

Inequality Indices as Tests of Fairness

Ravi Kanbur

Dyson School and Economics Department, Cornell University

Andy Snell

School of Economics, University of Edinburgh

March 2017

Abstract

Standard income inequality indices can be interpreted as a measure of welfare loss entailed in departures from equality of outcomes, for egalitarian social welfare functions defined on the distribution of outcomes. But such a welfare interpretation has been criticized for a long time on the grounds that these indices are snap shot outcomes-based measures which do not take into account the process generating the observed distribution. Rather than focusing on outcomes, it is argued, we should be interested in whether the underlying process is “fair.” Following this line of argument, this paper develops statistical tests for fairness within well-defined income distribution generating processes and a well specified notion of fairness. We find that the likelihood ratio (LR) test for fairness versus unfairness within two such processes are proportional to Theil’s first and second inequality indices respectively. The LR values may either be used as a test statistic or to approximate a Bayes factor that measures the posterior probabilities of the fair version of the processes over that of the unfair. The answer to the process versus outcomes critique is thus not to stop calculating inequality measures, but to interpret their values differently—to compare them to critical values for a test of the null hypothesis of fairness, or to use them directly as a measure of the chance that the process was fair relative to the chance it was unfair. We also apply this perspective to measurement of “inequality of opportunity.”

Keywords:- Snapshot Inequality Indices, Process Versus Outcomes, Fair Versus Unfair Process, Likelihood Ratio Tests of Fairness, Bayes Factor

JEL Codes:- A10,A13,C01,C12,D63

1 Introduction

The standard procedure for measuring income inequality in a society is to take its snapshot observed distribution of income and to calculate an inequality index from it.¹ Such indices have also been interpreted as a measure of welfare loss entailed in departures from equality of outcomes, for egalitarian social welfare functions defined on the distribution of outcomes.² However, this procedure faces the well-known critique that the observed distribution is nothing but a snapshot outcome of a process, and that it is the process which matters for normative assessment. In particular, it is the “fairness” of the underlying process which is held to be the appropriate normative standard, not whether the observed inequality of outcomes is high or low.³ But if we take the process versus outcomes critique seriously, does this mean that we stop calculating total inequality, since it no longer has normative validity in and of itself?

We argue in this paper that even within the process frame overall indices of inequality still maintain their relevance, but now as statistical tests of fairness. In particular, we motivate the use of Theil’s two canonical indices of inequality as tests of a null of a "fair" income process versus an alternative of an unfair one in the same class. We find that the likelihood ratio tests for two distinct income processes are proportional to these two well used indices respectively. We then suggest that instead of presenting the indices as raw numbers one could present the p-values the raw numbers imply; a high Theil would imply a low p-value which in turn would indicate that the probability that incomes were generated by the relevant null fair process was low. We also apply this perspective to measurement of “inequality of opportunity.” In an extension we show that if the null and alternative are treated as equally likely prior "models" then the Theil indices are also proportional to the log of the relative posterior probabilities of unfairness to fairness under these two processes.

An early proponent of the process versus outcome line of argument was Milton Friedman (1962), who brought in the consequences of risk taking for interpreting observed inequality:

“Another kind of inequality arising through the operation of the market is also required, in a somewhat more subtle sense, to produce inequality of treatment. It can be illustrated most simply by a lottery. Consider a group of individuals who initially have equal endowments and who agree voluntarily to enter a lottery with very unequal prizes. The resultant inequality is surely required to permit the individuals in question to make most of their initial equality. Much of the inequality of income produced by payment in accordance with product reflects equalizing differences or the satisfaction of men’s taste for uncertainty.”

Friedman’s contribution simultaneously highlights the issue of process versus outcomes, and the fact that even when the process implies ex ante equality

¹For an overview and survey of standard methods and interpretation see Cowell(2011).

²The classic reference here is Atkinson(1970).

³Early critiques are by Sen(1979) and Dworkin(1981). A recent survey is by Roemer and Trannoy (2015).

(free lottery choice by identical individuals), the outcome may well show (misleadingly, in his view) inequality among individuals. Even though the process itself is fair, inherent randomness may show spurious inequality in outcomes.

Suppose we wish to evaluate the process in Friedman's example. This would require an evaluation of whether the lotteries faced by the different individuals were indeed identical. If we could directly observe the lottery choices, that would be the end of the matter. But this is usually not the case. All we can observe are in fact the outcomes. The task is then to try and infer from these outcomes the nature of the process which generated them. It is clear that in order to do this we will have to provide a minimal structure to the class of processes. It is only within a given class of processes that we will be able to infer more specific properties of the process which gave rise to the outcomes we observe. But we hope to show that making these assumptions can provide considerable insight into the relationship between outcomes and process.

Returning to Friedman's example, it is obvious that in reality endowments are not equal and individuals face different lotteries. The final outcomes are thus the result of these inequalities across lotteries, as well as the inequality caused by the fact that even when a group of individuals faces the same lottery, and so are equal ex ante, there will be winners and losers ex post within that group. That portion of observed inequality which can be attributed to the initial ex ante differences across lotteries might be referred to as "inequality of opportunity." We are then led to ask whether we can test for whether there are these ex ante differences ("unfairness") or not. This is clearly related to Roemer's (1998) famous formulation of attributing variation in outcomes to variation in "circumstance" (factors outside an individual's control) and "effort" (factors within an individual's control). Roemer's formulation takes the fraction of observed variation attributable to circumstance to be a measure of the "inequality of opportunity."⁴The question then arises as to whether this Roemer ratio can be a test statistic for the hypothesis of "fairness" within a specified income distribution process. In this paper we show that such an interpretation is indeed possible for a class of processes.

