The Gini coefficient as a measure of experienced inequality

Samuel Bowles and Wendy Carlin[1]
2 October 2021

Thanks to Rajiv Sethi and Debraj Ray for your stimulating and clarifying memos (Ray, 2021, Sethi, 2021). Our paper (Bowles and Carlin, 2020) and the notes below present an alternative way of looking at the Gini coefficient(s).[2]

Before turning to the substantive issues we begin with three pieces of background– about how two years ago we came to work on this, about Gini and his coefficient, and about ethics and the measurement of inequality.

Background

Our CORE colleague, Antonio Cabrales, gave his students in the CORE intro course at UCL the conventional definition of the Gini coefficient, namely (here and throughout using the notation of our paper for consistency),

$$1)\quad G^L = \frac{\sum_{i=1}^{i=n}\sum_{j=1}^{j=n}\left|y_i - y_j\right|}{2n^2\underline{y}}$$

along with the usual summary (following Gini) that its value is 1 if one person has all the wealth and zero if wealth is equally distributed. The students could not solve the toy example problems with small n that they had been assigned.

After a certain amount of 'check your calculations' and a lot of head scratching we realized that of course they could not make the problems work because G^L is not 1 when one person (in a finite population) has all the wealth; for example, where n=2, G^L = 0.5.

Our head scratching included a return to Gini's original paper in which he defined what he called his "concentration ratio" as the sum of the absolute differences among the (unique non-identical) pairs, which we call Δ, or

$$2)\quad \Delta \equiv \sum_{i=j+1}^{i=n}\sum_{j=1}^{j=n-1}\left|y_i - y_j\right|$$

divided by the total number of such pairs, relative to mean wealth multiplied by one half or

---

[1] Bowles: Santa Fe Institute and CORE; Carlin: UCL, SFI, CEPR and CORE
[2] Thanks to Debraj for drawing (Thon, 1982) to our attention.

3) $$G = \frac{\Delta}{\left(\frac{n(n-1)}{2}\right)} \frac{1}{\underline{y}} \frac{1}{2} = \frac{\Delta}{n(n-1)} \frac{1}{\underline{y}}$$

from which we can see that it is the mean difference among all pairs in the population (the first term in the expression in the middle) divided by the mean value of *y*, giving us the "relative mean difference" times one half. A feature of this measure is, Gini pointed out, that it satisfies the condition that it varied from one ("maximum concentration") to zero ("minimum concentration") as also shown by (Deaton, 1997)

And he showed that this quantity is equal to the area between the (then newly invented) Lorenz curve and the perfect equality line divided by one-half in an infinite population. He provided the appropriate step functions for both the perfect equality "line" and the Lorenz curve for an example of a finite $n = 14$ population. Consistent with our paper, we designate the Lorenz curve-based spatial version of the Gini coefficient as G^L and Gini's version of the concentration ratio in his equation 11 (from his 1914 paper, which is also our equation 3 above) as *G*. As Debraj pointed out, Gini also provided alternative measures (and many more have been proposed since). In their notes, Debraj and Rajiv use G' to refer to our *G*, and *G*, to refer to our G^L.

Maybe the differences that we have aired in these memos stem from our differing perspectives on the relationship between one's ethics and the choice of an appropriate measure of inequality. Debraj writes "Every measure must be evaluated by the core ethical axioms they satisfy, and we then need to reach into our own ethical system to see which set of axioms fit the best."

Referring to the Lorenz curve Debraj writes: " Its foundation is laid by a fundamental set of axioms: population neutrality, income neutrality, as well as the transfers principle of Pigou and Dalton … as contributions to a *welfare economics* foundation for inequality measurement, these are the key axioms, and all further explorations begin from them—or should."

We agree that this is a way to choose among competing measures. But we have been motivated, instead, primarily by the desire to measure inequality as it is experienced by the members of a society, hence our term for *G*: *experienced inequality*. This could be the basis of a normative evaluation; surely one's ethical stance on some distribution of wealth cannot be indifferent to how it is experienced by members of the society. But it might also be an entirely descriptive measure for understanding such things as subjective wellbeing, stress, political attitudes, and the like.

