



# Measuring inequality beyond the Gini coefficient may clarify conflicting findings

Kristin Blesch <sup>1,2,3</sup>✉, Oliver P. Hauser <sup>4,5</sup>✉ and Jon M. Jachimowicz <sup>6</sup>✉

**Prior research has found mixed results on how economic inequality is related to various outcomes. These contradicting findings may in part stem from a predominant focus on the Gini coefficient, which only narrowly captures inequality. Here, we conceptualize the measurement of inequality as a data reduction task of income distributions. Using a uniquely fine-grained dataset of  $N = 3,056$  US county-level income distributions, we estimate the fit of 17 previously proposed models and find that multi-parameter models consistently outperform single-parameter models (i.e., models that represent single-parameter measures like the Gini coefficient). Subsequent simulations reveal that the best-fitting model—the two-parameter Ortega model—distinguishes between inequality concentrated at lower- versus top-income percentiles. When applied to 100 policy outcomes from a range of fields (including health, crime and social mobility), the two Ortega parameters frequently provide directionally and magnitudinally different correlations than the Gini coefficient. Our findings highlight the importance of multi-parameter models and data-driven methods to study inequality.**

Economic inequality is at high levels around the world and continues to rise in many countries<sup>1,2</sup>. A wealth of prior research has explored the outcomes of such high inequality levels. While initial research has often found negative associations with wide-ranging policy outcomes (for an overview, see ref. <sup>3</sup>), subsequent work has arrived at more conflicting findings. For example, different studies have found that the relationship between economic inequality and obesity is both positive<sup>4</sup> and negative<sup>5</sup>. Similarly, different studies have found that economic inequality is associated with both lower and higher subjective well-being (for a meta-analysis, see ref. <sup>6</sup>). Finally, some studies have found that economic inequality is related to less prosociality<sup>7</sup>, which other studies do not corroborate<sup>8</sup>. While these studies differ in their conclusions, they all share one attribute: they measure and operationalize inequality through a single-parameter measure, predominantly the Gini coefficient. Here we suggest that this singular focus on the Gini coefficient may lie at the heart of several of these conflicting findings of inequality and its correlates. We demonstrate not only that single-parameter inequality measures such as the Gini coefficient are unable to capture crucial information contained in income distributions but also that moving beyond these types of measures by replacing them with more comprehensive measures can help resolve extant tensions in the field.

Across the social sciences, the Gini coefficient is by far the most popular measure of economic inequality<sup>9</sup> and is often used to inform policy debates<sup>10</sup> and justify political decisions (for example, see ref. <sup>11</sup>). Several reasons exist for the widespread use of the Gini coefficient, including its ease of interpretation<sup>12–14</sup> and access, as many official bureaus of statistics across the world publish this summary statistic regularly<sup>15</sup>. This could result in a self-sustaining feedback loop: because of the widespread availability of the Gini coefficient, researchers frequently use this inequality measure in their work; this may lead statistics bureaus to continue providing this measure

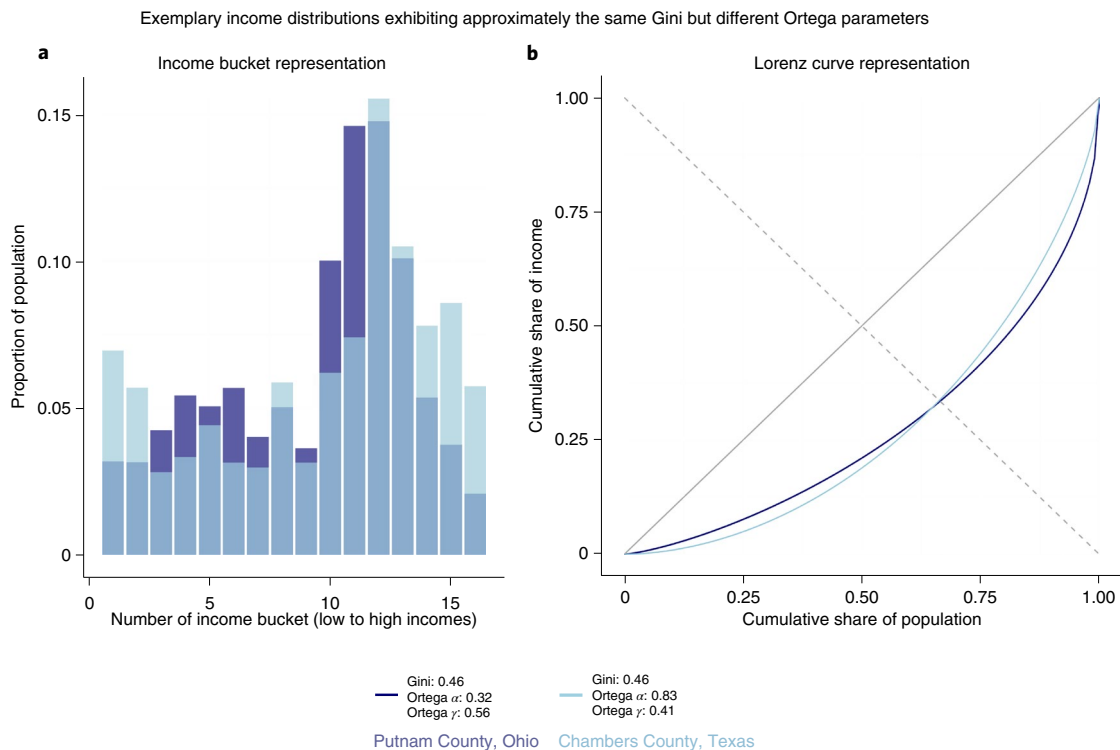
to researchers instead of exploring potential alternative or additional inequality measures. The widespread prevalence of the Gini coefficient may even give the impression that this measure is the *only* or *best* way to capture inequality<sup>16</sup>.

However, several drawbacks of the Gini coefficient are well-known<sup>12,17,18</sup>. One of the main criticisms pertains to its inability to adequately distinguish between different income distributions that result in the same Gini coefficient<sup>9,12,14,17,19,20</sup>. This shortcoming becomes particularly apparent when investigating income distributions through the lens of Lorenz curves, which map absolute income distributions on a relative scale (see our discussion of Fig. 1 below; see also ref. <sup>21</sup>). While in some cases, different distributions resulting in the same Gini coefficient may be a desirable property, we argue that focusing only on the overall concentration of inequality—as captured by the Gini coefficient<sup>13</sup>—is insufficient to fully appreciate how inequality affects important policy outcomes. Note that this problem does not plague the Gini coefficient alone: all inequality measures require some decisions around the compression of information, which is particularly salient in single-parameter measures of inequality, as they attempt to condense a lot of information into a single parameter<sup>17</sup>. As a result, critical aspects of income distributions may be missed by the Gini coefficient, which, we propose, may partially underlie prior mixed findings in associations with policy outcomes. We argue that this shortcoming, among others<sup>12,18</sup>, therefore necessitates alternative approaches to more comprehensively capture the non-equal distribution of income.

Alternatives to the Gini coefficient in prior literature predominantly build on two streams of research. First, some prior work suggests replacing the Gini coefficient with other single-parameter measures of inequality<sup>9,22,23</sup>. Many such alternative measures have desirable properties, such as the Zanardi index, which refines the Gini coefficient to capture asymmetries in an income distribution<sup>24</sup>. However, there is no clear consensus on what alternative measure

<sup>1</sup>Seminar for Statistics, ETH Zurich, Zurich, Switzerland. <sup>2</sup>Faculty of Mathematics and Computer Science, University of Bremen, Bremen, Germany. <sup>3</sup>Leibniz Institute for Prevention Research & Epidemiology—BIPS, Bremen, Germany. <sup>4</sup>Department of Economics, University of Exeter Business School, University of Exeter, Exeter, UK. <sup>5</sup>Behavioural and Experimental Data Science, Institute for Data Science and Artificial Intelligence, University of Exeter, Exeter, UK.

<sup>6</sup>Organizational Behavior Unit, Harvard Business School, Harvard University, Boston, MA, USA. ✉e-mail: [blesch@leibniz-bips.de](mailto:blesch@leibniz-bips.de); [o.hauser@exeter.ac.uk](mailto:o.hauser@exeter.ac.uk); [jjachimowicz@hbs.edu](mailto:jjachimowicz@hbs.edu)



**Fig. 1 | Plotting the distributions of income for Putnam County, Ohio, and Chambers County, Texas.** **a**, Income bucket representation: the proportion of earners per income bucket is shown for two counties that have approximately the same Gini coefficient (0.46). **b**, Lorenz curve representation: the same income distributions are plotted as Lorenz curves, which reveals that while overall levels of inequality are the same for both distributions (that is, the same area under the curve), where inequality is concentrated differs between the counties.

to use<sup>25</sup>, in part because no clear criteria have been established to decide which measure is best. The second approach suggests using the Gini coefficient in combination with another measure<sup>17</sup>. For example, Sitthiyot and Holasut<sup>14</sup> suggest using the Gini coefficient and the income share held by the top and bottom 10% of the population as a joint inequality measure. The attempt to use multiple measures of inequality simultaneously to capture more features of the income distribution is intuitively appealing; however, this approach lacks a systematic analysis to ascertain whether it truly captures all relevant information contained in an income distribution. Indeed, for this approach to be informative, multiple measures of inequality need to convey mostly *unique* information about the income distribution.

Here we propose a systematic approach to identify which inequality measure is the most appropriate for a given dataset by capturing the greatest amount of relevant information about an empirical income distribution. Any measure of inequality requires researchers to define how to bundle relevant information present in the income distribution into key parameters, and to determine what attributes of the income distribution these should represent. Our starting point is the notion that income distributions that exhibit inequality are, by definition, non-equal, and that capturing their shape is of key interest in measuring inequality. Put differently, we conceptualize the path from income distributions to inequality measures as a data reduction task. We note that this approach does not focus on axiomatic properties that need to be satisfied for an appropriate measure of inequality but instead is a bottom-up and data-driven approach that draws on the shape of actual income distributions. Our goal is to bundle relevant information present in income distributions into a reasonable number of numerical values for later use as measures of inequality to evaluate whether the different attributes of income distributions captured through this approach explain meaningful variance in important outcomes.

To do so, we employ a jointly theoretically derived and data-driven approach to systematically determine how many and what kind of parameters should be used to capture relevant information contained in an income distribution. We first draw on prior research to examine theoretical models that have been proposed to model income distributions. Next, we combine data from several sources to create a unique dataset with  $N=3,056$  real-world income distributions at the US county level, including uniquely fine-grained information on top-income earners. This allows us to combine maximum likelihood estimation (MLE) with a systematic evaluation framework based on information criteria to determine the optimal parameters necessary to characterize income distributions. Finally, we move to real-world applications: we study the correlations of the best-fitting model in our dataset with 100 wide-ranging policy outcomes, allowing us to shed light on extant tensions in the literature and highlighting the importance of moving beyond just evaluating how much inequality exists towards considering where inequality is concentrated.

To illustrate the benefits of moving beyond existing inequality measures, consider the two income distributions depicted in Fig. 1, which are based on real-world data from two US counties (Putnam County, Ohio, and Chambers County, Texas). We chose these counties because, when measured by the Gini coefficient, they seem to exhibit the same level of inequality (that is, a Gini of approximately 0.46). However, when considering the income bucket representation (Fig. 1a) and especially the Lorenz curve representation of incomes (Fig. 1b), it becomes evident that the distribution of income differs between the counties. Figure 1a shows that, at different income levels, the two counties share different levels of overlap in the number of people earning a certain amount of money. While income bucket representations are simple and easy to understand, they are less suitable for comparing different income distributions.