The plan of the paper is as follows. Section 2 sets out two income generating processes which will be the main focus of the paper, making clear what fairness or "equity" means within each. It then derives the Likelihood Ratio (LR) tests for fairness/equity under these income processes, and shows the links between the resulting test statistics and two canonical measures of inequality used in the literature, both due to Theil. Of particular interest is that the Mean Log Deviation (Theil's second measure), which is fast becoming the workhorse measure in inequality and inequality of opportunity analysis, can be interpreted as a test statistic for fairness. Section 3 extends the analysis to Roemer's conceptualization and specification of illegitimate and legitimate income variation. Section 4 moves from a classical to a Bayesian perspective, and extends the analysis to link the Theil indices to the Bayes factor which measures the posterior probability of fairness divided by that of unfairness. Section 5 presents

⁴This approach is not without its critics. See for example Kanbur and Wagstaff(2015).

an illustrative application of the results to data from the National Longitudinal Survey of Youth from the USA. Section 6 concludes.

2 Two Lotteries and Likelihood Ratio Tests for Fairness

In this section we introduce two income distribution processes which will frame our discussion and analysis of testing for fairness. We envisage them both in terms of “helicopter drops”, but what is being dropped differs between the two cases. In the first case the helicopter hovering over the population drops dollars which attach to individuals at random. In the second case there is initially a uniform flow of dollars to individuals on the ground before the helicopter comes on to the scene. What the helicopter does is to drop *stops* to this flow, stops which attach to individuals at random.

We refer to the two processes above as “The Helicopter Money Drop” and “The Helicopter Money Stop” respectively. These two processes generate lotteries for individuals and we can identify “fairness” within each process in a specified way. We now develop tests for fairness in each case and relate the tests to conventional inequality indices.

2.1 The "Helicopter Money Drop"

We envisage a helicopter dropping Y units of currency ("dollars" say and henceforth referred to as just "income units") onto a population one at a time during the year. The null is that the probability of each person receiving an income unit is equal and the alternative is that these probabilities are unequal.

Let y_i ($i = 1, 2, \dots, n$) denote individual i 's income, $Y = \sum_{i=1}^n y_i$ denote the total number of income units, $s_i = \frac{y_i}{Y}$ the income share of the i th individual and $\mu = \frac{Y}{n}$ be average income per head again measured in income units. Consider the process where each of the Y income units are allocated across the n individuals. Let the probability that individual i receives one unit be p_i . Each unit is assumed to be distributed independently so that the pdf of y_1, y_2, \dots, y_n is multinomial with likelihood

$$L(y_1, y_2, \dots, y_{n-1}; p_1, p_2, \dots, p_{n-1}) = \frac{Y!}{y_1! y_2! \dots y_n!} p_1^{y_1} p_2^{y_2} \dots p_n^{y_n} \quad (1)$$

where $p_n = 1 - \sum_{i=1}^{n-1} p_i$ and $y_n = Y - \sum_{i=1}^{n-1} y_i$.

As noted, we consider the null hypothesis that ex ante, each individual has an equal chance of receiving an income unit. Hence under H_0 we have the $n-1$ restrictions

$$H_0 : \quad p_1 = p_2 = \dots = p_{n-1} \quad (2)$$

and an alternative H_1 that one or more of these restrictions are violated. This null hypothesis encapsulates precisely and in analytical terms what we mean by

fairness/equity in the context of the current paper. We now show that the LR test is proportional to Theil's first inequality index.

The log likelihood (l) is

$$l = \ln(Y!) - \sum_{i=1}^n \ln y_i! + \sum_{i=1}^n y_i \ln p_i \quad (3)$$

A form for the LR test of (2) is found by comparing the values of (3) obtained when $\hat{p}_i = \frac{1}{n}$ (the null "estimates") with that obtained under $\hat{p}_i = \frac{y_i}{Y}$ (the alternative estimates). Hence we have

$$\begin{aligned} LR_1 &= 2(l_n - l_1) = 2 \left(\sum_{i=1}^n y_i \ln \frac{y_i}{Y} - \sum_{i=1}^n y_i \ln \frac{1}{n} \right) \quad (4) \\ &= 2 \left(\sum_{i=1}^n y_i (\ln s_i + \ln n) \right) = 2Y \sum_{i=1}^n s_i \ln n s_i \\ &= 2Y.T_1 = 2n\mu T_1 \quad (5) \end{aligned}$$

where T_1 is Theil's "first" index of inequality, $l_n(l_1)$ are the maximised log likelihoods assuming n (1) probability parameters and LR_1 is the likelihood ratio test that generates it.

Notice however that Y is scale dependent; if Y was originally dollars then going from a helicopter drop of dollars to one of cents for a given income level would increase Y and μ 100 fold. For fixed n it would also decrease the variance of income shares 100-fold. Both the dollar and cent drops are fair processes and there is no way a priori to choose the level of bundles being distributed. If this were the end of the story our contribution would be solely taxonomical; we would have shown that Theil's first inequality indicator (T_1) may be interpreted as a test of fairness under our a fair (Helicopter drop) income generation process. However we can in fact use T_1 as a test statistic by estimating μ under the null.

To implement the test based on T_1 we will need to estimate μ . We do so under the null and in a way that fully exploits the null property of (asymptotically) normally distributed income shares.