Debraj writes "there are axioms and there are axioms." Our view of experienced equality suggests one to add to the list. **Axiom: Economic inequality is social**, that is, it is a relationship between or among people.

Experienced inequality.

We all agree that there is no single right way to measure inequality. Which measure one uses depends of course on the question for which the inequality measure is to provide an answer. This often turns on what it is about inequality that one wishes to capture. In our paper we cite the polarization index that Joan Esteban and Debraj developed as an example of a different measure of inequality developed for a specific purpose, to illuminate social conflict.(Esteban and Ray, 1994) We also cited one of Sam's works on the polygyny threshold, which shows that the Gini coefficient (either variant) fails to capture the expected relationship between wealth inequality and polygyny, suggesting the need for an alternative measure.

In our paper, we illustrated the Gini coefficient, that is, *G*, as a statistic describing inequality on a complete network, the edges of which are the differences in wealth between all pairs of the network nodes. We found it helpful both in teaching and in testing our own intuitions to let inequality be about the edges of the network rather than the nodes, that is, about pairwise differences in wealth, not how much wealth each individual has. An example of the two approaches is in Figure 1, with the inequalities counted in equation (3) on the left and equation (1) on the right.

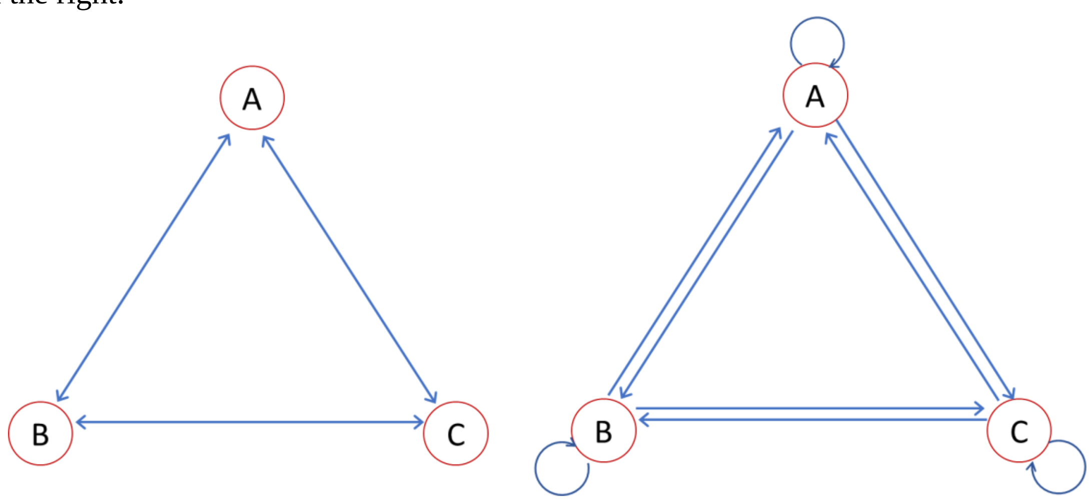

Figure 1. Experienced differences (left panel) and the edges used in the conventional Lorenz curve-based measure (right panel). If the nodes A, B, and C in Figure 1 have wealth 10, 4, and 3 the Gini coefficient given by equation (3) using the network representation in the left panel is 0.412. Using the network representation on the right (that is, equation (1)), however, the Gini is estimated as 0.274, which needs to be multiplied by $n/(n-1)$ or 1.5 to get the Gini coefficient restricted to the differences among actual pairs in the population (excluding the three "self-on-self" zero differences), as pointed out by (Yitzhaki and Schechtman, 2013)

We made the case that represented in this way, *G* captures important aspects of the way *inequality is experienced* by the members of a society in which people are aware of the wealth levels of everyone else. We also find the approach insightful as it allows us to measure experienced inequality in a society in which the relevant comparison set of individuals is not all others, but instead all others to whom one is connected in the network. For example, in a star network with a wealthy individual at the center, there is much more experienced inequality than in a complete network for the same set of nodes and levels of wealth at each node.