**Table 1 | Theoretically derived parametric theoretical models considered in our empirical analyses**

Originates from	Lorenz curve $\eta(u)$
1. Pareto distribution	$1 - (1 - u)^{1-1/\alpha}$
2. Lognormal distribution	$\Phi(\Phi^{-1}(u) - \sigma)$
3. Gamma distribution	$G(G^{-1}(u; \sigma); \sigma + 1)$
4. Weibull distribution	$G(-\log(1 - u); \frac{1}{\alpha} + 1)$
5. Generalized gamma distribution	$G(G^{-1}(u; p); p + \frac{1}{\alpha})$
6. Dagum distribution	$B(u^{1/q}; q + \frac{1}{\alpha}, 1 - \frac{1}{\alpha})$
7. Singh-Maddala distribution	$B(1 - (1 - u)^{1/q}; 1 + \frac{1}{\alpha}, q - \frac{1}{\alpha})$
8. GB1 distribution	$B(B^{-1}(u; p, q); p + \frac{1}{\alpha}, q)$
9. GB2 distribution	$B(B^{-1}(u; p, q); p + \frac{1}{\alpha}, q - \frac{1}{\alpha})$
10. Kakwani and Podder <sup>62</sup>	$ue^{-\beta(1-u)}$
11. Rasche et al. <sup>63</sup>	$(1 - (1 - u)^\alpha)^{1/\beta}$
12. Ortega et al. <sup>26</sup>	$u^\alpha(1 - (1 - u)^\beta)$
13. Chotikapanich <sup>64</sup>	$\frac{e^{ku} - 1}{e^k - 1}$
14. Sarabia et al. <sup>58</sup>	$u^{\alpha+\gamma}[1 - a(1 - u)^\beta]^\gamma$
15. Abdalla and Hassan <sup>65</sup>	$u^\alpha(1 - (1 - u)^\beta)e^{ku}$
16. Rhode <sup>66</sup>	$u \times \frac{\beta-1}{\beta-u}$
17. Wang et al. <sup>67</sup>	$\delta u^\alpha [1 - (1 - u)^\beta] + (1 - \delta)[1 - (1 - u)^\beta]^\nu$

Rows 1-9: Lorenz curve models from distributional origin. Rows 10-17: functional forms proposed to model Lorenz curves. Model 14 is recognized as a family of Lorenz curves but not proposed as a Lorenz curve specifically. As this family is the most general form of the specific Lorenz curve that Sarabia et al.<sup>58</sup> propose, we use it as a four-parameter Lorenz curve<sup>44,68-71</sup>.  $\eta$  denotes the cumulative percentage of income,  $u$  denotes the cumulative percentage of the population,  $\Phi(\cdot)$  is the cumulative distribution function of the standard normal distribution,  $G(\cdot)$  is the incomplete gamma function ratio and  $B(\cdot)$  is the lower incomplete beta function ratio as defined in the Notation Preface in the Supplementary Information. Details on the parameter restrictions are given in Supplementary Section 1.

Figure 1b displays the corresponding Lorenz curves of the two counties, depicting the cumulative share that each percentile of the income distribution holds. Lorenz curves are particularly useful for comparing income distributions because they are scale-free and can be used regardless of the average income in a population. Lorenz curves also visually depict why the Gini coefficients of the income distributions are the same, given that it is proportional to the area spanned between the diagonal line and the Lorenz curve. This area is equally large for both counties, which is why they yield the same Gini coefficients. However, the Gini coefficient does not take into account that the Lorenz curve of Putnam County, Ohio, bends more intensely within the top of the income distribution, whereas the Lorenz curve of Chambers County, Texas, is more strongly bent within the bottom of the income distribution. For example, we can see from the estimated Lorenz curves that the top 10% of the population in Putnam County, Ohio, possess 38.7% of the total income in that county, whereas in Chambers County, Texas, the top 10% hold 32.1% of total income. Given their useful features, we subsequently aim at representing income distributions using Lorenz curves.

**Results**

**Fitting Lorenz curves.** We begin by sourcing an extensive range of proposed Lorenz curve models in the literature as a starting point for our data-driven approach (Table 1), arriving at a total of 17 different models. Then, using MLE (Methods), we estimate and subsequently evaluate the fit of each model in every county in our dataset with a Borda count voting procedure, assigning more points to better-fitting models. The Borda count enables us to identify the ‘winning’ model among the proposed Lorenz curve models across

**Table 2 | Borda count result using AIC<sub>c</sub> as information criterion**

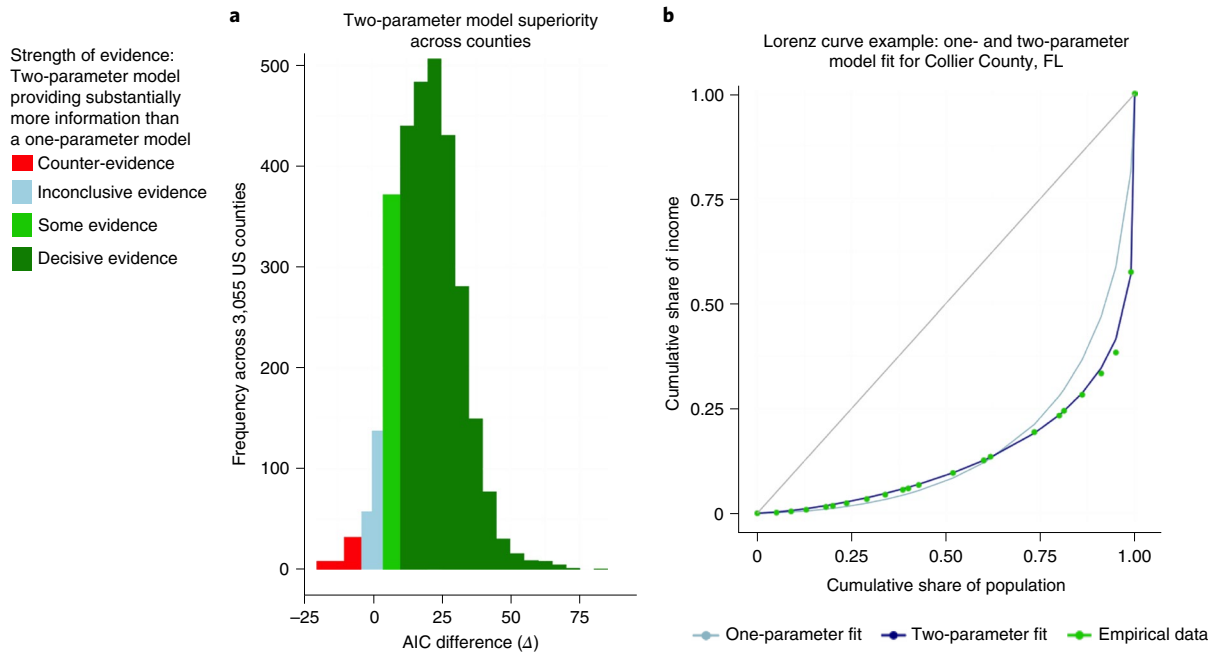
Number of parameters	Model	Borda score
2	Ortega	42,597
3	GB2	41,906
2	Dagum	38,791
5	Wang	38,187
2	Singh-Maddala	36,274
3	Abdalla-Hassan	35,354
4	Sarabia	32,272
2	Rasche	32,131
1	Lognormal	24,749
2	Generalized gamma	23,178
3	GB1	22,852
1	Gamma	13,926
1	Weibull	11,400
1	Pareto	9,522
1	Rhode	7,296
1	Chotikapanich	4,071
1	Kakwani-Podder	1,110

In each county, the Lorenz curve models were scored according to the Borda count procedure. The model with the highest Borda score wins. Models modelling Lorenz curves with one parameter represent single-parameter inequality measures (for example, the Gini coefficient).

all counties. Our analyses reveal that multi-parameter Lorenz curve models outperform almost all single-parameter Lorenz curve models considered in our analysis when using the Akaike information criterion corrected for small sample sizes (AIC<sub>c</sub>) as a measure of goodness of fit; in addition, we find that the two-parameter Ortega model is the overall winner of the Borda count (Table 2). We conclude that the two-parameter Ortega model provides the best overall fit to capture the information contained in the income distributions across US counties in this fine-grained dataset.

**Strength of evidence.** While the Borda count is a mechanism that aggregates results in a way that provides an overall model winner across all counties, we are also interested in how strong the evidence in favour of the two-parameter Ortega model is. That is, we aim to quantify how much more information we can capture by using a two-parameter model instead of a single-parameter model using AIC<sub>c</sub> differences (Methods). Taken together, the Borda count and AIC<sub>c</sub> difference analysis function as complementary building blocks in evaluating whether a two-parameter model performs well across counties while providing substantially more information than a one-parameter model within counties. On the basis of the finding that the two-parameter Ortega model won in the voting procedure, we are particularly interested in using the AIC<sub>c</sub> to determine whether the two-parameter Ortega model provides more relevant information about the income distribution than single-parameter Lorenz curve models, which function as representatives of single-parameter measures such as the Gini coefficient. We therefore compare AIC<sub>c</sub> values of the Ortega Lorenz curve model with the best-performing single-parameter Lorenz curve model in the Borda count contest (that is, the lognormal Lorenz curve model). We subsequently expand this analysis and also compare the Ortega model with the higher-parameter GB2 and Wang models, which were the closest runners-up in our analyses.

Figure 2a depicts the frequency of  $\Delta_{\text{lognormal,Ortega}}$  values across counties. As this figure shows, for the vast majority of cases, the lognormal single-parameter model (which is reflective of the



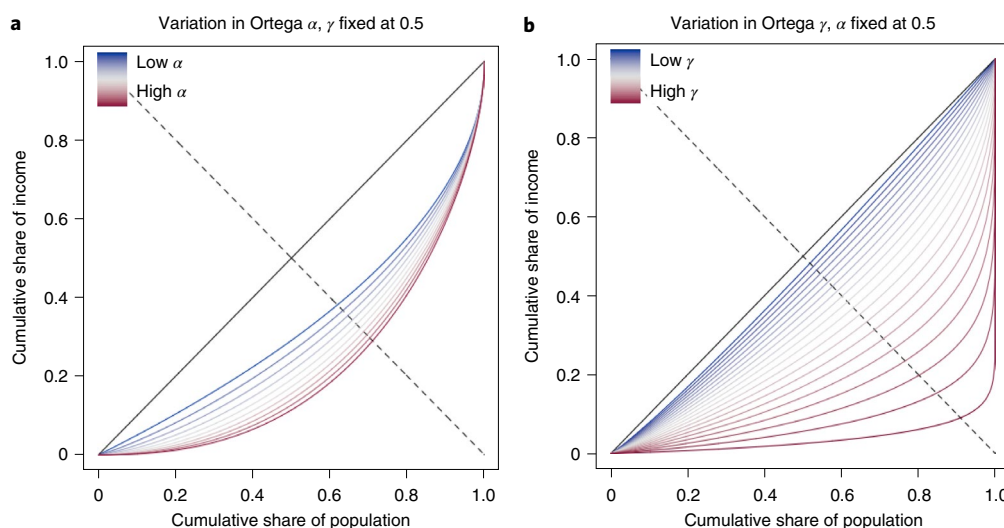
**Fig. 2 | The strength of evidence in favour of the two-parameter Ortega model.** **a**, The histogram plots the  $AIC_c$  differences ( $\Delta_{ij}$ ) between the one-parameter lognormal model ( $i$ ) and the two-parameter Ortega ( $j$ ). To categorize the strength of evidence, we define the following ranges:  $\Delta_{ij} > 10$  implies decisive evidence that model  $j$  is superior to model  $i$ ;  $\Delta_{ij} \in [4, 10]$  implies some evidence;  $\Delta_{ij} \in [-4, 4]$  implies inconclusive evidence; and  $\Delta_{ij} < -4$  implies counter-evidence (that is, evidence in favour of model  $i$  over  $j$ ). **b**, An example to illustrate the goodness of fit of one-parameter versus two-parameter models of Lorenz curves to empirical data. For the two-parameter model, we fitted the Ortega Lorenz curve model using the empirical data points and MLE, plotted next to the empirically best-fitting one-parameter model (the lognormal Lorenz curve model).