We specialise the fairness process to one where the number of income units per head μ is fixed but Y and n are large. We show in the annex that

$$LR_1^* = \sqrt{n}\{T_1 - b(\mu)\} \rightarrow N(0, a(\mu)) \quad (6)$$

$$\text{where } b(\mu) = \sum_{j=1}^{\infty} \frac{(2j)!!}{\mu^{j-1}(2j)!} \quad (7)$$

$$\text{and where } !! \text{ denotes double factorial} \quad (8)$$

where " \rightarrow " means "tends in distribution to" as $Y, n \rightarrow \infty$ with $\frac{Y}{n} = \mu$. We may estimate the mapping from μ to the variance term $a(\mu)$ by numerical

simulation. If μ were known we could implement (6) directly using a one tailed test (rejecting the null when we see above average values of T_1). But of course μ is unknown and we must estimate it. To implement (6) we replace μ with a (\sqrt{n}) consistent estimate $\hat{\mu}$. Under the null income shares are (asymptotically normal) that is in large samples, $ns_i \sim N(1, \frac{1}{\mu})$. Now consider the quantity $\tilde{s} = \sum_{i=1}^n |ns_i - 1|/n$. Under the null we can show that

$$\hat{\mu} = 2(\pi\tilde{s}^2)^{-1} \tag{9}$$

is a (\sqrt{n}) consistent estimator of μ . We propose to use $\hat{\mu}$ to estimate $b(\mu)$ in order to implement the test. Note that

$$\begin{aligned} \sqrt{n}\{T_1 - b(\hat{\mu})\} &= \sqrt{n}\{T_1 - b(\mu)\} + \sqrt{n}\{b(\mu) - b(\hat{\mu})\} = \\ &\sqrt{n}\{T_1 - b(\mu)\} + \sqrt{n}\left\{\frac{1}{\hat{\mu}} - \frac{1}{\mu}\right\}.b'(\hat{\mu}) + o(1) \end{aligned}$$

where $'$ denotes derivative with respect to $\frac{1}{\mu}$. The random variable $\frac{1}{\hat{\mu}}$ is the square of the sum of *iid* variates and so is asymptotically normal. Hence using $\hat{\mu}$ in LR_1^* results in an asymptotically mean zero normal variate whose variance is unknown but finite and depends only on μ . Given that $\hat{\mu}$ is consistent under the null we may use the estimated value in numerical simulations to estimate the corresponding variance of LR_1^* .

2.2 The "Helicopter Money Stop"

We now consider a second income process. Suppose that within any year each of our n individuals receive the same amount of income each "hour" and will continue to receive this hourly amount subject to a fixed hazard of exiting the income receipt process for the rest of the year. Under a fair null these hazards - the probability of "stopping" - are the same each hour for each individual regardless of how many hours they have survived the hazard. Suppose further we allow the time interval and hazard probability to both tend to zero. The result of this fair process is that each individual's income is a draw from an exponential pdf with the same mean β . This result maps back into the exponential distribution's natural role in modelling inter-arrival times. The income process is memory less in that if we see an individual with income x and who has survived the exit hazard their expected income will be $x + \beta$.

One way to interpret this fair process in developed economies would be as a metaphor for the number of remunerated working hours available to an individual within a year. Alternatively in developing economies where crop yield (and hence income) rely on (and are proportional to) rainfall, it could be the number of rainy days in the year.

Consider the *LR* test of a null that y_i has pdf $EXP(\beta)$ versus the alternative that y_i has pdf $EXP(\beta_i)$. Under the null we have

$$l_1 = -n \log \widehat{\beta}_0 - \sum_{i=1}^n \frac{y_i}{\widehat{\beta}_0} \quad (10)$$

where $\widehat{\beta}_0 = \bar{y}$ is the mle of β under H_0 and \bar{y} is average income (per head). Under the alternative, we have

$$l_n = - \sum_{i=1}^n \log \widehat{\beta}_i - \sum_{i=1}^n \frac{y_i}{\widehat{\beta}_i} \quad (11)$$

where $\widehat{\beta}_i$ is the mle of β for individual i under the alternative. A quick glance at the form of l_n shows that $\widehat{\beta}_i$ is just y_i so that l_n simplifies to

$$l_n = - \sum_{i=1}^n \log y_i - n \quad (12)$$

Straightforwardly then the likelihood ratio test of equal means is proportional to the log of the average minus the average of the logs i.e.

$$LR_2 = 2(l_n - l_1) = 2n(\log \bar{y} - \frac{\sum_{i=1}^n \log y_i}{n}) = 2nT_2 \quad (13)$$

Where T_2 is Theil's "second" inequality index and unlike the LR test in section 2.1 is scale free. Of course T_2 is the mean log deviation (MLD) which is fast becoming the workhorse of applied inequality measurement.

As was the case with T_1 we cannot appeal to standard likelihood theory to ascertain the limiting distribution of LR in the helicopter stop process because the model under the alternative is saturated. And in any event we wish to allow n - the number of parameters under the alternative - to be large for our asymptotics. However in the annex we show that

$$LR_2^* \sqrt{n}(T_2 - c) \rightarrow N(0, d) \quad (14)$$

where $c = \sum_{n=2}^{\infty} (-1)^n !n$ and where $!n$ denotes subfactorial. Both c and d were computed via numerical simulation as .57785 and .64973 respectively. Again we look to reject in the right hand tail only - we wish to reject for large values of T_2 not small.