To sharpen our intuitions about inequality seen as pairwise differences among members of a population, suppose that every day, individuals are randomly paired to interact – economically, socially, in religious observance, and so on – with another member of the society. In the complete network representation, one of each individual's edges is selected at random. We are interested in the frequency over a great many such random pairings with which a member of the population interacts with a person of similar or different wealth, as we believe that the nature of these interactions will differ in important ways if the wealth differences are significant.

To explore looking at inequality this way, let's consider an economy with just two wealth levels which without loss of generality, we will set to zero and some positive number which is total wealth *y* divided equally among *r* rich members of the total population of *n*. Total wealth is some given level of mean wealth multiplied by the size of the economy or $y \equiv \underline{y}n$. Then the only unequal pairs in the population are the *r* wealth holders interacting with the *n-r* wealthless individuals with both members of the pair experiencing a wealth difference of *y/r*. So, we have $\Delta = r(n-r)\,\underline{y}n/r$ and equation (3) becomes

$$4) \quad G = \frac{\Delta}{n(n-1)}\frac{1}{\underline{y}} = \frac{r(n-r)\,\underline{y}n/r}{n(n-1)}\frac{1}{\underline{y}} = \frac{n-r}{n-1}$$

from which we confirm that irrespective of population size, $G = 1$ when one person has all the wealth and that inequality declines as wealth is redistributed so as to be shared among a larger number of rich, that is, increasing *r*.

We can also see that holding constant the number of the wealthy, *r*,

$$5) \quad \frac{dG}{dn} = \frac{1-G}{n-1} > 0 \ \text{ for } G < 1$$

So *G* increases as the propertyless class increases (holding constant the size of the wealthy class). In the limiting case where a single person owns all of the wealth, increasing the number of propertyless does not affect the *G*. It is this limiting case where $r = 1$ so $G = 1$ that Debraj labels the "absurd consequence" of measuring inequality by *G*.

Finally, to focus on the edges of the network rather than the nodes we calculate the fraction δ of all pairs in which the two have a different wealth level,

6) $\delta = \dfrac{r(n-r)}{n(n-1)/2} = \dfrac{2r}{n}G$

from which we see that if the population is divided into classes of equal size, then G *is* the fraction of interactions in which the members of the pair have different wealth (because 2*r/n* = 1).

Note also that if there is just one wealth owner and one wealthless person, all of the interactions are unequal $\delta = 2/n = 1$ (there is just a single interaction). But as the propertyless class increases in size the fraction of all interactions that are unequal falls.

How does it come about, then, that *G* increases with *n* (for a fixed number of rich households)? The answer is that as *n* rises (with *r* fixed) mean wealth remains unchanged (by assumption), but the mean wealth of the rich increases, and so the wealth difference between members of an unequal pair (that is $\underline{y}n/r$) also increases.

Is invariance to population replication a desideratum for a measure of inequality?

An attractive feature of the now-common variant of the Gini coefficient (equation 1, G^L) is (for populations of any size) equal to the area between the perfect equality line and the Lorenz curve divided by one-half . But as equation (1) makes clear, it includes the "equality" of one's own wealth with one's own wealth (what we term the "fictive zeros") in the measure of societal inequality, violating the "economic inequality is social" axiom. Because societal inequality as the term is routinely used concerns relationships between people and is meaningless in a "perfectly equal" one-person society (where $G^L = 0$), we find this aspect of G^L to be a reason not to use it where it differs appreciably from *G*.

The payoff to the inclusion of the self-on-self comparisons is said to be that, defined in this way, the measure conforms to another axiom ("population symmetry" in Rajiv's memo), namely, that the measure of inequality should be invariant to replication of its members, so that inequality would be the same (maximal concentration) in a two-person society in which one held all of the wealth and a four person society in which two people equally shared all of the wealth, or one thousand equally shared the wealth with another thousand wealthless, and so on.

Thus, that the counterintuitive fictive zeros are the price of invariance to population replication, an axiom that is thought to be sufficiently intuitively appealing to justify setting aside the problem of including the self-on-self comparisons in the inequality measure.