Gini coefficient) captures substantially less information than the two-parameter Ortega model. Put differently, we find decisive evidence that the two-parameter Ortega model captures substantially more information on the actual distribution of income in 80% of all US counties, providing further evidence that a two-parameter Lorenz curve model outperforms single-parameter models. For an illustration of how well the Ortega model fits the empirical data relative to the single-parameter model, see Fig. 2b.

We also use the  $AIC_c$  differences analysis to evaluate how two runners-up in our prior analysis, the higher-parameter GB2 and Wang models, performed compared with the model winner, Ortega. Comparing the three-parameter GB2 model with the two-parameter Ortega model, we find inconclusive evidence for whether one model outperforms the other (that is, in 2,413 of 3,056 total US counties, the absolute value of the  $AIC_c$  difference is below 4; Supplementary Fig. 12). While this is no indication of the two-parameter Ortega model performing better, we favour the two-parameter Ortega model for its simplicity (that is, two parameters are easier to interpret than three). In an  $AIC_c$  comparison of the five-parameter Wang model and the two-parameter Ortega model, we find that the Wang model outperforms the Ortega model for some counties but find the opposite for other counties (Supplementary Fig. 13). A closer look at the results reveals that the Ortega model more consistently ranks among the top-performing models, whereas the Wang model shows great performance in some counties and only mediocre performance in others (that is, the Wang model wins in plurality voting (Supplementary Fig. 5) but does not maintain a leading position in the Borda count (Supplementary Fig. 6)). Because our stated aim is to find a model that performs well across all counties, the two-parameter Ortega model is our preferred choice (details on the analysis and relevant figures are given in Supplementary Section 8). That said, other scholars may benefit from using different success criteria in choosing which model to use.

**Robustness analyses.** To evaluate the reliability of our results, we tested the robustness of estimates across estimation methods and goodness-of-fit measures. We estimate Lorenz curves through a nonlinear least squares (NLS) approach (Supplementary Section 9) and compare the NLS results with those obtained by MLE, ruling out the possibility that our results are driven by our choice of estimation procedure. Our analyses reveal only small relative differences between MLE and NLS estimates; for example, the median relative difference between MLE and NLS estimates for Ortega parameter  $\alpha$  across US counties is 0.0234 (see Supplementary Section 10 for more details). Additionally, we ran a simulation study to investigate the ability of the  $AIC_c$  to detect the true data-generating model when only a few empirical observations are available (Supplementary Section 5). We find a high true-model detection rate for our sample size of 19–23 data points per county; that is, if Ortega were the true data-generating model and 19 data points were available, we would on average correctly detect Ortega as the true model in 98.5% of all cases (Supplementary Table 4). This provides additional confidence in the reliability of the  $AIC_c$  given our specific setting. To rule out the possibility that our results are influenced by the choice of information criterion itself, we also conducted analyses with different information criteria. For example, we replicated our analyses using the Bayesian information criterion (BIC) instead of the  $AIC_c$  to check whether Borda voting results that determine the winning model are robust to other measures of performance and found similar results (Supplementary Section 6, in particular Supplementary Fig. 7). Furthermore, we conducted an analysis of BIC differences instead of  $AIC_c$  differences and found that the results in favour of the two-parameter Ortega model are robust across information criteria (Supplementary Section 7).

In sum, our robustness checks demonstrate the reliability of our results, suggesting that the two-parameter Ortega model does consistently well across all additional analyses in our dataset. Because



**Fig. 3 | Using simulations to systematically vary the two Ortega parameters to identify their impacts on the shape of the income distribution.** **a**, The disproportionate change exhibited by the Lorenz curve when the Ortega parameter  $\alpha$  varies within the range of 0.01 to 1.5 leads to a more pronounced change for lower income percentiles. (The dashed off-diagonal line facilitates the recognition that the Lorenz curve is stretched more intensely in lower income percentiles.) **b**, Conversely, when the Ortega parameter  $\gamma$  varies within the range 0.01 to 0.99, the Lorenz curve exhibits a disproportionate change in the top income percentiles. For comparison, the empirical estimates across counties for  $\alpha$  range from 0.12 to 1.23; for  $\gamma$ , they range from 0.3 to 0.93.

we are proposing a data-driven approach to studying inequality, our finding that the Ortega parameters are a close approximation to the real data critically depend on the dataset used. Note that the methodology we propose might yield different inequality measures in other datasets. This has implications for future researchers: we encourage scholars to use and apply our methods, not the resulting measures we find here, to income distributions in other settings, including in different countries around the world.

**Ortega parameters.** Parametric Lorenz curve models characterize the shape of income distributions using their parameters. As a result, those parameters themselves can be used as inequality measures. Because our two-parameter Ortega model emerged as the best-fitting model in our previous analysis, we now turn to investigating the characteristics of these two parameters as measures of inequality in more depth (Methods). To provide better insight into the role of both Ortega parameters in capturing the income distribution, we simulate a number of Ortega-type Lorenz curves. We vary one Ortega parameter while keeping the other fixed and visually evaluate the changes in the Lorenz curves' shapes, thereby examining how each Ortega parameter individually affects the Lorenz curve. Note that, as described in our Methods, we use a transformation of the second Ortega parameter  $\beta$  to  $\gamma$ , which reflects  $1-\beta$ .

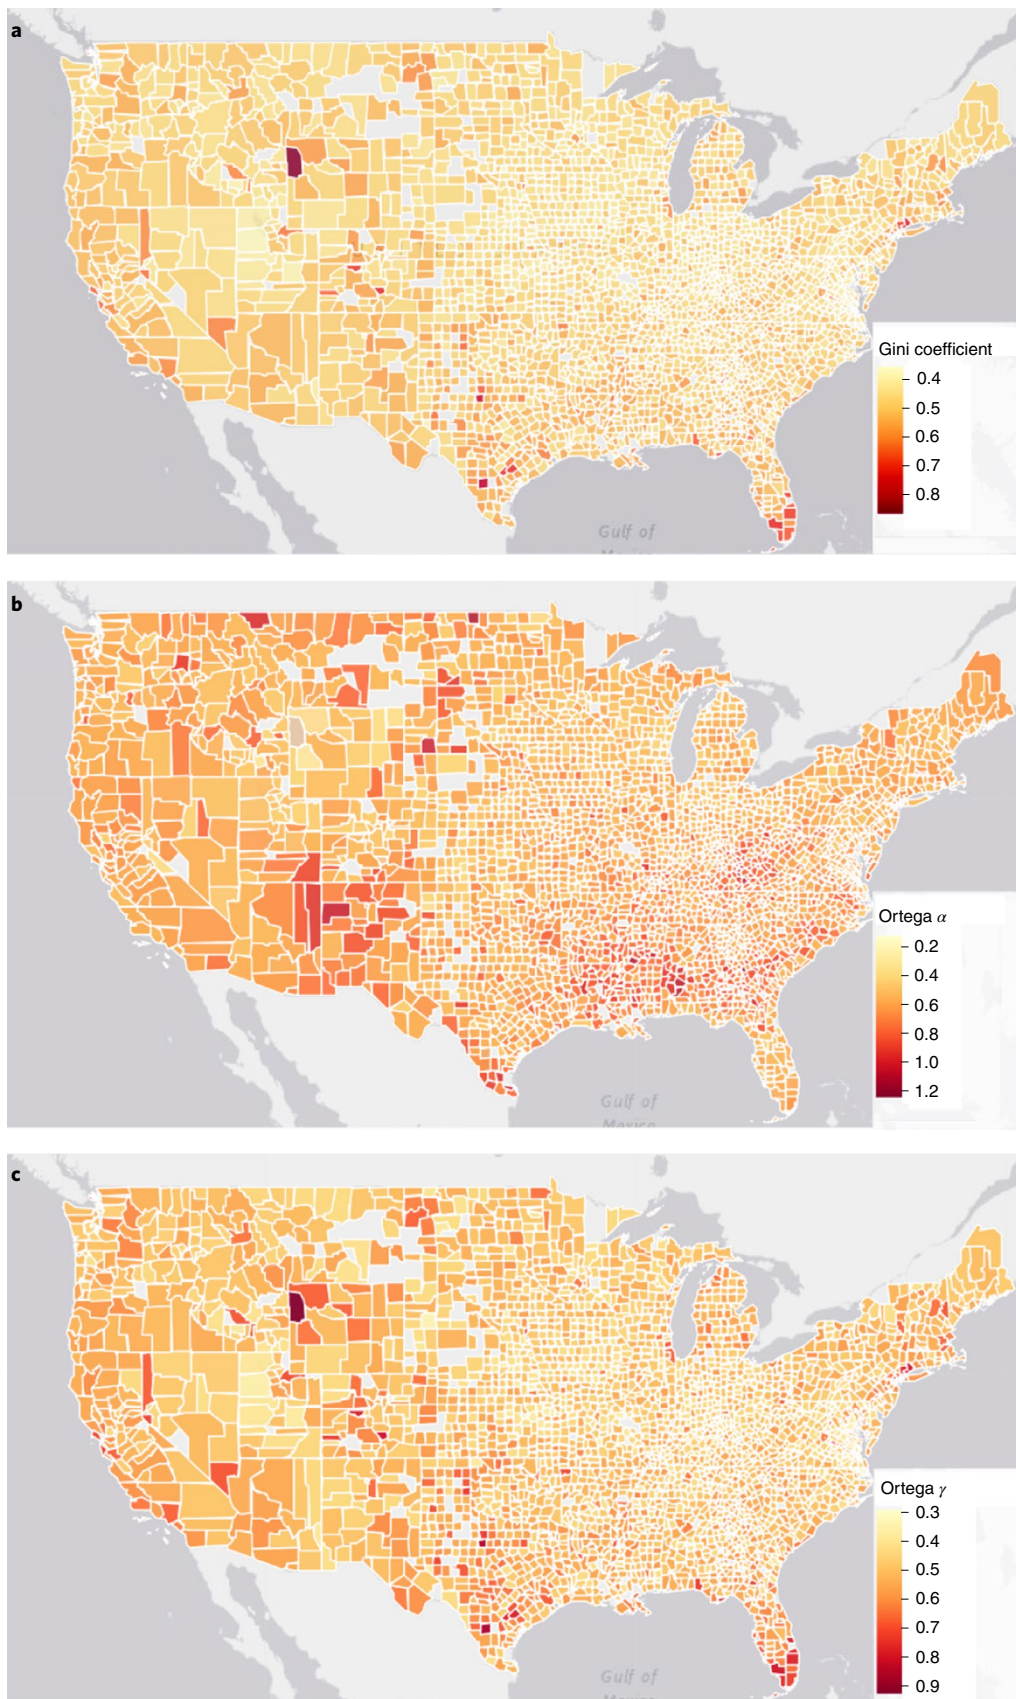
Contrasting the two parameters reveals that the first Ortega parameter,  $\alpha$ , captures inequality with a more pronounced focus on concentrations at the bottom of the income distribution, while the second Ortega parameter,  $\gamma$ , reflects an emphasis towards inequality concentrated at the very top of the income distribution (Fig. 3). Specifically, when  $\gamma$  is held constant,  $\alpha$  stretches the Lorenz curve on the left side of the income distribution (that is, at lower incomes; Fig. 3a), while variation in  $\gamma$  with a constant  $\alpha$  influences the shape of the Lorenz curve on the far right side (that is, at the highest income levels; Fig. 3b). This indicates that the parameters focus on different spectra of the Lorenz curve. That is, while the parameters combined capture the degree of inequality overall, each individually reflects a focus on a different part of the income distribution. This interpretation is in line with the parameters being correlated and affected by each other's alternating values, yet individually providing additional valuable information about the relative extent

of bottom- or top-concentrated inequality. Interested readers may find the interactive R shiny tool we created useful, which displays changing  $\alpha$  and  $\gamma$  parameters to better understand the effects of each parameter (available at <http://www.measuringinequality.com/>).