3 Testing for Inequality of Opportunity

Above we showed that both of Theil's indices of inequality for a single population can be interpreted as test statistics of a null of a particular "fair" income generation process. Here we extend these results to the case where we are interested in comparing the extent of inequality between two or more subgroups of a population with that obtaining in the population as a whole. Roemer(1998) and others identify mutually exhaustive and exclusive characteristics of the population that

are considered not to be a result of individual decision making. Examples might be the (childhood) ethnicity/race of the individual or her parents' social class. Inequality arising because of these characteristics are termed "illegitimate" and the idea is to assess the extent to which inequality in the population as a whole can be ascribed to membership of these subgroups.

To assess the extent of illegitimate inequality Roemer presents a ratio of T_2 indices; the numerator is a measure of the extent of illegitimate inequality and the denominator of total inequality. This ratio has become the workhorse of a burgeoning empirical literature on the measurement of inequality of opportunity. We analyse this ratio now and show that it maps back into test statistics for fairness. But we also argue that this ratio is not a natural test of the null of a fair process; in fact the spirit of our approach leads us to consider only the variation arising from the numerator of this ratio.

In what follows we assume the population divides into two subgroups. We focus on two groups purely for notational simplicity and clarity - extension to k subgroups is trivial and briefly discussed below. We start with the helicopter stop process null and its corresponding test statistic T_2 .

Suppose there are $n_1(n_2)$ individuals drawn from two mutually exclusive and exhaustive subgroups of a population. We will assume the sample is ordered so that group one observations appear first. It is well known that Theil's second inequality index T_2 - the mean log deviation or MLD - may be decomposed into a within group index plus a between group index as follows

$$\begin{aligned} T_2 &= \{p_1(\log \bar{y}_1 - \overline{\log y_1}) + p_2(\log \bar{y}_2 - p_2 \overline{\log y_2})\} \\ &\quad + \{\log \bar{y} - (p_1 \log \bar{y}_1 + p_2 \log \bar{y}_2)\} \\ &= \{T_2^L\} + \{T_2^I\} \end{aligned}$$

$$\text{where } \bar{y}_1 = \frac{\sum_{i=1}^{n_1} y_i}{n_1}, \bar{y}_2 = \frac{\sum_{i=n_1+1}^n y_i}{n_2} \text{ and } p_i = \frac{n_i}{n}, i = 1, 2.$$

The within group component T_2^L is the weighted sum of the T_2 indices computed from each subgroup separately. In Roemer's formulation it represents the "legitimate" variation in the Theil index - the income variation that is a result of personal effort rather than inherited circumstance. The between component T_2^I is the Theil index computed assuming that all individuals within a subgroup have income equal to the group's average. In Roemer's formulation T_2^I measures the "illegitimate" component of the Theil index - the income variation that is a result purely of inherited circumstances.

Using the results in the previous sections we can readily show that T_2^I is merely $2n$ times the likelihood ratio test of the helicopter stop rule null (equal mean stopping times for all individuals regardless of subgroup) against the alternative that the two subgroups have different mean stopping times i.e.

$$\begin{aligned} 2nT_2^I &= l_2 - l_1 \\ \text{where } l_2 &= \max_{\beta_1, \beta_2} \{-n_1 \log \beta_1 - n_2 \log \beta_2 - \sum_{i=1}^{n_1} \frac{y_i}{\beta_1} - \sum_{i=1}^{n_2} \frac{y_i}{\beta_2}\} \end{aligned}$$

and where l_1 is as defined in (10). Once more we have a different and concrete interpretation of a Theil index as a test statistic of a null of a fair income process against an alternative that circumstances outside the individual's control affect this process. The statistic is a χ_1^2 under the null so instead of presenting the numerical value of the index one may present its corresponding p-value. When comparing illegitimate Theil indices from two populations we could then compare p-values rather than Theil index levels.

However as we have noted Roemer uses a ratio of illegitimate to total variation as a measure of inequality of opportunity. This share - R_2 say - can be written as

$$R_2 = \frac{T_2^I}{T_2} = \frac{\ln \bar{y} - p_1 \ln \bar{y}_1 - p_2 \ln \bar{y}_2}{\ln \bar{y} - \ln y} \quad (15)$$

where T_2^I is the illegitimate between group Theil index. Given the decomposition above this has a clearcut interpretation; it is the proportion of income variation (as measured by Theil) attributable to illegitimate inequality. Driving our approach to its logical conclusion we could also interpret R_2 as a ratio of LR tests with the same null of fairness as per our helicopter stop process. Explicitly, a simple application of the analysis in section 2 shows that

$$R_2 = \frac{2nT_2^I}{2nT_2} = \frac{LR_2^I}{LR_2}$$

where LR_2 tests the fair helicopter stop rule null against an alternative that individuals have different means and LR_2^I is the LR test of the same null against the alternative that each subgroup has its own distinct mean (each individual draws from her particular subgroup's exponential pdf).

As we have already indicated, using the variation in numerator and denominator for a test runs against the spirit of our approach. Even if we could deduce the null distribution of R_2 , the test would have power properties that would make rejections and failures-to-reject hard to interpret. In particular the power of the test against illegitimate inequality is moderated by the extent of legitimate variation. Instead we suggest modifying the ratio to use only the variation in the numerator of R_2 . Explicitly, note that nR_2 can be written as

$$nR_2 = \frac{2nT_2^I}{2T_2} = \frac{n(\ln \bar{y} - p_1 \ln \bar{y}_1 - p_2 \ln \bar{y}_2)}{\ln \bar{y} - \ln y} = \frac{\chi_1^2}{c} + o(1)$$

where $c(= .57785)$ is the null limit of T_2 as defined above.

