(*) The Gini Coefficient

The *Lorenz curve*, $y = f(x)$, for a nation's economy describes the income (or wealth) distribution of that nation. For $0 \le x \le 1$, $y = f(x)$ gives the proportion of the national income earned by the lowest-earning $x \times 100\%$ of the population. For example, if the lowest-earning 10% of the population earn 1% of the national income, then $f(0.1) = 0.01$ and if the lowest-earning 50% of the population earn 22% of the national income, then $f(0.5) = 0.22$.

The Lorenz curve has the following characteristics.

- $f(0) = 0$ and $f(1) = 1$, because 0% of the population earns 0% of the income and 100% of the population earns 100% of the income.
- $f(x)$ is increasing because the bigger the proportion of the population, the more they earn.
- Lorenz curves are *concave up*, i.e., their derivatives (assuming that they are differentiable) are *increasing*.

Thus, the typical Lorenz curve looks like the (red) one in Figure 1 below.

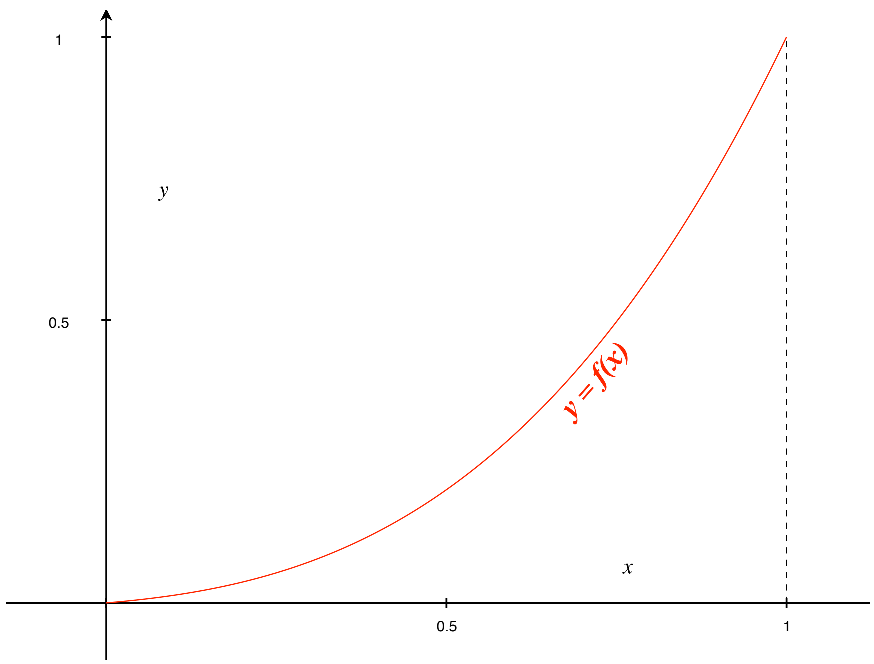

Figure 1: A typical Lorenz curve.

In the nation whose income distribution is described by the curve above, the lower-income earners earn less than the higher-income earners — i.e., there is *inequality* in the income distribution. In a (hypothetical) nation in which income is equally distributed among the entire population, the first 1% of the population earns 1% of the income, the next 1% of the population also earns 1% of the income, so the first 2% of the population earns 2% of the income, etc. For such a nation the Lorenz curve would be given by $y = x$, so that for each x, the lowest $x \times 100\%$ of the population earns $x \times 100\%$ of the income.

Comparing the Lorenz curve $y = f(x)$ of a nation to the curve of perfect equality $y = x$, we see that the more unequal the income distribution, the bigger the area of the region trapped between the curves $y = x$ and the Lorenz curve. This is depicted in Figure 2 below, where the Lorenz curve $y = g(x)$ describes a nation where income is more unequally distributed than the nation with Lorenz curve $y = f(x)$.