**Relationships with other inequality measures.** To provide another interpretation of the two Ortega parameters, we correlate them to simulated income ratios (see Supplementary Section 12 for the details). These analyses reveal a high partial correlation between  $\gamma$  and the 99/90 ratio ( $r=0.9088$ ) and between  $\alpha$  and the 90/10 ratio ( $r=0.9081$ ). That is, we can think of the Ortega parameters as shifting the line of differentiation between top and bottom inequality away from the median (that is, the 50th percentile) towards a higher percentile (for example, the 90th percentile). This is visually reflected in Fig. 3, which reveals a larger impact of Ortega  $\gamma$  on the very top income percentiles, whereas  $\alpha$  moderately bends the Lorenz curve within the lower income percentiles.

Combining the two Ortega parameters provides the degree of overall inequality. Analytically, we can calculate the Gini coefficient from an Ortega Lorenz curve (using the original notation of the Ortega parameters  $\alpha, \beta$ :  $\text{Gini}(\alpha, \beta) = \frac{\alpha-1}{\alpha+1} + 2B(\alpha+1, \beta+1)$ , where  $B()$  is the beta function<sup>26</sup>). This also implies that we cannot view one Ortega parameter alone as representative of the Gini coefficient and the other one as providing 'additional' information. Instead, the Ortega parameters individually allow us to differentiate where in the income distribution inequality is concentrated, and considering them jointly yields estimates of the overall level of inequality. Using both Ortega parameters allows us to distinguish between different sources of inequality, which the Gini coefficient cannot do because it condenses the same amount of information into a single value. To gain a better understanding what information is gained from using the Ortega parameters over the Gini coefficient, we have compiled Fig. 4, which depicts the values for both across the United States at the county level.

We conducted a number of additional analyses. First, we calculated derivatives, finding that the Ortega parameters have different rates of change depending on the section of the  $x$  axis (that is, the population share), with  $\gamma$  affecting the top percentile of the Lorenz curve most intensely (see Supplementary Section 12 for further



**Fig. 4 | Different representations of inequality across counties in the United States. a**, The Gini coefficient. **b**, The first Ortega parameter,  $\alpha$  (a measure of more bottom-concentrated inequality). **c**, The second Ortega parameter,  $\gamma$  (a measure of more top-concentrated inequality). An interactive version of this figure is available at [www.measuringinequality.com](http://www.measuringinequality.com). Sources: Esri, HERE, Garmin, © OpenStreetMap contributors, and the GIS User Community.

details, and in particular Supplementary Fig. 15). Additionally, the derivatives of the percentile ratios 90/50 and 50/10 calculated from the Ortega model showed that the functions are heavily influenced by a change of  $\gamma$  for the 90/50 ratio and  $\alpha$  for the 50/10 ratio. However, note that percentile ratios do not fully correspond to the Ortega parameters, suggesting that Ortega parameters  $\alpha$  and  $\gamma$  provide information similar to the percentile ratios plus additional valuable information. More specifically, note that the two Ortega parameters characterize the whole income distribution Lorenz curve, whereas percentile ratios give only point-wise information on how the underlying Lorenz curve behaves at certain points in the income distribution. As a result, redrawing the income distribution, especially when only a single percentile ratio is available, may still result in widely varying Lorenz curves (and therefore lead to concerns similar to those pertaining to single-parameter measures such as the Gini coefficient). An illustration of this is provided in Supplementary Fig. 14.

We also compared the Gini coefficients implied by the model parameters with those Gini coefficients calculated non-parametrically on the US county data (Supplementary Section 13 and Supplementary Fig. 16). This analysis demonstrates that one-parameter models substantially deviate from the ideal average deviation of zero, whereas two-parameter models are a major improvement (for example the one-parameter Pareto model implied Gini coefficient has a median deviation from the empirical Gini of  $-0.076$ , whereas the two-parameter Ortega model Gini yields a median deviation of  $0.004$ ). Across the two-parameter models, the Ortega model is the one closest to the deviation of zero with a substantial number of data points. With more parameters, precision further increases, but the improvements are much smaller than between the one- and two-parameter models.

**Policy-relevant outcomes.** Finally, given that we found the two-parameter Ortega model to aptly reflect the real-world income distributions in our dataset, we now turn to investigating whether the parameters of this model, used as inequality measures, are able to disentangle prior mixed findings on correlates of inequality. In so doing, we follow an established literature that correlates inequality measures—typically the Gini coefficient—with policy-relevant outcomes<sup>27,28</sup>. As is important to do in this established literature, we note the limitations of a correlational approach in these settings, such as the lack of causal claims and the need for theoretically derived predictions about the existence of any such relationships. Our goal is not to speak to any particular policy outcome but instead to illustrate how this approach can allow for more theoretically driven inquiry in the future. To do so, we calculate the correlations of the two Ortega parameters with a large number of policy-relevant variables at the county level intended to capture many different fields across the social sciences, and we compare them with the correlations between the Gini coefficient and those same variables.

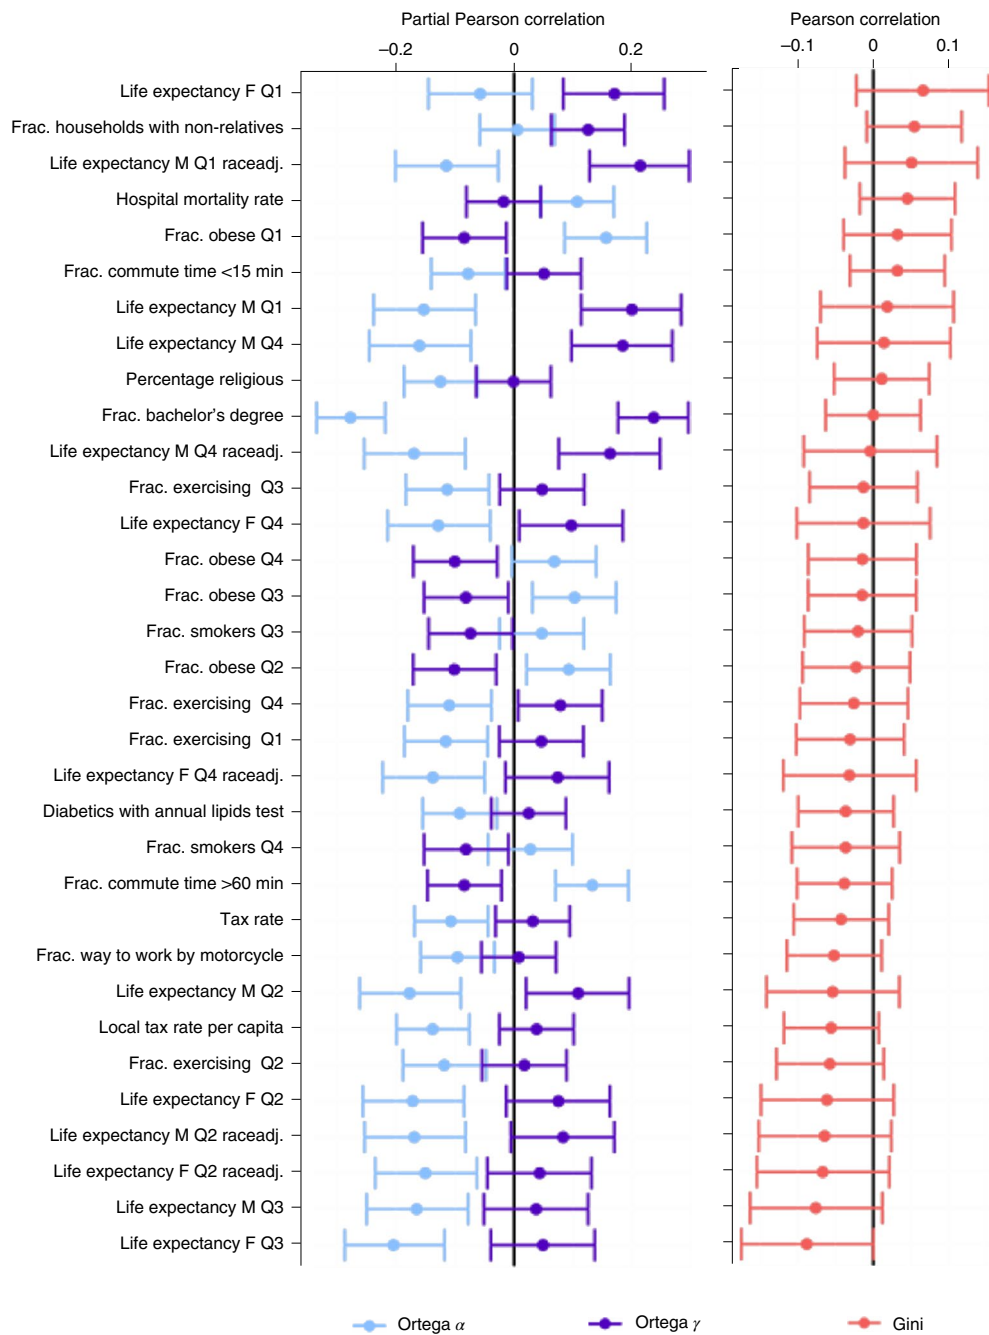
Our approach is exploratory and compares the use of two Ortega parameters with that of the Gini coefficient. Specifically, we investigate whether the two Ortega parameters detect statistically significant correlations that the Gini coefficient misses (that is, where the Gini coefficient does not have a statistically significant association). In other words, we might see cases where a county-level variable is not significantly correlated with the Gini coefficient but where there might be a statistically significant correlation with one (or two) of the Ortega parameters. In such cases, the two Ortega parameters may disentangle the effects of inequality associated with a certain spectrum of the income distribution that may be masked by the Gini coefficient. This may also apply to cases where we find a statistically significant correlation between the Gini coefficient and a county-level variable, and where one or both of the Ortega parameters are also significantly correlated. In such cases, our analyses would provide a clearer picture of the driver of the correlation between inequality

and that policy-relevant outcome—that is, whether that association stems from inequality concentrated at the top or the bottom of the income distribution, or both (see Supplementary Section 14 for the details). We emphasize that this is an exploratory analysis, and should not be viewed as a definitive determination that inequality is linked with a particular policy outcome.

