

Notes on the Gini Coefficient(s)*

Rajiv Sethi[†]

September 30, 2021

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

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

where y_1, \dots, y_n are non-negative income levels in a population of size n , and $\mu > 0$ is mean income in 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. But 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 favor an alternative version of the index that simply excludes such pairs, and hence involves just $n(n-1)$ rather than n^2 comparisons (Jasso, 1979; Deaton, 1997; Bowles and Carlin, 2020). The resulting measure of inequality is

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

Clearly

$$G = \frac{n-1}{n} G'$$

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

*These notes grew out of a very stimulating conversation with Debraj Ray, Julia Schwenkenberg, Sam Bowles, and Wendy Carlin.

[†]Department of Economics, Barnard College and the Santa Fe Institute.

Both G and G' have been axiomatized by Thon (1982). Four principles together yield G , namely equidistance (all order-preserving and equalizing transfers between people at adjacent income levels have the same effect on the index), the transfer principle (order-preserving transfers from rich to poor reduce inequality), population symmetry (pooling two identical populations results in the same inequality as in the component populations), and constant population comparability (at any given population, the range spanned by the index should not depend on total income).

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."

So one way to decide on 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 income was 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 the case with G' (which equals one in the former case and is close to one-half in the latter).

One could argue, however, that the pooling of two or more identical populations could well result in a composite that is qualitatively different, and can 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, who must accommodate each other in social and political life, and who may establish rights and responsibilities that apply only to themselves but nevertheless operate as a constraint on behavior and may 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 income as it does to a group in which one of a 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 just 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 is 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 (in order of increasing income) and the cumulative share of income on the vertical. A Lorenz curve is obtained by interpolating these points to get a non-decreasing and convex function on the unit interval.¹ If all 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 provide 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 either explicitly or implicitly assumed.

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 method of interpolation 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 piece-

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

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

wise 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.

However, there exist several methods for nonlinear interpolation that can generate Lorenz curves with all the required properties (Gastwirth and Glaubergerman, 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 we have data at the individual level for a finite population. A key step involves the fitting of an underlying density function to the available data points. The piecewise linear interpolation corresponds to a piecewise uniform density. Other densities map on to other Lorenz curves, including curves that are constructed to be continuously differentiable.

In his response to this note, Debraj Ray makes the 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 the other required properties that generates a partial order with which G' is consistent. To see 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 has all income. By choosing a population n to be sufficiently large, and considering a distribution in which just two people in this population share all income equally, one can get a Lorenz curve that lies strictly below the first one for the same method of interpolation. This will be treated as more unequal under the Lorenz criterion (based on the chosen interpolation). But under G' the former has 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 that lies strictly below the conventional perfect equality line for finite populations, and is sensitive to the size of the population, approaching the conventional line in the limit. In addition, there is a line of perfect *inequality* that lies on the horizontal axis, and is 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 be

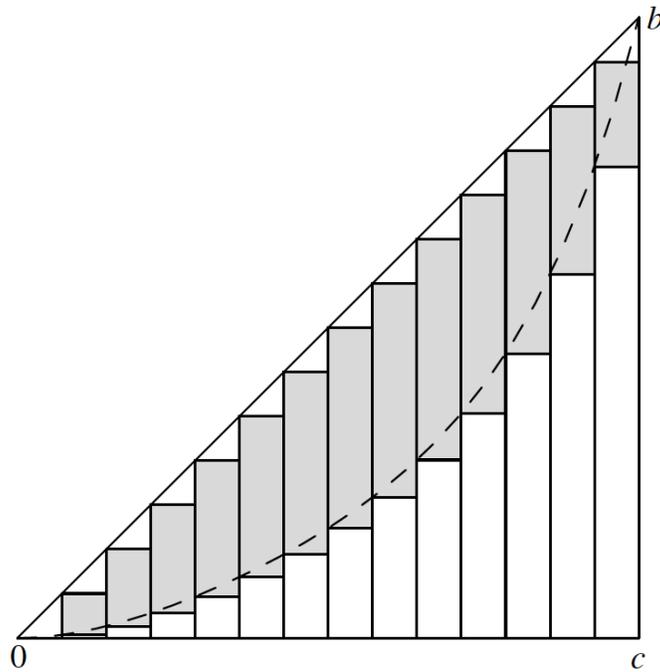


Figure 1: Geometric interpretation of G' in Gini (1914).

used 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 see this, one need only verify that this area measure satisfies strong comparability, since it clearly satisfies equidistance and the transfer principle. And this is also clearly true—the ratio must be zero when income is equally distributed, and must be one when a single individual has all income.

In fact, the equivalence of G' and this particular ratio of areas was recognized 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 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 (in order for it to satisfy population symmetry it would have to be piecewise linear, as noted above.)

Finally, consider which of the two measures has a stronger claim to be known as the

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

Gini coefficient. As it happens, both measures may be found in Gini (1912), although he appears to favor G' as a measure of income inequality (Ceriani and Verme, 2012). And in Gini (1914), translated and published as Gini (2005), 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 Allison (1979) has observed, "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, and presented with the arguments in favor of each. There is no uniquely correct Gini coefficient.

References

- Paul D Allison. Reply to Jasso. American Sociological Review, 44(5):870–872, 1979.
- Samuel Bowles and Wendy Carlin. Inequality as experienced difference: A reformulation of the Gini coefficient. Economics Letters, 186:108789, 2020.
- Lidia Ceriani and Paolo Verme. The origins of the Gini index: Extracts from *Variabilità e Mutabilità* (1912) by Corrado Gini. The Journal of Economic Inequality, 10(3):421–443, 2012.
- Frank A Cowell and Fatemeh Mehta. The estimation and interpolation of inequality measures. The Review of Economic Studies, 49(2):273–290, 1982.
- Angus Deaton. The Analysis of Household Surveys: A Microeconomic Approach to Development Policy. World Bank Publications, 1997.
- Joseph L Gastwirth and Marcia Glauber. The interpolation of the Lorenz curve and Gini index from grouped data. Econometrica, 44(3):479–483, 1976.
- Corrado Gini. Variabilità e Mutabilità: Contributo allo Studio delle Distribuzioni e delle Relazioni Statistiche. C. Cuppini, Bologna, 1912.
- Corrado Gini. Sulla misura della concentrazione e della variabilità dei caratteri. Atti del Reale Istituto Veneto di Scienze, Lettere ed Arti, 73:1203–1248, 1914.
- Corrado Gini. On the measurement of concentration and variability of characters. International Journal of Statistics, 63(1):1–38, 2005.
- Guillermina Jasso. On Gini’s mean difference and Gini’s index of concentration. American Sociological Review, 44(5):867–870, 1979.
- Max O Lorenz. Methods of measuring the concentration of wealth. Publications of the American statistical association, 9(70):209–219, 1905.
- Debraj Ray. Development Economics. Princeton University Press, 1998.
- Amartya Sen. On Economic Inequality. Oxford University Press, 1973.
- Dominique Thon. An axiomatization of the Gini coefficient. Mathematical Social Sciences, 2(2):131–143, 1982.