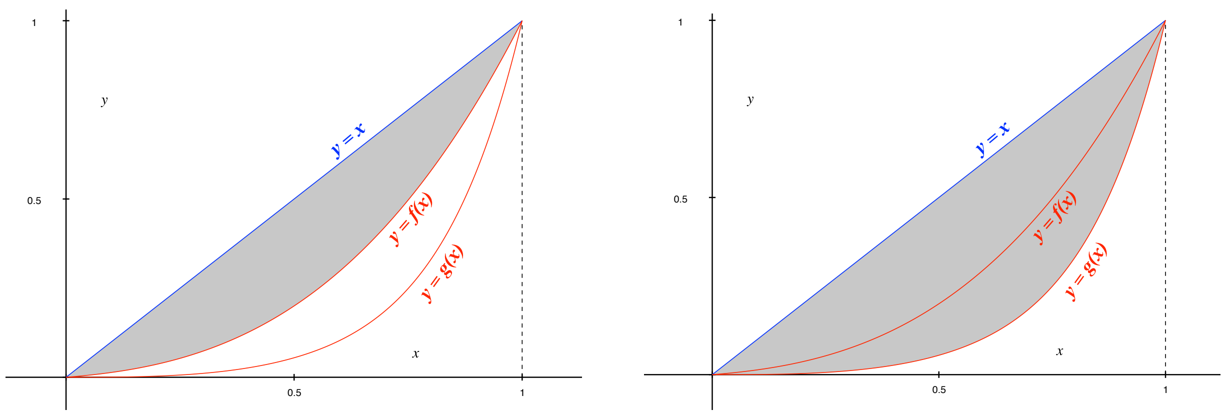

Figure 2: Two Lorenz curves compared to the curve of perfect equality.

The Italian sociologist Corrado Gini suggested measuring the inequality for a given nation by the ratio of two areas: (i) the area of the region between $y = x$ and $y = f(x)$ and (ii) the area of the triangle between $y = x$ and the interval $[0, 1]$ on the x-axis. This ratio is called the *Gini coefficient of inequality*, γ. I.e.,

area 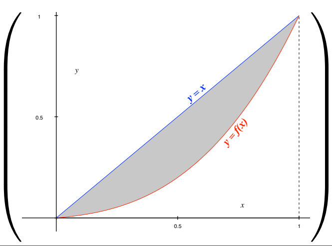

area 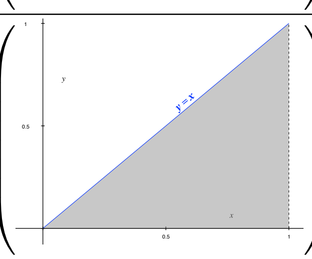

$$\gamma = \frac{\text{area}}{\text{area}} = \frac{\displaystyle\int_0^1 x - f(x)\,dx}{\frac{1}{2}} = 2\int_0^1 x - f(x)\,dx.$$

Example. Find the Gini coefficient of inequality for the nation with Lorenz curve given by $f(x) = 2^x - 1$.

$$\gamma = 2\int_0^1 x - (2^x - 1)\,dx = 2\int_0^1 x + 1 - 2^x\,dx = 2\left(\frac{x^2}{2} + x - \frac{2^x}{\ln 2}\bigg|_0^1\right) = 3 - \frac{2}{\ln 2} \approx 0.1146.$$

(*) **Average Value**

The average of n values, $a_1, a_2, \ldots, a_n$ is defined to be their sum divided by the number of values,[†]

$$Avg(a_1, \ldots, a_n) = \frac{1}{n}\sum_{j=1}^{n} a_j.$$

Suppose now that we want to find the average value of a function $f(x)$ on an interval $[a, b]$. The problem is that the set of values $\{f(x) : a \le x \le b\}$ whose average we want is an infinite set,[‡] so the definition above can't be applied directly in this case.

The solution is to use the original definition of average to find an approximate average of $f(x)$ on $[a, b]$, refine the approximation (i.e., make it more accurate somehow) and then declare the limit of the successive refinements (if it exists) to be the average value that we seek.

Step 1. To approximate the average value of $f(x)$ on $[a, b]$, we (i) choose a finite sample of points in the interval, $x_1, x_2, \ldots, x_n$, where

$$a = x_0 < x_1 < x_2 < \cdots < x_n = b,$$

and (ii) compute the average of the (finite) set $\{f(x_1), f(x_2), \ldots, f(x_n)\}$,

$$A_n = \frac{1}{n}\sum_{j=1}^{n} f(x_j).$$

If the points x_j are evenly spaced, i.e., if $x_j - x_{j-1} = (b-a)/n$, and n is large enough, the sense is that these n values do a good job of capturing the behavior of $f(x)$ in $[a, b]$ so that $Avg(f) \approx A_n$.[§]

Step 2. Refine the approximation by increasing the 'sample size'. I.e., make n bigger.