To conduct this exploratory correlational study, we surveyed publicly available datasets, yielding 100 variables in the fields of health, crime, socio-economic status, social mobility and urban structures, which we then correlate with the inequality measures. We draw on various data sources, including the American Community Survey (ACS) 2011–2015<sup>29</sup>, aggregated tax records, Social Security Administration data<sup>30</sup> and a combined census, tax records and IRS Statistics of Income dataset<sup>31</sup>. Using these datasets, we calculate Pearson correlations between the inequality measures and county-level policy-relevant variables. Note that because we have 100 variables with which to correlate the inequality measures, we use a Bonferroni correction throughout the analysis, adjusting the  $\alpha=0.05$  significance level to  $\alpha=0.0005$  to account for multiple hypothesis testing<sup>32</sup>. Note that these are two-tailed tests. Since the Ortega model is a two-parameter model and should thus be interpreted jointly, we control for one Ortega parameter while correlating the other Ortega parameter with the county-level characteristics, and vice versa. For each Ortega parameter, we therefore calculated partial Pearson correlations, whereas we calculated simple Pearson correlations for the Gini coefficient.

Our analysis reveals that in 33 of 100 cases, at least one of the Ortega parameters was able to detect a statistically significant correlation (after applying the Bonferroni correction), but the Gini coefficient failed to do so. Figure 5 illustrates this subsample of cases, which include, among other things, cases of obesity, commuting time and the fraction of people with a bachelor's degree. For a further 59 cases, the Gini coefficient had a statistically significant correlation, as did at least one of the Ortega parameters, shedding light on whether concentrations of income at the bottom or top of the income distribution are driving this correlation (see Supplementary Table 11 for a general overview of our analysis results and Supplementary Section 14 for more details).

**Examples of applications to policy.** We highlight three examples that result from this analysis to illustrate how this approach can provide insights. Consider the association between economic inequality and obesity: prior research has found that the relationship between economic inequality and obesity is inconsistent, at times finding a positive relationship<sup>4</sup> and at times a negative relationship<sup>5</sup>. Our dataset provides detailed information about the percentage of people within any given county that have a body mass index classified as obese (body mass index 30+). Note that the Gini coefficient has no statistically significant correlation with obesity in our data, either aggregated across income levels (Bonferroni-corrected 99.95% confidence interval,  $(-0.088, 0.055)$ ) or separated by income quartile (Fig. 5). Both Ortega parameters, however, show statistically significant partial correlations in opposite directions (Bonferroni-corrected 99.95% confidence interval for  $\alpha$ ,  $(0.170, 0.306)$ ; for  $\gamma$ ,  $(-0.298, -0.161)$ ). Recall that a higher  $\alpha$  reflects a more pronounced bottom-concentrated inequality and a higher  $\gamma$  denotes higher inequalities at the very top of the income distribution. Our analysis reveals opposite effects for bottom- and top-concentrated inequality, such that greater bottom-concentrated inequality is associated with more obesity and higher top-concentrated inequality is associated with less obesity. The Gini coefficient, in contrast, fails to differentiate those diverging effects and finds a null association. Using the two Ortega parameters, we can differentiate between the *opposing* effects driving the relationship between inequality and obesity, with both theoretical and empirical implications for researchers and policymakers alike.



**Fig. 5 | A two-parameter Ortega approach reveals significant correlations between inequality and policy outcomes across  $N = 3,049$  US counties that the Gini coefficient misses in our dataset.** Point estimates of the Pearson correlations (Gini coefficient) and partial Pearson correlations (Ortega parameters) with policy outcomes are visualized with the bounds of the 0.9995 confidence interval, using a Bonferroni correction. The figure shows the subsample of covariates (33 of 100) for which the Pearson correlations with the Gini coefficient were not significant but that exhibited at least one statistically significant partial correlation with the Ortega parameters. M, male; F, female; Q, income quartile; frac., fraction; raceadj., race adjusted.

The correlation between economic inequality and educational outcomes provides a second example of the utility of our approach. Consider that a prior meta-analysis<sup>33</sup> found a wide range of results for the relationship between educational outcomes and economic inequality, both positive and negative. In our analysis, we find that the relationship between the Gini coefficient and an educational outcome such as the share of the population holding a bachelor's degree is not statistically significant but that both Ortega parameters show statistically significant associations in opposite directions (Fig. 5). More specifically, we find that higher bottom-concentrated inequality

is associated with a lower share of bachelor's degrees in the population and that higher top-concentrated inequality is associated with a greater share of bachelor's degrees. Viewed through this lens, a focus on the Gini coefficient alone obscures that educational outcomes are related to inequality—but in *opposing* ways for inequality concentrations at the bottom and top of the income distribution. Both examples highlight that a single inequality measure such as the Gini coefficient may mask effects that are revealed by the two Ortega parameters.

Finally, the two Ortega parameters may also clarify a relationship between inequality and its correlates even in cases where the



relationship between the Gini coefficient and correlates is statistically significant. For example, consider the association between economic inequality and the fraction of the population receiving social security income. In this case, the Gini coefficient and one of two Ortega parameters,  $\alpha$ , are significantly and positively correlated with the fraction of social security income recipients (Supplementary Fig. 21). In other words, a higher level of Ortega parameter  $\alpha$ , suggesting greater bottom-concentrated inequality, is associated with a higher percentage of social security income recipients. Although the Gini coefficient was positively correlated with the percentage of social security recipients as well, by using the Ortega parameters, we can see that this relationship is driven primarily by bottom-concentrated inequality.

## Discussion

Our jointly theoretically derived and data-driven analysis shows that single-parameter measures of inequality such as the widely used Gini coefficient may miss crucial information contained in income distributions. The two-parameter Ortega model, which we found shows a superior fit in our dataset of US county-level income distributions, reveals where inequality is concentrated along the income distribution. This information could enable researchers to generate and evaluate substantially more refined theories that relate economic inequality to social, political or psychological phenomena. That is, future theorizing may need to move beyond considering total levels of inequality to instead account for different inequality concentrations across the income distribution. It is likely that individuals psychologically experience inequality concentrated among low-income individuals (that is, a relatively larger gap between lower-income individuals and the rest of the population) very differently from inequality concentrated among high-income individuals (that is, a relatively larger gap between higher-income individuals and the rest of the population). For example, prior research has often found that individuals misperceive levels of inequality<sup>34,35</sup>; the approach detailed here may shed light on whether people perceive certain kinds of inequality more accurately than others, such as whether they are more accurate in estimating inequality concentrated among lower- than among higher-income individuals<sup>36,37</sup>. It is also likely that different concentrations of inequality are associated with distinct structural factors that may either impede or support the progress of different populations. More broadly, our exploratory correlational study provides tentative initial evidence for the variety of ways in which correlates of inequality may be empirically disentangled using multi-parameter measures of inequality, highlighting the need for future theory to develop a better understanding of why inequality concentrations that are more pronounced in a certain region along the income distribution may produce disparate effects. To aid in these endeavours, we are making our datasets and methodology—including Ortega parameter estimates for 3,056 US counties and 50 US states—publicly available for other researchers to use at [www.measuringinequality.com](http://www.measuringinequality.com).

Across academic, policy and public spheres, inequality has received growing attention in recent years. For example, a recent survey<sup>38</sup> suggested that a majority of Americans think there is too much economic inequality. At the same time, public support for measures to redress inequality depends on a variety of factors<sup>39</sup>. Our results highlight that one way to understand the diverging beliefs about inequality and preferences for redistribution is to focus on what kind of inequality respondents were dissatisfied with the most. This becomes clearer when discussing potential measures taken to redress inequality. For example, reducing top-concentrated economic inequality could be achieved by raising top income taxes, and reducing bottom-concentrated inequality may involve raising the minimum wage. Our approach and findings suggest that moving beyond the overall concentration of inequality as reflected in the Gini coefficient may be fruitful in pinpointing both how different

kinds of inequality affect outcomes and how to make meaningful change to redress inequality.

One limitation of our research is that our results are restricted to a specific dataset of counties in the United States, and our insights may not generalize to other datasets including in countries. To more broadly generalize beyond the current research, similar high-quality data including from other countries needs to be made publicly available by statistics bureaus. Most datasets that are available to researchers do not contain sufficient information to conduct the kinds of analyses we have demonstrated here. We hope that our work encourages statistics bureaus to publish more detailed inequality data and that they take note of the kind of information that should be included in publicly available data to ensure maximum usability and information content. This may include data on additional inequality measures, including the two Ortega parameters, as well as additional and fine-grained information on income distributions that would allow subsequent research to build on and extend the current research. In additional exploratory simulations reported in the Supplementary Information—which we suggest should be interpreted with caution—we outline three criteria that datasets on income distributions used for the method we detail here should meet: (1) data granularity, with at least 15 data points per Lorenz curve (Supplementary Section 5); (2) at least two data points of top income shares above the 90th percentile (Supplementary Section 15); and (3) at least 60 Lorenz curves (and ideally, many more; Supplementary Section 15). Currently, publicly available information on income distributions is far more limited and commonly falls short of satisfying all three criteria. For example, the World Bank database<sup>40</sup> only has data available on income quintiles as well as the top and bottom 10% (that is, a total of seven data points).

To fully take advantage of our research, we highlight that it is important for both additional inequality measures and for more fine-grained inequality distribution data to be made publicly available: while the two-parameter Ortega model was the best-fitting model in our dataset (which uniquely meets these three criteria), it is possible that in other datasets (including in other countries), a different model might outperform the Ortega model. Alternating model winners that provide the best fit to the data at hand might depend on the amount of available data (that is, how many data points are available might affect whether higher- or lower-parameter models best represent the data) and the actual shapes of different income distributions at different levels of analysis and in different areas of the world. Our research provides both a toolbox and an impetus for future work to move beyond single-parameter measures of inequality, which can be readily adapted as more granular inequality data become available. To better understand inequality and its correlates may require us to move beyond the widely available Gini coefficient.

## Methods

**Modelling the distribution of income.** As is the case for many constructs in the social sciences, there is no self-defining concept of inequality<sup>41</sup>, which leads to definitions of inequality being highly dependent on normative judgements<sup>42</sup>. To conduct research that does not impose normative judgements, we follow the etymological definition of income inequality—that is, the non-equal distribution of income. Through this lens, the measurement of inequality necessitates capturing the shape and form of income distributions. We use a parametric model that allows us to attain a “multidimensional view of the level of inequality which you can't get from a summary statistic directly” (ref. <sup>43</sup>, p. 196). Through a parametric model, we can subsequently redraw the income distribution when given only its parameters and compare it with the actual income distribution.

We also considered using non-parametric approaches—that is, methods that do not require parametric assumptions at any step—in our analyses. However, when evaluating non-parametric inequality measures, such as generalized entropy measures, we face one major disadvantage that renders them ill-fitting given the goals of our analysis: non-parametric summary statistics do not allow a comparison of their output with a ‘real’ income distribution (that is, the empirical data), which would enable us to ascertain the extent to which the measure is a good or bad approximation of actual data. And although there are some non-parametric

procedures available to model the distribution of income, a recent study finds that these methods “fail to represent income distributions accurately” (ref. <sup>44</sup>, p. 964). We therefore rely only on parametric approaches in our analysis.

**Lorenz curves.** The well-established Lorenz curve framework is helpful for modelling income distributions parametrically for the purpose of measuring inequality. The Lorenz curve is a graphical representation that visualizes the distribution of economic quantities across the population on a relative scale instead of using absolute terms. That is, the Lorenz curve displays which part of the population contributes what share to the total income of a whole population. To calculate the relative quantities for the distribution of income, the population is ordered from lowest- to highest-income individuals (or income groups), and then the share of total income held by the respective proportion of the population is determined. Subsequently, the proportions of total income are cumulated ( $y$  axis) and plotted against the cumulative share of the low- to high-income ordered population ( $x$  axis). The resulting curve shows where along the income distribution what share of total income is held.