This statistic is a standard LR test. It follows a χ_1^2 distribution. It is this statistic - extended to the case of k groups - that we compute in our empirical work in section 5.

We now turn to consider the helicopter drop process (where Theil's first inequality index T_1 is the relevant data quantity). Following previous logic we could write T_1 as (proportional to) the sum of two LR tests as follows

$$2n\mu T_1 = 2\{l_2 - l_1\} + 2\{l_n - l_2\} = \{2n\mu T_1^I\} + \{2n\mu(T_1 - T_1^I)\} \quad (16)$$

$$\text{where } T_1^I = \sum S_i \ln(\rho_i \cdot S_i) \quad (17)$$

where S_i is now group i 's share of the total income Y and where $\rho_i = n/n_i$. The first statistic in (16) - $2n\mu T_1^I$ - tests a null of equal means against an alternative each group has its own mean. If we treat all of the people in each group as a single person then we see it is a measure similar to T_1 the difference being the replacing of n in T_1 with ρ_i . The second statistic in (16) tests the null that each subgroup has its own mean versus an alternative that each individual has her own separate mean. If we accept that some proportion of income inequality is due to "legitimate" reasons then our approach suggests we should focus on the first component only. Analogous to nR_2 above we propose to compute the ratio nR_1 where

$$nR_1 = \frac{2n\mu T_1^I}{2\mu T_1} = \frac{nT_1^I}{T_1} = \frac{\chi_1^2}{const} + o(1) \quad (18)$$

$$\text{where } const = 2\mu \cdot plim T_1 = \mu b(\mu)$$

where $b(\mu)$ is as defined above. Unlike before we do not need to estimate μ as long as it is larger than unity. The parameter μ is the average number of income "packets" an individual receives. In our sample below, average income is around \$15,000 so a value of unity for μ would imply the Helicopter is dropping packets of this size. Whilst technically this is fair by our reasoning such large packets would certainly guarantee - ex ante - a very high variation in incomes after the Helicopter Drop has occurred. A sensible prior view therefore would be that μ is larger than one. Finally on this point in the empirics below our estimate of μ (under the null) is around 2.5. We therefore assert that $\mu b(\mu) \approx \frac{1}{2}$ to compute (an approximation to) nR_1 as

$$nR_1^* \approx 2 \frac{nT_1^I}{T_1}$$

and compare it to critical values of the χ_1^2 distribution. Finally note that because $T_1 - T_1^I > 0$ R_1 is always between zero and one the decomposition of T_1 into T_1^I and a remainder term allows us to interpret R_1 as the proportion of T_1 that results from illegitimate income variation.

Extending the above to k mutually exclusive and exhaustive groups (for example white and male, white and female, non-white and male, non-white and female) is trivial. Under the alternative the log likelihood is "additively separable"; the k parameters appear separately in k separate additive terms. The formulae presented above for the helicopter stop rule remain intact and unchanged. For the helicopter drop (16) we must replace the number 2 with k in the formula. Of course in this case the resulting tests are chi-squared statistics with $k - 1$ rather than 1 degrees of freedom.

We close by noting that the statistics above are analogous to the " nR^2 " F-test of significance of the regression of income on group dummies. Under the null of no illegitimate variation this regression has no explanatory power whilst under the unfair alternative it does.

4 A Bayesian Perspective

Up until now we have motivated use of Theil's indices using a purely classical view that we wish to make a "decision" as to whether or not our income process is fair. An alternative motivation for their use comes through their relationships with Bayes' factor as we now show.

If we treat the "fair" and "unfair" processes (null and alternatives above) as models and if we assign equal prior probability to their respective "truth" we may interpret the likelihood ratio in terms of the ratio of posterior probabilities of fair to unfair models. In particular Kass and Raftery (1995) show that for large n that the log of the Bayes factor ($\log(BF)$) is approximately

$$\log(BF) \approx LR - k \cdot \log(n)$$

where LR is the difference in the log likelihoods between unfair and fair models and k is the number of extra parameters moving from the fair to unfair models.

There are several things about the use of $\log(BF)$ that make it attractive in our context particularly when we wish to compute Theil indices from different datasets with possibly different sample sizes (n). Firstly it is not a decision making tool. Instead - as we have noted above - it gives us the posterior probability that the process is unfair divided by the posterior probability it is fair⁵. It does not therefore require that one or other of our "models" be true - it just offers a measure of the models' relative concordance with the data. By contrast the p-values of the classical tests (above) are only valid as probability statements when the null is absolutely true. Additionally even if the estimated parameters of the processes were close to one another in the two models in a quantitative sense, the p-value will still tend to zero with n . Hence the p-value may mask this quantitative closeness. An additional problem is that comparisons of results are difficult across different datasets if they have different numbers of datapoints. We may of course solve this problem by simply bootstrapping the data and computing estimated p-values over a range of n ⁶. However this procedure is computationally inconvenient. A further - admittedly minor - issue associated with large n is that p-values may be too small to compute accurately.

Set against these arguments in favour of the BF is the fact that to compute it for our helicopter drop process we need an estimate of the scaling parameter μ and this is only identified under the null Whilst this (nuisance) parameter

⁵Making the posterior probability of the alternative the numerator and that of the null the denominator is a convention we adopt to match the spirit of the LR test.

