G' = \frac{1}{2\mu n(n-1)} \sum_{i=1}^{n} \sum_{j=1}^{n} |y_i - y_j| \, . \tag{2}$$

Clearly,

$$G = \frac{n-1}{n} G', \tag{3}$$

so the two measures are close for large n and converge in the limit. For small populations, however, the difference between the two can be substantial.

Both G and G' have been axiomatized by Thon (1982). Four principles together yield G, namely, the transfer principle (all order-preserving and equalizing transfers reduce inequality), population symmetry (pooling identical populations results in the same level of inequality as in the

* *Address for Correspondence:* Department of Economics, Barnard College, Columbia University, 3009 Broadway, New York, NY 10027, United States. E-mail: rs328@columbia.edu

component populations), constant population comparability (at any given population, the range spanned by the index should not depend on total income), and equidistance (all order-preserving and equalizing transfers between people at adjacent income levels have the same effect on the index).

To obtain an axiomatization for G', one can dispense with population symmetry and strengthen the comparability axiom to strong comparability, which requires that the range spanned by the index depends neither on total income nor on population. As Thon (1982, p. 140) puts it: “One might indeed want to postulate that the range of an inequality index is to be the same over the redistribution of any total income not only between a given number of people but between any number of people”.

One way to determine whether G or G' is preferred is to consider which of the two axioms – population symmetry or strong comparability – one is more willing to discard. Allison (1979) makes a case for G on the grounds that reasonable measures ought to satisfy population symmetry. That is, pooling two (or more) identical populations should result in the same level of inequality in the pooled population as existed in the component groups. If one person in a group of two has all the income, measured inequality should be the same as if all the income were shared equally by a thousand people in a group of two thousand. This is the case with G (which equals one-half in each case) but not with G' (which equals one in the former case and is close to one-half in the latter).

It could be argued, however, that the pooling of two or more identical populations could well result in a composite that is qualitatively different and may reasonably be held to have a very different level of inequality. In the example above, there is a group of people in the pooled population who share equality, and they must accommodate each other in social and political life. They may establish rights and responsibilities that apply only to themselves but nevertheless operate as a constraint on their behaviour and could eventually spread to society more broadly. Some will consider these arguments extraneous and irrelevant, but they may be persuasive to others.

By the same token, one can construct examples that seem to suggest that G' is a poor measure of inequality. It assigns the same value to a group in which one of two people has all the income, as it does to a group in which only one out of two thousand does ($G' = 1$ in each case, while G is one-half in the former and close to one in the latter). Worse, it assigns greater inequality to the former society (in which only one of two people has positive income) than it does to a society in which only two people out of a million have positive incomes. It could be argued that income is far more concentrated when the elite is small relative to the total population and that this ought to be reflected in the inequality measure.

There is another important argument to be considered, which pertains to the consistency of the measures with the partial order on arbitrary income distributions induced by the Lorenz criterion.

Any income distribution associated with a finite population can be represented by a set of points in two dimensions, with the cumulative share of the population on the horizontal axis (in order of increasing income) and the cumulative share of income on the vertical axis. A Lorenz curve is obtained by interpolating these points to obtain a non-decreasing and convex function on the unit interval.[1] If all the incomes are equal, there is only one interpolation that satisfies these criteria – the line of perfect equality or identity function. All other distributions yield curves that lie below this line, meeting it at the end points.

The process of interpolation associates with each finite population distribution an infinite population counterpart. This allows us to compare distributions, regardless of total income or population, by simply comparing their Lorenz curves. If a curve lies closer to the line of equality at all points, it corresponds to a distribution with lower inequality.[2] The result is a partial order on the set of all income distributions.

However, the particular partial order thus obtained clearly depends on the method of interpolation. Given two distributions, one method of interpolation may yield a clear ranking, while a second may involve intersecting curves. So, when one speaks of a Lorenz order, or consistency with the Lorenz criterion, there has to be a method of interpolation assumed either explicitly or implicitly.

The assumed interpolation is usually piecewise linear (Ray 1998). This always generates a curve that has the properties necessary for interpretation as an income distribution in an infinite population. In addition, it is the only interpolation method that respects the population symmetry axiom. That is, with piecewise linear interpolation, the merging of two identical populations results in a distribution which lies on the (interpolated) Lorenz curve corresponding to the component populations.

The standard Gini coefficient G is a completion of the partial order generated by piecewise linear interpolation. Specifically, G is the ratio of the area between the line of perfect equality and the Lorenz curve thus constructed, and the total area below the line of perfect equality. It is in this sense that G is Lorenz consistent.