In prior literature, Lorenz curve models originated from two distinct streams of research. One approach begins with a suggested statistical distribution of income and derives the respective Lorenz curve. For a random variable  $x$  representing the income of a member of the population with cumulative distribution  $F(x)$ , we can use the following formula given by ref. <sup>45</sup>: let  $F^{-1}(t) = \inf\{x : F(x) \geq t\}$  be the inverse of  $F(x)$  (that is, the quantile function), and let  $\mu = \int x dF(x)$  be the finite mean; then, the Lorenz curve is defined as  $L(u) = \mu^{-1} \int_0^u F^{-1}(t) dt$ ,  $0 \leq \mu \leq 1$ . A second stream of research proposes functional forms to satisfy relevant properties that qualify them as Lorenz curves. These properties are inspired by the real-world implications that a Lorenz curve model should have—for example, being bounded between zero and one, such that 0% of the population has 0% of the total income and 100% of the population possesses the total income. For a complete list of properties that need to be satisfied to qualify for a Lorenz curve, see refs. <sup>18,46–48</sup>.

Our study bridges the two approaches, and a resulting comprehensive literature review of possible candidate models yields a total of 17 Lorenz curve models, which we subsequently test (Table 1; for more information, see Supplementary Section 1). These vary in the number of parameters they use, from one to five. Note that the single-parameter Lorenz curve models such as the lognormal Lorenz curve model can be directly transformed into the Gini coefficient<sup>49</sup>; however, we cannot include the Gini coefficient as a model itself in Table 1 because it is a statistic—that is, a function of the data but not a statistical model that aims at describing the underlying data-generating process. For multiple-parameter Lorenz curve models, the Gini coefficient can also be calculated through a combination of parameters<sup>26</sup>. In reviewing these competing models, we ask: how many and which parameters are necessary for Lorenz curve models to capture relevant information contained in income distributions?

To answer this question, we fitted the Lorenz curve models presented in Table 1 to each of the  $N=3,056$  empirical Lorenz curves we obtained by combining two sets of US income data. Note that our approach is more extensive than comparable prior studies such as Chotikapanich and Griffiths<sup>50</sup>, who compare parametric model estimates across only five Lorenz curves, or Paul and Shankar<sup>51</sup>, who compare the fit of single-parameter models on only ten Lorenz curves. In addition, through the systematic application of goodness-of-fit analyses that we introduce for our specific question at hand, we determine the theoretical Lorenz curve model that most adequately describes the empirical Lorenz curves. The model winner reflects how many and what kind of parameters are best suited to capture the income distribution as depicted by Lorenz curves in the current data.

**US county-level datasets.** To arrive at a large dataset of income distributions, we combine two distinct data sources. The first is the ACS 2011–2015<sup>29</sup>, collected by the US Census Bureau from a representative sample of the US population (see Supplementary Section 2 for details about the dataset and data-cleaning procedures). The ACS data are particularly detailed for lower- and medium-income groups. The variables of interest are the ACS yearly estimates over the five-year period for the share of income earned by population quintile and the top 5% of income earners, the income aggregate per county, and the count of people that fall into certain ranges of income (income buckets). Within income buckets, we assumed a symmetrical distribution of income, such that we can calculate the share of income held by the fraction of the population within the respective bucket and draw a Lorenz curve (see Supplementary Section 3 for more details on this procedure). As with most grouped-income data, the top income bucket is an open interval. In our case, the ACS provides the number of households that have an annual income  $>US\$200,000$ , but no information is provided on how people are distributed within that bucket. This makes accurate estimates of top income shares inaccessible; however, this information is particularly important for our purposes of accurately depicting real-world income distributions.

To overcome this shortcoming at top income levels, we enrich the ACS data with more precise estimates for top income groups through data from the Economic Policy Institute<sup>52</sup> that contains income shares for the bottom 90%, 90–95%, 95–99% and top 1% of income earners in the United States for the year 2013. The data for this table were constructed using tax data from the Internal

Revenue Service’s Statistics of Income Tax Stats and therefore provide more reliable information on high-income shares. The Economic Policy Institute dataset consists of data on 3,064 US counties for which the ACS also provides data. We excluded the District of Columbia because of its special nature and seven counties because of data inconsistencies. Our final sample of empirical real-world Lorenz curves at the US county level covers 3,056 counties. Of a total of 3,143 US counties and county equivalences, our dataset therefore achieves a coverage of 97% of all counties in the United States.

**Estimation and goodness-of-fit analysis.** For the estimation and goodness-of-fit analysis, we combine elements that are well known from applied statistics and that are particularly suitable given the current context. Following Chotikapanich and Griffiths<sup>50</sup>, who introduced MLE for Lorenz curve estimation, we estimate the Lorenz curve parameters by maximizing a log-likelihood function that originates from a Dirichlet distribution with newly defined parameters that incorporate the Lorenz curve parameters (the details can be found in Supplementary Section 4). The MLE framework allows us to use the AIC, which is defined as

$$AIC = -2 \times \ell(\hat{\theta}) + 2p$$

where  $p$  is the number of parameters of the model and  $\ell(\hat{\theta})$  is the value of the log-likelihood function at the maximum likelihood estimate of the parameter vector  $\theta$ . While we have a large number of counties for which the models are estimated independently, the number of data points available to construct the Lorenz curve for a certain county ranges from 19 to 23, which makes it reasonable to adjust for small sample sizes in our estimation. Drawing on refs. <sup>53,54</sup>, the bias-corrected version of the AIC for small sample sizes can be written as

$$AIC_c = AIC + \frac{2p(p+1)}{n-p+1}$$

We chose the AIC because it is well defined within the MLE framework and offers a useful evaluation criterion that balances complexity and model fit, whereby high-complexity models incur a penalty<sup>55</sup>. That is, the AIC helps us distinguish whether an additional parameter (that is, a more complicated Lorenz curve model) captures further relevant information, while ensuring that models that do well in approximating the empirical Lorenz curves are not unnecessarily complex. One can think of the AIC as a way to improve the bias–variance trade-off between models—that is, a high-parameter model might overfit the data (high variance across counties), while a low-parameter model might incorporate a large bias (see Supplementary Section 5 for further discussion). At the same time, there are various ways to penalize for the use of many parameters; we therefore also use the BIC, which uses a different penalty term than the AIC, to evaluate the robustness of our results (Supplementary Sections 6 and 7).

We next determine maximum likelihood parameter estimates and  $AIC_c$  values for each of the 17 Lorenz curve models considered in each of the  $N=3,056$  counties. The lower the  $AIC_c$ , the better, which allows us to rank the models: for each county, the Lorenz curve model with the lowest  $AIC_c$  is assigned to rank 1, the Lorenz curve model with the second-lowest  $AIC_c$  value is assigned to rank 2 and so on. We subsequently aggregate the ranks across the  $N=3,056$  counties and use a common voting procedure—predominantly used to aggregate preferences of individuals on a group level—to help determine  $AIC_c$  model preferences across all counties. Specifically, we use the Borda count (see ref. <sup>56</sup> for more details), which scores choices through the summation of points assigned according to their ranks. That is, if there are  $n$  options to choose from, the option ranking first receives  $n$  points, the option ranking second receives  $n-1$  points, ..., and the least favoured option receives 0 points. Those points are then summed across observations (in our case, across counties), and the option that receives the most points wins the Borda count. (For alternative voting procedures and a discussion of why the Borda count voting procedure is our preferred choice, see the additional analyses in Supplementary Section 6.)

**$AIC_c$  differences.** We analyse  $AIC_c$  differences that allow us to evaluate the extent to which the single-parameter model may miss information contained in income distributions compared with the two-parameter model. While the absolute  $AIC_c$  values themselves are not meaningful, because they contain arbitrary constants and are affected by sample size, differences between  $AIC_c$  values are free of such constants, as they affect all  $AIC_c$  values equally<sup>57</sup>. To calculate  $AIC_c$  differences, we generalize and extend prior work<sup>27</sup> by defining  $AIC_c$  differences as follows:

$$\Delta_{ij} = AIC_{c,i} - AIC_{c,j}$$

where  $AIC_{c,i}$  is the  $AIC_c$  value of the model  $i$  and  $AIC_{c,j}$  the  $AIC_c$  value of model  $j$ . Hence,  $\Delta_{ij}$  is the information loss experienced when fitting model  $i$  rather than model  $j$ . Information loss will thus act as a criterion for the strength of evidence for or against a model: if  $\Delta_{ij}$  is small, we do not lose much information when fitting model  $i$  instead of  $j$  to our data. In this case, there would be support (or evidence) for model  $j$  in capturing as much information as model  $i$ . The larger  $\Delta_{ij}$  gets, the less plausible it is for model  $i$  to be as good an approximation of the data as model  $j$ ; that is, the larger  $\Delta_{ij}$ , the more certain we are that model  $j$  provides a

substantially better model for our data. Using a conservative estimate<sup>57</sup>, we can define the following ranges:  $\Delta_{ij} > 10$  implies decisive evidence that model  $j$  is superior to model  $i$  in capturing relevant information from the empirical income distribution;  $\Delta_{ij} \in [4, 10]$  implies some evidence;  $\Delta_{ij} \in [-4, 4)$  implies inconclusive evidence; and  $\Delta_{ij} < -4$  implies counter-evidence (that is, evidence in favour of model  $i$  over  $j$ ).

**Ortega parameters.** We use the parameters of the Ortega model directly as measures of inequality. The parameters of a Lorenz curve model characterize the shape of the resulting Lorenz curve. In other words, we argue that key information from the income distribution can be condensed into parameters that act as measures of inequality. For the Ortega model, there are two parameters available for fitting to the data, and we aim to explore what kind of information each Ortega parameter captures. While Ortega et al.<sup>56</sup> did not detail the theoretical origins of the proposed functional form, others acknowledge that the Ortega Lorenz curve model coincides with a model inside the hierarchical family of Pareto Lorenz curves<sup>58</sup>. In particular, the Lorenz curve associated with the Pareto distribution takes the form

$$L(u) = 1 - (1 - u)^{1 - \frac{1}{a}}, \text{ where } a > 1$$

Applying a previously proposed generalization<sup>58</sup> such that  $L_2(u) = u^\alpha \times L_1(u)$  and defining  $\beta = 1 - \frac{1}{a}$  results in the Ortega Lorenz curve of the form:

$$L(u) = u^\alpha \times (1 - (1 - u)^\beta), \text{ where } 0 \leq \alpha; 0 < \beta \leq 1$$

There is therefore a close connection between the Pareto parameter  $a$  and the second Ortega parameter  $\beta$ . In fact, when the first Ortega parameter equals zero, an analytical solution for the first Ortega parameter and the Pareto parameter relationship can be found:  $\beta = 1 - \frac{1}{a}$  (for more technical details on the derivation of the relationship between the Pareto distribution and the second Ortega parameter  $\beta$ , see Supplementary Section 11).