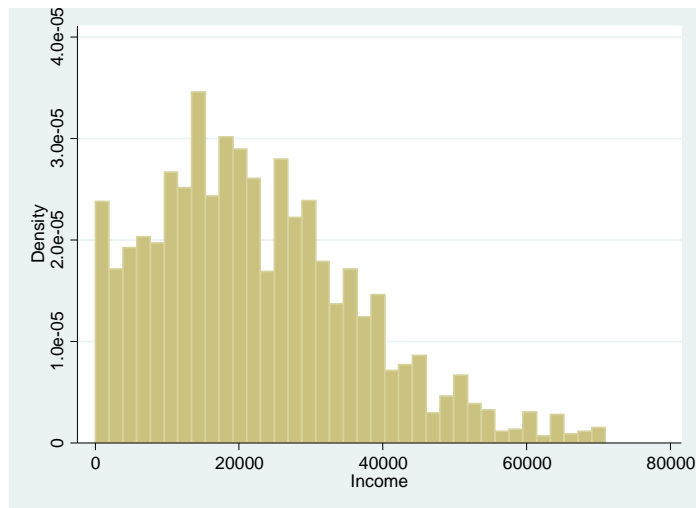
⁶A bootstrap based simply on random resampling of the data is valid under our null hypothesis of fairness because under the null incomes are indeed independent of one another.

takes the same value in both models the estimate we obtain is only consistent under the null. By contrast nR_1^* may be computed without an estimate of μ as can nR_2 (which has no scaling parameter). Therefore we only compute the BF 's for the analysis of illegitimate income variation.

5 Empirical Illustration

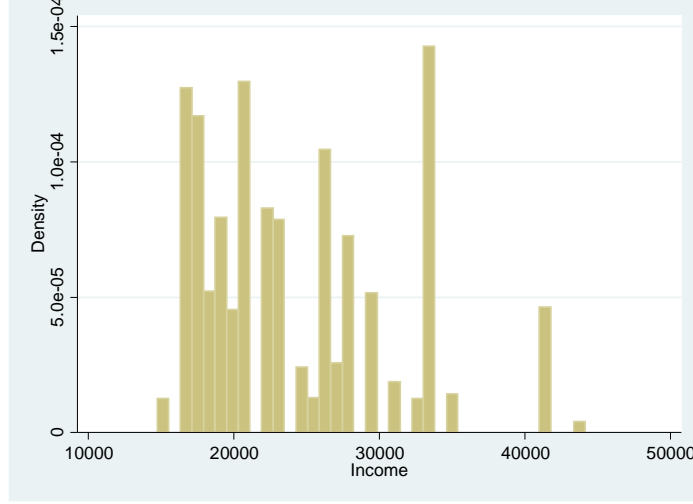
To provide an illustrative application of the above tests we drew a sample of incomes from the 1979 NLSY for the year 1995. The 1979 NLSY is a nationally representative sample of people who were between the ages of 14 and 22 in 1979 so by the year 1995 nearly all members of the sample will be young adults who have completed their education. As always in a single cross sectional snapshot we are using current income as a proxy for permanent income which is obviously flawed. In particular many people in our sample recorded zero income for the year. Either they are not in the formal workforce (e.g. married workers in the home) or they are unemployed. Given the purpose of our exercise is only illustrative, we do not attempt to adjust income of individuals to take account of the disconnect between permanent and current income and nor do we try and estimate household income. We simply drop those people who recorded zero income from the sample. After dropping data we were left with 5736 individuals.

Figure 1: Figure 1: Raw Income Distribution



We compute the four test statistics LR_1^* , LR_2^* , nR_1^* and nR_2 . To compute the last two of these - based on the illegitimate variation in incomes - we control for race (white/nonwhite), gender, region of upbringing (South/ Non South), whether or not the individual has a work hindering health condition and whether

Figure 2: Figure 2: Distribution of Incomes at Group Means



or not the child was raised by biological parents. This gives ($k =$) 32 mutually exclusive and exhaustive groups of people. Figure 1 shows the distribution of raw incomes (total variation) whilst Figure 2 shows the distribution when each individual is assumed to earn their group mean income (illegitimate variation). The graphs are radically different. In terms of an "eyeball" metric, the empirical pdf of raw incomes in Figure 1 show some resemblance to an exponential pattern and also resemble a left truncated normal pdf. Figure 2 on the other hand shows a distribution of illegitimate income that neither resembles exponential nor normal.

<i>statistic</i>	LR_1^*	LR_2^*	nR_1^*	nR_2	BF_1	BF_2
	-3.30	-22.81	407.66	333.19	129.97	55.56
<i>p - value</i>	$1 - .0005$	$1 - 1.15 \times 10^{-115}$	5.6×10^{-38}	4.37×10^{-23}		

Notes: LR_i^* are standard normal and nR_1^* and nR_2 are χ_{32}^2 variates⁷.

Table 1 presents the four test values and their p-values. We see that LR_1^* (helicopter drop process) and LR_2^* (helicopter stop process) are negative implying the gap in log likelihoods (which is always strictly positive) is less than the respective expected values under the null. Our rejection region is always in the right tail (only higher than expected Theils can reject not lower). Put another way, the Theil statistics lie below their null expected value. We showed in an earlier version of this paper that the LR_i^* tests are positively related to

⁷We approximate the p-values using

$p = \frac{f(x)^2}{2f'(x)}$ This is the area of the triangle whose height is the height of the pdf at the statistic's value x , and whose hypotenuse is the tangent to the pdf at x . It is an underestimate.

the difference between empirical and theoretical moments. We conjecture that the statistics are both negative because the tails of the empirical distribution are less fat than their normal and exponential counterparts would predict.