Step 3. Take a limit as $n \to \infty$. This leads to the following definition: *If the function $f(x)$ is continuous on the interval $[a, b]$ then the average value of $f(x)$ in the interval is given by*

$$Avg(f) = \lim_{n\to\infty}\left(\frac{1}{n}\sum_{j=1}^{n} f(x_j)\right).$$

Step 4. Observe that choosing the points $x_1, x_2, \ldots, x_n$ to be evenly spaced, as we did, means that these points divide the interval $[a, b]$ into n equally sized subintervals that have (common) width

$$\Delta x_j = \frac{b-a}{n}.$$

[†]Technically, this is the *arithmetic mean* of the n values. The word *average* can have many different interpretations, but this is the most common one.

[‡]A largish size of infinity at that.

[§]For this 'sense' to be justifiable, we assume that $f(x)$ is a *continuous function* in $[a, b]$.

Now, rewriting the definition of the average value, above we see that

$$\begin{aligned}
Avg(f) &= \lim_{n\to\infty}\left(\frac{1}{n}\sum_{j=1}^{n} f(x_j)\right) = \lim_{n\to\infty}\left(\sum_{j=1}^{n} f(x_j)\frac{1}{n}\right)\\
&= \frac{1}{b-a}\left[\lim_{n\to\infty}\left(\sum_{j=1}^{n} f(x_j)\frac{b-a}{n}\right)\right]\\
&= \frac{1}{b-a}\left[\lim_{n\to\infty}\left(\sum_{j=1}^{n} f(x_j)\Delta x_j\right)\right] = \frac{1}{b-a}\int_a^b f(x)\,dx,
\end{aligned}$$

since the limit in the square brackets in the third line is equal to the definite integral $\int_a^b f(x)\,dx$, *by definition*. So, to *compute* the average value of $f(x)$ on $[a, b]$, we compute the definite integral $\int_a^b f(x)\,dx$ (and divide by the length of the interval).

Example. Find the average value of the function $f(x) = \dfrac{x}{\sqrt{4x+9}}$ on the interval $[0, 4]$.

$$Avg(f) = \frac{1}{4-0}\int_0^4 \frac{x\,dx}{\sqrt{x^2+9}} = \ldots$$

$\ldots$ make the substitution $u = x^2 + 9$ and $du = 2x\,dx$, so $x\,dx = \dfrac{1}{2}\,du$ and the limits of integration change: $x = 0 \implies u = 9$ and $x = 4 \implies u = 25\ldots$

$$\ldots = \frac{1}{4}\int_9^{25} \frac{\frac{1}{2}\,du}{\sqrt{u}} = \frac{1}{8}\int_9^{25} u^{-1/2}\,du = \frac{1}{8}\cdot\left.\frac{u^{1/2}}{1/2}\right|_9^{25} = \frac{1}{4}(5-3) = \frac{1}{2}$$

You can think of the average value of $f(x)$ in $[a, b]$ as the average *height* of the graph $y = f(x)$ in the interval $[a, b]$. It is the height of the rectangle with base $[a, b]$ that has the same area as the area of the region bounded by $y = f(x)$, $y = 0$, $x = a$ and $x = b$, as illustrated below.