Note that the Pareto parameter  $a$  has previously been used as a measure of inequality, more widely known as the Pareto index. Indeed, prior research has found that the Pareto index is particularly useful for modelling the upper tail of the income distribution<sup>59,60</sup>, denoting the frequency with which top incomes occur. More formally, this is described as the breadth of the Pareto distribution, corresponding to the shape parameter  $a$  within the Pareto distribution<sup>61</sup>. This means that the smaller the  $a$ , the thicker the right tail of the Pareto distribution<sup>59</sup>. One might therefore suspect that the lower the second Ortega parameter  $\beta$ , the more inequality is concentrated at the top of the income distribution—that is, that there are more occurrences of top incomes. To ease interpretation, we transform the Ortega parameter  $\beta$  as follows:

$$\gamma := 1 - \beta$$

The newly defined parameter  $\gamma$  now implies the more intuitive interpretation that a higher  $\gamma$  indicates a higher concentration of inequality at the top of the income distribution. Note that  $\gamma$  is bounded by  $0 \leq \gamma < 1$ . In contrast, for the first Ortega parameter,  $a$ , prior literature has not suggested an interpretation. We therefore turn to simulations to study  $a$  in more detail. These simulations reveal that an increase in  $a$  stretches the left side of the Lorenz curve towards the  $x$  axis (that is, at lower incomes), suggesting that  $a$  captures inequality that is more pronounced among the bottom and middle percentiles of the distribution (see Supplementary Section 12 for the details).

**Reporting summary.** Further information on research design is available in the Nature Research Reporting Summary linked to this article.

### Data availability

All data to reproduce the findings discussed in this paper are available at <http://www.measuringinequality.com/>.

### Code availability

All code to reproduce the findings discussed in this paper are available at <http://www.measuringinequality.com/>.

Received: 2 August 2021; Accepted: 13 July 2022;

Published online: 29 August 2022

### References

- Davies, J., Lluberas, R. & Shorrocks, A. in *Credit Suisse Wealth Report 1–157* (Credit Suisse, 2016); <https://www.credit-suisse.com/media/assets/corporate/docs/about-us/research/publications/global-wealth-databook-2016.pdf>
- Cornia, G. A. *Falling Inequality in Latin America: Policy Changes and Lessons* (Oxford Univ. Press, 2014).
- Wilkinson, R. & Pickett, K. *The Spirit Level: Why Greater Equality Makes Societies Stronger* (Bloomsbury, 2011).
- Pickett, K. E., Kelly, S., Brunner, E., Lobstein, T. & Wilkinson, R. G. Wider income gaps, wider waistbands? An ecological study of obesity and income inequality. *J. Epidemiol. Community Health* **59**, 670–674 (2005).
- Kim, D., Wang, F. & Arcan, C. Geographic association between income inequality and obesity among adults in New York State. *Prev. Chronic Dis.* **15**, E123 (2018).
- Ngamaba, K., Panagioti, M. & Armitage, C. Income inequality and subjective well-being: a systematic review and meta-analysis. *Qual. Life Res.* **27**, 577–596 (2018).
- Côté, S., House, J. & Willer, R. High economic inequality leads higher-income individuals to be less generous. *Proc. Natl Acad. Sci. USA* **112**, 15838–15843 (2015).
- Schmukle, S. C., Korndörfer, M. & Egloff, B. No evidence that economic inequality moderates the effect of income on generosity. *Proc. Natl Acad. Sci. USA* **116**, 9790–9795 (2019).
- De Maio, F. Income inequality measures. *J. Epidemiol. Community Health* **61**, 849–852 (2007).
- Tackling High Inequalities Creating Opportunities for All* (OECD, 2014); <https://www.oecd.org/unitedstates/Tackling-high-inequalities.pdf>
- Hearing document: Congress must act to reduce inequality for working families. In *Hearing before the Committee on the Budget House of Representatives* Congress Hearing No. 116-14 (116th Congress, 2019); [https://budget.house.gov/sites/democrats.budget.house.gov/files/document/Inequality%20Post%20Hearing%20Report\\_Final.pdf](https://budget.house.gov/sites/democrats.budget.house.gov/files/document/Inequality%20Post%20Hearing%20Report_Final.pdf)
- Charles-Coll, J. Understanding income inequality: concept, causes and measurement. *Int. J. Econ. Manage. Sci.* **1**, 17–28 (2011).
- Giorgi, G. & Glijarano, C. The Gini concentration index: a review of the inference literature. *J. Econ. Surv.* **31**, 1130–1148 (2016).
- Sithiyot, T. & Holasut, K. A simple method for measuring inequality. *Palgrave Commun.* **6**, 112 (2020).
- SI Index* (World Bank, 2021); <https://data.worldbank.org/indicator/SI.POV.GINI>
- Cowell, F. in *Handbook of Income Distribution* (eds Atkinson, A. & Bourguignon, F.), Ch. 2 (Elsevier, 2000).
- Liu, Y. & Gastwirth, J. L. On the capacity of the Gini index to represent income distributions. *METRON* **78**, 61–69 (2020).
- Fellman, J. Income inequality measures. *Theor. Econ. Lett.* **08**, 557–574 (2018).
- Clementi, F., Gallegati, M., Gianmoena, L., Landini, S. & Stiglitz, J. Mis-measurement of inequality: a critical reflection and new insights. *J. Econ. Interact. Coord.* **14**, 891–921 (2019).
- Atkinson, A. & Bourguignon, F. *Handbook of Income Distribution* 1st edn, Vol. 1 (Elsevier, 2000); <https://EconPapers.repec.org/RePEc:eee:income:1>
- Davies, J., Hoy, M. & Zhao, L. Revisiting comparisons of income inequality when Lorenz curves intersect. *Soc. Choice Welfare* **58**, 101–109 (2022).
- Davydov, Y. & Greselin, F. Comparisons between poorest and richest to measure inequality. *Sociol. Methods Res.* **49**, 526–561 (2020).
- Atkinson, A. B. On the measurement of inequality. *J. Econ. Theory* **2**, 244–263 (1970).
- Zanardi, G. Della asimmetria condizionata delle curve di concentrazione. lo scentrimento. *Riv. Ital. Econ. Demogr. Stat.* **18**, 431–466 (1964).
- Gwatkin, D. Health inequalities and the health of the poor: what do we know? What can we do? *Bull. World Health Organ.* **78**, 3–18 (2000).
- Ortega, P., Martín, G., Fernández, A., Ladoux, M. & García, A. A new functional form for estimating Lorenz curves. *Rev. Income Wealth* **37**, 447–452 (1991).
- De Dominicis, L., Florax, R. J. & De Groot, H. L. A meta-analysis on the relationship between income inequality and economic growth. *Scott. J. Polit. Econ.* **55**, 654–682 (2008).
- Kondo, N. et al. Income inequality, mortality, and self-rated health: meta-analysis of multilevel studies. *Br. Med. J.* **339**, 1178–1181 (2009).
- American Community Survey, 2011–2015: 5-Year Period Estimates* (US Census Bureau, 2016); <https://data2.nhgis.org/main>
- Chetty, R. et al. The association between income and life expectancy in the United States, 2001–2014. *JAMA* **315**, 1750–1766 (2016).
- Chetty, R. & Hendren, N. The impacts of neighborhoods on intergenerational mobility II: county-level estimates. *Q. J. Econ.* **133**, 1163–1228 (2018).
- Abdi, H. Bonferroni and Šidák corrections for multiple comparisons. *Encycl. Meas. Stat.* **3**, 103–107 (2007).
- Abdullah, A., Doucouliagos, H. & Manning, E. Does education reduce income inequality? A meta-regression analysis. *J. Econ. Surv.* **29**, 301–316 (2015).
- Hauser, O. P. & Norton, M. I. (Mis)perceptions of inequality. *Curr. Opin. Psychol.* **18**, 21–25 (2017).
- Knell, M. & Stix, H. Perceptions of inequality. *Eur. J. Polit. Econ.* **65**, S0176268020300756 (2020).
- Phillips, L. T. et al. Inequality in people's minds. Preprint at *PsyArXiv* <https://doi.org/10.31234/osf.io/vawh9> (2020).

37. Jachimowicz, J. M. et al. Inequality in researchers' minds: four guiding questions for studying subjective perceptions of economic inequality. *J. Econ. Surv.* <https://doi.org/10.1111/joes.12507> (2022).
38. *Most Americans Say There Is Too Much Economic Inequality in the U.S., but Fewer Than Half Call It a Top Priority* (Pew Research Center, 2020); [https://www.pewresearch.org/social-trends/wp-content/uploads/sites/3/2020/01/PSDT\\_01.09.20\\_economic-inequality\\_FULL.pdf](https://www.pewresearch.org/social-trends/wp-content/uploads/sites/3/2020/01/PSDT_01.09.20_economic-inequality_FULL.pdf)
39. Brown-Iannuzzi, J. L., Lundberg, K. B. & McKee, S. E. Economic inequality and socioeconomic ranking inform attitudes toward redistribution. *J. Exp. Soc. Psychol.* **96**, 104180 (2021).
40. *Income Share Held by Highest 20%* (World Bank, 2021); <https://data.worldbank.org/indicator/SI.DST.05TH.20>
41. Cowell, F. *Measuring Inequality* 3rd edn (Oxford Univ. Press, 2011); <https://EconPapers.repec.org/RePEc:oxp:obooks:9780199594047>
42. Campano, F. & Salvatore, D. *Income Distribution* (Oxford Univ. Press, 2006); <https://EconPapers.repec.org/RePEc:oxp:obooks:9780195300918>
43. Slottje, D. J. Using grouped data for constructing inequality indices: parametric vs. non-parametric methods. *Econ. Lett.* **32**, 193–197 (1990).
44. Jorda, V., Sarabia, J. M. & Jäntti, M. Inequality measurement with grouped data: parametric and non-parametric methods. *J. R. Stat. Soc. Ser. A* **184**, 964–984 (2021).
45. Gastwirth, J. L. A general definition of the Lorenz curve. *Econometrica* **39**, 1037–1039 (1971).
46. Krause, M. Parametric Lorenz curves and the modality of the income density function. *Rev. Income Wealth* **60**, 905–929 (2014).
47. Basmann, R., Hayes, K., Slottje, D. & Johnson, J. A general functional form for approximating the Lorenz curve. *J. Econ.* **43**, 77–90 (1990).
48. Kakwani, N. On a class of poverty measures. *Econometrica* **48**, 437–446 (1980).
49. Krause, M. Parametric Lorenz curves and the modality of the income density function. *Rev. Income Wealth* **60**, 905–929 (2014).
50. Chotikapanich, D. & Griffiths, W. E. Estimating Lorenz curves using a Dirichlet distribution. *J. Bus. Econ. Stat.* **20**, 290–295 (2002).
51. Paul, S. & Shankar, S. An alternative single parameter functional form for Lorenz curve. *Empir. Econ.* **59**, 1393–1402 (2020).
52. Sommeiller, E., Price, M. & Wazeter, E. *Income Inequality in the US by State, Metropolitan Area, and County* Tech. Rep. (EPI, 2016); <https://www.epi.org/publication/income-inequality-in-the-us/#epi-toc-1>
53. Sugiura, N. Further analysts of the data by Akaike's information criterion and the finite corrections. *Commun. Stat. Theory Methods* **7**, 13–26 (1978).
54. Hurvich, C. M. & Tsai, C.-L. Regression and time series model selection in small samples. *Biometrika* **76**, 297–307 (1989).
55. Akaike, H. A new look at the statistical model identification. *IEEE Trans. Autom. Control* **19**, 716–723 (1974).
56. Brams, S. J. & Fishburn, P. C. in *Handbook of Social Choice and Welfare* Vol. 1 (eds Arrow, K. J. et al.) 173–236 (Elsevier, 2002); <http://www.sciencedirect.com/science/article/pii/S157401100280008X>
57. Burnham, K. P. & Anderson, D. R. Multimodel inference: understanding AIC and BIC in model selection. *Sociol. Methods Res.* **33**, 261–304 (2004).
58. Sarabia, J. M., Castillo, E. & Slottje, D. An ordered family of Lorenz curves. *J. Econ.* **91**, 43–60 (1999).
59. Benhabib, J. & Bisin, A. Skewed wealth distributions: theory and empirics. *J. Econ. Lit.* **56**, 1261–1291 (2018).
60. Jenkins, S. P. Pareto models, top incomes and recent trends in UK income inequality. *Economica* **84**, 261–289 (2017).
61. Arnold, B. C. *Pareto Distribution* (American Cancer Society, 2015); <https://doi.org/10.1002/9781118445112.stat01100.pub2>
62. Kakwani, N. C. & Podder, N. On the estimation of Lorenz curves from grouped observations. *Int. Econ. Rev.* **14**, 278–292 (1973).
63. Rasche, R. H., Gaffney, J., Koo, A. Y. C. & Obst, N. Functional forms for estimating the Lorenz curve. *Econometrica* **48**, 1061–1062 (1980).
64. Chotikapanich, D. A comparison of alternative functional forms for the Lorenz curve. *Econ. Lett.* **41**, 129–138 (1993).
65. Abdalla, I. M. & Hassan, M. Y. Maximum likelihood estimation of Lorenz curves using alternative parametric model. *Metodoloski Zv.* **1**, 109–118 (2004).
66. Rohde, N. An alternative functional form for estimating the Lorenz curve. *Econ. Lett.* **105**, 61–63 (2009).
67. Wang, Z., Ng, Y.-K. & Smyth, R. A general method for creating Lorenz curves. *Rev. Income Wealth* **57**, 561–582 (2011).
68. Arnold, B. C. & Sarabia, J. M. *Majorization and the Lorenz Order with Applications in Applied Mathematics and Economics* 1st edn (Springer International, 2018).
69. Kleiber, C. & Kotz, S. A characterization of income distributions in terms of generalized Gini coefficients. *Soc. Choice Welfare* **19**, 789–794 (2002).
70. Dagum, C. Wealth distribution models: analysis and applications. *Statistica* **3**, 235–268 (2006).
71. McDonald, J. Some generalized functions for the size distribution of income. *Econometrica* **52**, 647–663 (1984).