The "illegitimacy" statistics for the two processes nR_1^* and nR_2 both resoundingly reject their respective nulls with miniscule p-values. As we have already noted, small p-values are the consequence of large samples; with over 5000 observations, even local departures from the null will result in very large test statistics.

Finally the Table also shows the two Bayes factors for fairness versus illegitimate variation. Here we see a reinforcement of the corresponding test results for nR_1^* and nR_2 ; the unfair alternative of illegitimate variation is very much more likely in posterior terms than the null of legitimate variation.⁸

6 Discussion and Conclusion

Consider a standard inequality index such as Mean Log Deviation (MLD), which is of course Theil's second inequality index, T_2 . One interpretation of this inequality index is that it is related to the loss in social welfare from inequality of income. To see this, let social welfare W be the sum of log incomes across n individuals in a society

$$W = \sum_{i=1}^n \ln y_i$$

Then, following Atkinson (1970), the equally distributed equivalent income Y_{ede} is the mean income which, if distributed equally, will give the same level of social welfare as the current unequal distribution of income. Denoting μ as the current mean, it can be shown that:

$$Y_{ede} = \mu e^{-T_2}$$

Thus, within a particular social welfare framework, T_2 measures the social cost of inequality in a very precise sense. Other inequality indices can also be interpreted in this way.

However, this approach is vulnerable to the critique that it is focused only on evaluation of outcomes, not of the process. There is a large philosophical literature (for example, Sen, 1979, and Dworkin, 1981) which opens up this question. There is an equally large literature on income mobility (for example, Chetty et. al; 2014, Corak, 2013, and Kanbur and Stiglitz, 2016) which

⁸Classical tests and Bayes factors are conceptually and statistically completely different data measures and as a result do not always point in the same direction vis a vis their relative support for the null and alternative. For example in their seminal paper on Bayes factors Kass and Raftery (1995) note "...a dramatic example with $n = 113,566$...A substantively meaningful model that explained 99.7% of the deviance was rejected with a standard chi-squared test with a p-value of about 10^{-120} but was nevertheless favored by the Bayes factor" (op cit. p 789). In other empirical work not reported here we have indeed found a miniscule p-value for the classical test "overturned" by a negative BF.

emphasizes distributional processes and in particular random shocks which can affect income status. Moreover, the recent literature on equality of opportunity (Roemer, 1998, Trannoy and Roemer, 2015) focuses on that portion of overall inequality which can be attributed to factors outside and within individual control.

Returning to the Friedman example in the introduction, suppose the observed inequality was the outcome of a lottery equally available and freely chosen by individuals who were identical *ex ante*. A normative assessment might conclude that the process was fair even though it produced unequal outcomes. The real question, in this perspective, is the extent to which the inequality of outcomes is the result of *ex ante* inequalities, differences in the lotteries faced by different individuals.

But if we accept the process perspective, does this mean that we jettison all use of inequality indices, for example the mean log deviation measure T_2 ? The answer given in this paper is no. The observed T_2 , while no longer acceptable as a measure of welfare loss as above, can nevertheless serve as a test statistic for “fairness” in a well-defined sense, in the context of a class of income distribution processes. In particular, we have shown that for the “helicopter stop” process, T_2 can indeed be used as a test for fairness. Similarly, for the “helicopter drop” process, we have shown that Theil’s first index T_1 , suitably transformed, can also serve as a test of fairness. Further, we have also shown how we can test inequality of opportunity using the same Theil indices. The response to the process critique is thus not to stop calculating inequality indices, but to use them as tests of fairness or as estimates of the chance the process was unfair relative to the chance it was fair. The comparison is not whether the index is high or low, but whether or not it rejects fairness or alternatively whether or not it implies that fairness is more probable in a posterior sense.

It is a natural question to ask if more generic fair processes may be tested in this way. To put the question another way, would our respective tests reject their respective nulls if incomes were generated by other “fair” income processes? That is would we reject “fairness” in favour of unfairness in such cases? The answer to this question is ambiguous. Whilst the two “fair” processes we outlined above lead to unique and distinct pdf’s for income which form the null of our respective tests, the issue of rejection is one of power. In particular LR_1^* and LR_2^* will tend to reject when the empirical distribution has a fat right hand tail and this is a characteristic shared by the alternative hypotheses in both the Helicopter Drop and Stop processes. At the same time it is entirely possible that LR_1^* (LR_2^*) may reject when the fair null of the Helicopter Stop(Drop) process holds true.

To expand on the previous point further consider an attempt to generalise fairness by moving away from the specifics of the processes we have identified. One obvious generalisation is to envisage each individual’s income as an independent draw from the same single but unspecified/unknown pdf. It is easy to see that testing this generic fair null is impossible using a single cross section of incomes; for any set of income realisations we can always find a pdf for which the null is not rejected and a pdf for which it is. For example - and trivially -

we could never reject the null that incomes were independent draws from the observed empirical pdf itself! To gain traction from a single cross section there must be a specific income process that is the subject of the test.