However, from the perspective of experienced inequality, invariance to population replication is a bug not a feature of equation (1). Let's think about three economies with 2, 4, and 6 people, in each of which half of the population owns all of the wealth in equal shares and as before, total wealth is proportional to population size. Is it sufficiently obvious that the three economies are "equally unequal" and that the level of equality is unaffected by population replication that violation of this axiom disqualifies an inequality measure?

In all three societies G is also equal to the fraction of all interactions that are between people of differing wealth levels (because $2r = n$, from equation (5), $\delta = G$). This falls from 100 percent when n=2 to two thirds for n= 4 and 3/5 for n= 6. The wealth difference between the members of the unequal pairs is unchanged as the population replicates (because n/r by design is unchanging). The only thing that has changed is that in the larger population people experience interactions with others of the same wealth a substantial fraction of the time. Our intuition is that G correctly shows that experienced inequality in the larger societies is less, and that as a result, the invariance to population replication axiom is unappealing or at least not so intuitive that its violation would disqualify G.

To understand what the invariance to population replication means think about the above example, but as the number of rich and poor increase from 2 each to 4 each, let the 2 poor and the 2 rich form single households, and similarly when the population grows to n = 6. Then because of the fictive formation of ever larger single households, one rich and the other poor, between household inequality would remain unchanged. The reason is that the additional within class pairings of equals – which leads G to fall as n and r both rise proportionally – would all be within these fictive households and hence would be ignored.

Let's now reconsider what Debraj finds "even more absurd" about G, namely that a thousand-person society with just two wealth holders is more equal than a 2-person society with one person holding all the wealth. We agree that these extreme cases challenge the intuitions. In the former the vast majority of interactions are among people the same wealth $(1-\delta = 1-2/1000 = 0.998)$ while in the latter none are. This does not settle which is "really" more unequal – where they exist, the wealth differences in the thousand-person society (500 times mean wealth rather than twice) are much greater than in the two-person society. But the example does suggest that it is far from absurd that the former would be perceived in some sense as more equal than the latter.

For more than one wealth holder, both G and G^L increase as the number of propertyless increases. For example, for a given r = 2, G is 0.5 for n = 3 and 0.95 for n = 20, while G^L is respectively 0.333 and 0.90. It would be helpful to clarify the intuition for inequality increasing with increasing numbers of the propertyless. Once clearly articulated, we could then assess whether the best way to capture that intuition is to count self-on-self comparisons as if they were real social relationships in which wealth inequality is entirely absent, or instead use a measure that captures the very same intuition but does not count the fictive zeros.

The above discussion also recommends that for some purposes a network representation of inequality if capturing the experience of inequality is an objective. We have motivated our examples by a random pairing environment equivalent to a complete network. But societies differ greatly in who interacts with (or even is aware of) whom. For example, if the social structure in question is a star with the wealth holder at the center, then the fraction of one's interactions that are with someone of a different wealth level (all of them, δ=1) is a constant as the size of the network grows, and as a result δ will be greater (for n>2) than in the complete network that we have used above as our illustration.

We are not saying that Debraj's intuition is incorrect (how could it be?) But it is far from obvious to us. And competing and quite different intuitions may also be appealing.

Bowles, Samuel and Wendy Carlin. 2020. "Inequality as Experienced Difference: A Reformulation of the Gini Coefficient." *Economics Letters*, 186, 1-3.

Deaton, Angus. 1997. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. Washington, D.C.: World Bank.

Esteban, Joan and Debraj Ray. 1994. "On the Measurement of Polarization." *Econometrica*, 62, 819-52.

Ray, Debraj. 2021. "Notes on "Notes on the Gini Coefficients(S)"." *unpublished manuscript*.

Sethi, Rajiv. 2021. "Notes on the Gini Coefficient(S)." *unpublished manuscript*.

Thon, Dominique. 1982. "An Axionatization of the Gini Coefficient." *Mathematical Social Sciences*, 2, 131-43.

Yitzhaki, S and E Schechtman. 2013. *The Gini Methodology: A Primer on a Statistical Methodology* New York: Springer.