There exist, however, several methods for nonlinear interpolation that can generate Lorenz curves with all the properties required (Gastwirth and Glauberman 1976, Cowell and Mehta 1982). These methods have been developed to deal with empirical applications involving binned data, but can also be applied to the case when data is available at an individual level for a finite population. A key step involves the fitting of an underlying density function to the available data points. Piecewise linear interpolation corresponds to a piecewise uniform density. Other densities map onto other Lorenz curves, including curves that are constructed to be continuously differentiable.

In his response to this note, Debraj Ray makes an important point that G' is inconsistent with the standard Lorenz ranking (based on piecewise linear interpolation). One can go further; there is no method of interpolation consistent with convexity and other required properties that generates a partial order with which G' is consistent. To illustrate this, consider any method of interpolation with the necessary properties. Corresponding to this, there will be some continuous Lorenz curve associated with the two-person distribution in which one person holds all the income. By choosing a population n to be sufficiently large and considering a distribution in which only two people in this population share all the income equally, one can obtain another Lorenz curve that lies strictly below the first curve for the same method of interpolation. This distribution will be treated as more unequal under the Lorenz criterion (based on the chosen interpolation), but under G', the former exhibits maximal inequality, while the latter does not.

Although G' is inconsistent with the Lorenz criterion (for any interpolation), it does have an interesting geometric interpretation. Suppose that one uses a step function for interpolation rather than a continuous convex function. In this case, the "line of perfect equality" is replaced by a step function, lying strictly below the conventional perfect equality line for finite populations, and is sensitive to population size, approaching the conventional line at the limit. In addition, there is a line of perfect inequality that lies on the horizontal axis, remaining independent of population size.

Unlike the conventional Lorenz curve, this step-function interpolation (being non-convex) cannot be interpreted as an income distribution in an infinite population. Nevertheless, one may ask whether the area between the step function corresponding to perfect equality and the step function corresponding to the observed income distribution can serve as a measure of inequality that ranks all distributions regardless of total income or population size. Indeed, it can, and the ratio of this area to the total area below the perfect equality step function is precisely equal to G'. To demonstrate this, one need only verify that this area measure satisfies strong comparability, since it clearly satisfies equidistance and the transfer principle. This is also evident – the ratio must be zero, when income is equally distributed, and one, when a single individual has all the income.

In fact, the equivalence of G' and this particular ratio of areas was recognised by Gini himself and is self-evident from Figure 1, which appears in Gini (1914).[3] This figure depicts G' exactly, based on Gini's own

Figure 1
Geometric interpretation of G' in Gini (1914).

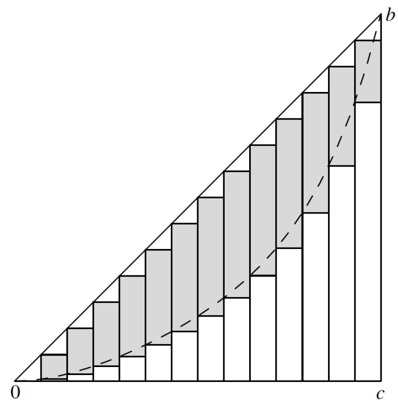

definition, for a population with $n = 14$ and a particular distribution of income. Notice that Gini uses a smooth, nonlinear interpolation (the dashed line) to construct a "Lorenz curve" for this finite population. This curve fails to satisfy the population symmetry axiom: merging two identical fourteen-person populations of this kind would result in points that do not lie on the curve as drawn.[4]

Finally, consider which of the two measures has a stronger claim to be known as the Gini coefficient. As it happens, both measures may be found in Gini (1912), although he appears to favour G' as a measure of income inequality (Ceriani and Verme 2012). Moreover, in Gini (1914), translated and cited as Gini (2005, p. 6), he is explicit about the requirement of strong comparability, stating that for any population size, his index of concentration "ranges from 1, in the case of perfect concentration, to 0, in the case of equidistribution".

Nevertheless, as observed by Allison (1979, p. 870), "both versions of the Gini index have found their way into the statistical literature, and neither one can be said to be incorrect". It is probably best if students are at least made aware of the existence and historical origins of both versions and presented with the arguments in favour of each. There is no uniquely correct Gini coefficient.

Notes

[1] Lorenz (1905) placed population shares on the vertical axis and income shares on the horizontal, resulting in concave curves.

[2] When comparing two distributions with the same population size, no interpolation is required.

[3] The figure shown here is taken from the translation (Gini 2005); I thank Sam Bowles for bringing it to my attention.

[4] To satisfy population symmetry, the curve would have to be piecewise linear, as noted above.