The tests we specify for our two "fair" nulls are not unique. Both test a null that shares are drawn from a particular pdf - normal and exponential respectively - and there are several tests that could have been used (e.g. tests of "excess" kurtosis where "excess" is with respect to the null distribution). Although our statistics arise naturally from Likelihood Ratio principles we should note the caveat that the usual rationale for LR tests - that they have optimal power against simple alternatives (via the Neyman Pearson Lemma) - is only relevant to the tests of illegitimate inequality rather than the tests of overall inequality LR_1^* and LR_2^* . In an earlier version of this paper we tentatively explored other "fair" income processes such as ones where individuals draw incomes from the same uniform pdf. This led to test statistics that were similar to but not the same as existing inequality indices. A more comprehensive investigation along these lines is an interesting agenda for future research.

Annex Proof of (6) and (14)

Define $s_i^* = ns$ then we have

$$T_1 = \frac{\sum_{i=1}^n s_i^* \log(s_i^*)}{n}$$

This shows that T_1 is the sum of n *iid* variates each with finite variance and so is asymptotically $(y, n \rightarrow \infty, \frac{Y}{n} = \mu)$ normal. Now expand each $s_i^* \log(s_i^*)$ element in a Taylor series around $s_i^* = 1$ (the null mean of s_i^*) to get

$$\begin{aligned} T_1 &= 0 + \frac{\sum_{i=1}^n (s_i^* - 1)}{n} + \sum_{j=2}^{\infty} \sum_{i=1}^n \frac{(s_i^* - 1)^j}{j!n} \\ &= \sum_{j=1}^{\infty} \frac{\sigma^{2j}}{(2j)!} + o(1) \\ \text{where } \sigma^{2j} &= \frac{(2j)!!}{\mu^j} \end{aligned}$$

The last equality follows from the asymptotic normality of shares. In the last sum, "!!" stands for double factorial. μT_1 is the sum of *iid* finite variance variates and so is asymptotically normal. We have just established then that its asymptotic mean is $b(\mu)$ as given in the text. We need to now establish that its variance is a constant and in particular that it does not depend on μ . We may write

$$\mu T_1 = \frac{1}{n} \sum_{i=1}^n (\mu s_i^*) \log(\mu s_i^*) - \mu \log \mu$$

The variates μs_i^* are asymptotically *iid* normal with unit variance. All central moments of μs_i^* are therefore defined constants. The asymptotic pdf of $(\mu s_i^*) \log(\mu s_i^*)$ is therefore free of nuisance parameters.

To prove (14) first of all note that under the null the scale variate of the exponential washes out of T_2 . Then expanding in a Taylor series around $y_i = 1$ gives

$$T_2 = \sum_{j=2}^{\infty} (-1)^j \frac{\widehat{\sigma(j)}}{j!}$$

where $\widehat{\sigma(j)} = \frac{\sum_{i=1}^n (y_i - 1)^j}{n}$ is the j th sample central (around 1) moment of income. All sample moments are consistent for their theoretical counterparts and the weights are geometrically declining so the *plim* of this expression simply

requires substitution of actual for sample moments of the exponential pdf with unit scale parameter. This gives the expression for c in the text.

Asymptotic normality is established by noting that each sample moment is the sum of n *iid* variates so each sample moment is asymptotically normal.

References

- [1] Atkinson, Anthony B. (1970). "On the Measurement of Inequality." *Journal of Economic Theory* vol 2, 244-263.
- [2] Bourguignon, F., Ferreira, F. H. G., and Walton, M. (2006). Equity, efficiency and inequality traps: A research agenda. *The Journal of Economic Inequality*, 5(2):235_256.
- [3] Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. "Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States", *Quarterly Journal of Economics* 129(4), pages 1553-1623.
- [4] Cowell, F. A. (2011). *Measuring Inequality*. Oxford University Press, Oxford.
- [5] Corak, Miles. 2013. "Income Inequality, Equality of Opportunity, and Intergenerational Mobility." *Journal of Economic Perspectives*, Volume 27, Number 3, Summer 2013, pages 79–102.
- [6] Dworkin, R. (1981). What is Equality? Part 2: Equality of Resources. *Philosophy & Public Affairs*, 10(4), pages 283-345.
- [7] Ferreira, F. H. G. and Gignoux, J. (2011). The Measurement of Inequality of Opportunity: Theory and an Application to Latin America. *Review of Income and Wealth*, 57(4), pages 622-657.
- [8] Kanbur, Ravi and Joseph E. Stiglitz. 2016. "Dynastic Inequality, Mobility and Equality of Opportunity" *Journal of Economic Inequality*, Vol. 14, No. 4, pages 419-434.
- [9] Kanbur, R. and Wagstaff, A. (2016). How useful is inequality of opportunity as a policy construct? In Basu, K. and Stiglitz, J. E., editors, *Inequality and Growth: Patterns and Policy: Volume I: Concepts and Analysis*, pages 131-150. Palgrave Macmillan UK, London.
- [10] Robert E. Kass, Adrian Raftery, 1995. "Bayes Factors", *Journal of the American statistical Association*, vol 90, pages773-795.
- [11] Roemer, John. 1998. *Equality of Opportunity*, Cambridge, MA: Harvard University Press.
- [12] Roemer, J. E. and Trannoy, A. (2015). Equality of Opportunity. In Atkinson, A. B. and Bourguignon, F., editors, *Handbook of Income Distribution*, volume 2, pages 217-300. Elsevier, Amsterdam.
- [13] Sen, A. (1979). *Equality of What? The Tanner Lectures on Human Values*, 22 May 1979.