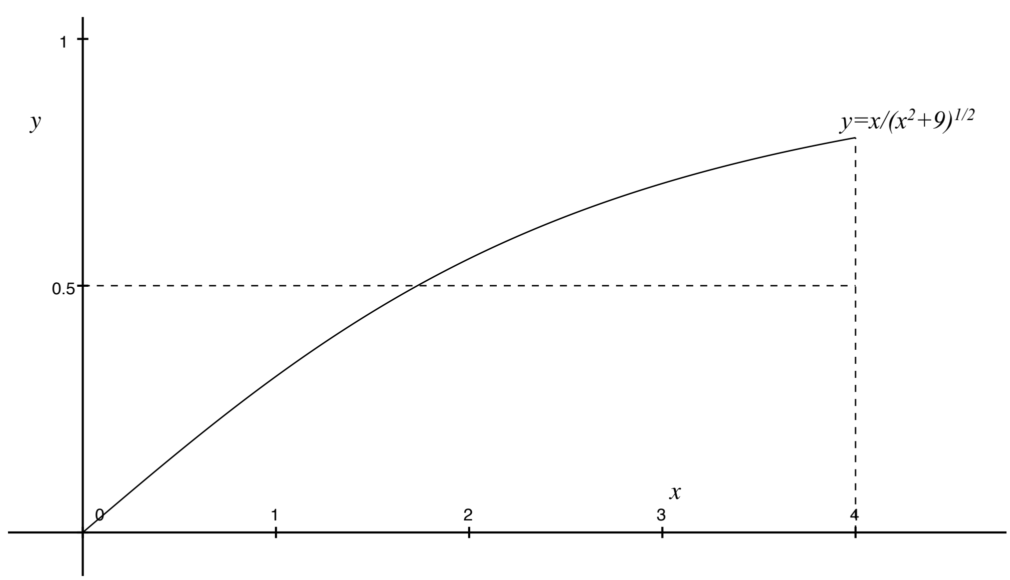

(*) **Consumers' and Producers' Surplus**

The market for an ordinary good can be (partially) described by the supply and demand curves, $p = g(q)$ (supply) and $p = f(q)$ (demand), for this good. Typically, the demand is higher than supply when the output is small, but as the output grows, the supply curve increases and the demand curve decreases, and they meet at the point of market equilibrium, (q^*, p^*). I.e., at the point of market equilibrium we have

$$f(q^*) = p^* = g(q^*)$$

and for $0 \leq q \leq q^*$ we have

$$(\text{demand}=)\ f(q) > p^* > g(g)\ (=\text{supply}).$$

We can divide the interval $[0, q^*]$ into subintervals,

$$[q_0, q_1],\ [q_1, q_2],\ \ldots, [q_{j-1}, q_j],\ \ldots, [q_{n-1}, q_n],$$

where $q_0 = 0$ and $q_n = q^*$, and think of them as market segments, $S_1, S_2, \ldots, S_n$. Each segment corresponds to a part of the total quantity q^* supplied and consumed by this market. We can think of the $(q_j - q_{j-1}) = \Delta q_j$ units of the good in segment S_j as being produced by suppliers in this segment and consumed by consumers in this segment.

Now, looking at the supply curve, we see that producers in this segment would sell these Δq_j units for a price of about $g(q_j)$, but since the price they are taking is p^*, this segment of producers is enjoying the *surplus* (of money)

$$PS_j \approx (p^* - g(q_j))\Delta q_j.$$

The total surplus for producers in this market is equal to the sum of the surpluses of the n segments,

$$PS = \sum_{j=1}^{n} PS_j \approx \sum_{j=1}^{n} (p^* - g(q_j))\Delta q_j.$$

Dividing the market into more and more (smaller and smaller) segments, yields (in principal) a more accurate approximation and we conclude that the producers' surplus for such a market is given by

$$PS = \lim_{n\to\infty} \sum_{j=1}^{n} (p^* - g(q_j))\Delta q_j = \int_0^{q^*} (p^* - g(q))\, dq,$$

where (to remind you), p^* is the equilibrium price for the market, q^* is the equilibrium quantity and $p = g(q)$ is the supply curve.

In the same way, looking at the demand curve, we see that consumers in this segment would be willing to pay about $f(q_j)$ per unit for the Δq_j units they are consuming, but they are only paying p^* per unit, so they are enjoying the surplus

$$CS_j \approx (f(q_j) - p^*)\Delta q_j.$$

As before, we see that the total surplus for consumers in this market is equal to the sum of the surpluses of the n segments,

$$CS = \sum_{j=1}^{n} CS_j \approx \sum_{j=1}^{n} (f(q_j) - p^*)\Delta q_j.$$

Dividing the market into more and more (smaller and smaller) segments, yields (in principal) a more accurate approximation and we conclude that the consumers' surplus for this market is given by

$$CS = \lim_{n\to\infty} \sum_{j=1}^{n} (f(q_j) - p^*)\Delta q_j = \int_0^{q^*} (f(q) - p^*)\, dq,$$

where in this case, $p = f(q)$ is the demand curve.

These formulas can be remembered using the figure below. In a sense, the consumer's surplus is equal to the area between the demand curve and the equilibrium price and the producers' surplus is equal to the area between the equilibrium price and the supply curve. This sense is correct as long as we use the correct *units* on the vertical and horizontal axes — (dollars/unit) on the p axis and (# of units) on the q axis — in which case height×width becomes (\$/unit) × (# units) = \$.