## Acknowledgements

We thank S. Bhatia, S. Davidai, T. Graeber and J. Tan for helpful discussions and comments that substantially improved this paper; I. Zahn for technical support; and M. Kalisch for his advice on statistics. We also acknowledge funding from the German Academic Scholarship Foundation (to K.B.), Harvard Business School (to J.M.J.), University of Exeter Business School (to O.P.H.) and the UKRI Future Leaders Fellowship (to O.P.H.). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

## Author contributions

K.B. led the data collection and statistical analysis under the supervision of J.M.J. and O.P.H. All authors wrote and edited the paper.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41562-022-01430-7>.

**Correspondence and requests for materials** should be addressed to Kristin Blesch, Oliver P. Hauser or Jon M. Jachimowicz.

**Peer review information** *Nature Human Behaviour* thanks the anonymous reviewers for their contribution to the peer review of this work.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2022

## Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

### Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size ( $n$ ) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided  
*Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g.  $F$ ,  $t$ ,  $r$ ) with confidence intervals, effect sizes, degrees of freedom and  $P$  value noted  
*Give  $P$  values as exact values whenever suitable.*
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's  $d$ , Pearson's  $r$ ), indicating how they were calculated

*Our web collection on [statistics for biologists](#) contains articles on many of the points above.*

### Software and code

Policy information about [availability of computer code](#)

Data collection

Data analysis

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

### Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

## Human research participants

Policy information about [studies involving human research participants and Sex and Gender in Research](#).

Reporting on sex and gender	<input type="text" value="We did not collect original data but accessed existing data."/>
Population characteristics	<input type="text" value="See above."/>
Recruitment	<input type="text" value="See above."/>
Ethics oversight	<input type="text" value="Accessed publicly available data."/>

Note that full information on the approval of the study protocol must also be provided in the manuscript.

## Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences       Behavioural & social sciences       Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

## Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	<input type="text" value="Data on US county-level income distributions was accessed through publicly available information, and then compared against 17 different previously proposed models of inequality using a ranking procedure and a comparison based on AIC_c differences. Once the winning model was identified (the two-parameter Ortega), we used its parameters as measures of inequality and compared its correlations to 100 policy-relevant variables (also accessed from publicly available information) to correlations with the Gini coefficient."/>
Research sample	<input type="text" value="We did not collect original data but accessed existing data."/>
Sampling strategy	<input type="text" value="see above"/>
Data collection	<input type="text" value="see above"/>
Timing	<input type="text" value="see above"/>
Data exclusions	<input type="text" value="see above"/>
Non-participation	<input type="text" value="see above"/>
Randomization	<input type="text" value="n/a"/>

## Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

### Materials & experimental systems

n/a	Included in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology and archaeology
<input checked="" type="checkbox"/>	<input type="checkbox"/> Animals and other organisms
<input checked="" type="checkbox"/>	<input type="checkbox"/> Clinical data
<input checked="" type="checkbox"/>	<input type="checkbox"/> Dual use research of concern

### Methods

n/a	Included in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/> MRI-based neuroimaging

---

**Supplementary information**

---

**Measuring inequality beyond the Gini coefficient may clarify conflicting findings**

---

In the format provided by the authors and unedited

# 1 **Supplementary Information**

## 2 **Measuring Inequality Beyond the Gini Coefficient May Clarify Conflicting Findings**

3 **Kristin Blesch, Oliver P. Hauser, Jon M. Jachimowicz**

4 **E-Mail Addresses: [blesch@leibniz-bips.de](mailto:blesch@leibniz-bips.de), [o.hauser@exeter.ac.uk](mailto:o.hauser@exeter.ac.uk), [jjachimowicz@hbs.edu](mailto:jjachimowicz@hbs.edu)**

### 5 **This PDF file includes:**

- 6     Supplementary Figures 1 to 24
- 7     Supplementary Tables 1 to 11
- 8     SI References



9 **Notation Preface**

- 10 • Gamma function:  $\Gamma(\nu) = \int_0^\infty \exp^{-t} t^{\nu-1} dt$
- 11 • Lower incomplete gamma function ratio:  $G(x, \nu) = \int_0^x t^{\nu-1} \exp(-t) dt / \Gamma(\nu)$
- 12 • Lower incomplete beta function ratio:  $B(x; a, b) = \frac{\int_0^x t^{a-1} (1-t)^{b-1} dt}{\int_0^1 t^{a-1} (1-t)^{b-1} dt}$

13 **1. Functional forms of Lorenz curve models**

14 **Properties.** To ensure that the proposed functional form can serve as a Lorenz curve model, certain properties of Lorenz curves  
 15 should be satisfied. As described in (1-3), general properties of the Lorenz curve  $L$  with respect to the cumulative percentages  
 16 of the population  $p$  are the following:

- 17 1.  $L(u)$  is monotone increasing
- 18 2.  $L(u) \leq p$
- 19 3.  $L(u)$  is convex
- 20 4.  $L(0) = 0$  and  $L(1) = 1$

21 More formally, the following theorem (cited by (4, 5) but attributed to Pakes 1981) determines what functions qualify as  
 22 Lorenz curves:

23 **Theorem 1 (Lorenz curve)**

24 *A function  $L(u)$ , continuous on  $[0, 1]$  and with second derivative  $L''(u)$  is a Lorenz curve if and only if  $L(0) = 0, L(1) =$   
 25  $1, L'(0^+) \geq 0, L''(u) \geq 0$*

**Supplementary Table 1. 1-9. Lorenz curve models from distributional origin. 10.-17. Functional forms proposed to model Lorenz curves. Model 14 is recognized as a family of Lorenz curves but not proposed as a Lorenz curve specifically. As this family is the most general form of the specific Lorenz curve that Sarabia proposes, we use it as a four-parameter Lorenz curve (see (4, 6-9)).  $\eta$  denotes the cumulative percentage of income,  $u$  denotes the cumulative percentage of the population.  $\Phi(\cdot)$  is the cumulative distribution function of the standard normal distribution,  $G(\cdot)$  is the incomplete gamma function ratio,  $B(\cdot)$  is the lower incomplete beta function ratios as defined in SI Section 1.**

Originates from	Lorenz curve $\eta(u)$	# Par.	Parameter restrictions
1. Pareto distribution	$1 - (1 - u)^{1-1/\alpha}$	1	$\alpha > 1$
2. Lognormal distribution	$\Phi(\Phi^{-1}(u) - \sigma)$	1	$\sigma > 0$
3. Gamma distribution	$G(G^{-1}(u; \sigma); \sigma + 1)$	1	$\alpha, \sigma > 0$
4. Weibull distribution	$G(-\log(1 - u); \frac{1}{\alpha} + 1)$	1	$\alpha > 0$
5. Gen. Gamma distr.	$G(G^{-1}(u; p); p + \frac{1}{a})$	2	$a, p > 0$
6. Dagum distribution	$B(u^{1/q}; q + \frac{1}{a}, 1 - \frac{1}{a})$	2	$q > 0; a > 1$
7. Singh-Maddala distr.	$B(1 - (1 - u)^{1/q}; 1 + \frac{1}{a}, q - \frac{1}{a})$	2	$q, a > 0, q > \frac{1}{a}$
8. GB1 distribution	$B(B^{-1}(u; p, q); p + \frac{1}{a}, q)$	3	$p, q, a > 0$
9. GB2 distribution	$B(B^{-1}(u; p, q); p + \frac{1}{a}, q - \frac{1}{a})$	3	$p, q, a > 0; q > \frac{1}{a}$
10. Kakwani/Podder [1973] (10)	$ue^{-\beta(1-u)}$	1	$\beta > 0$
11. Rasche et al. [1980] (11)	$(1 - (1 - u)^\alpha)^{1/\beta}$	2	$0 < (\alpha, \beta) \leq 1$
12. Ortega et al. [1991] (12)	$u^\alpha(1 - (1 - u)^\beta)$	2	$\alpha \geq 0; 0 < \beta \leq 1$
13. Chotikapanich [1993] (13)	$\frac{e^{ku} - 1}{e^k - 1}$	1	$k > 0$
14. Sarabia et al. [1999] (14)*	$u^{\alpha+\gamma}[1 - a(1 - u)^\beta]^\gamma$	4	$0 \leq a \leq 1; 0 < \beta \leq 1;$ $0 \leq \alpha; \gamma \geq 1$
15. Abdalla/Hassan [2004] (15)	$u^\alpha(1 - (1 - u)^\delta e^{\beta u})$	3	$\alpha \geq 0; 0 \leq \beta \leq \delta \leq 1$
16. Rhode [2009] (16)	$u \cdot \frac{\beta-1}{\beta-u}$	1	$\beta > 1$
17. Wang et al. [2011] (17)	$\delta u^\alpha[1 - (1 - u)^\beta] + (1 - \delta)[1 - (1 - u)^{\beta_1}]^\nu$	5	$\alpha \geq 0; \nu \geq 0; \alpha + \nu \geq 1;$ $0 < (\delta, \beta, \beta_1) \leq 1$

## 26 2. Detailed Description of Data Cleaning

27 **General Procedure to Match the Datasets.** Data from both sources (American Community Survey (ACS) 2011-2015 (18),  
28 Economic Policy Institute (EPI) (19)) were collected at the US county level, which allows us to calculate the Lorenz curve  
29 representation of the income distribution using the following procedure: recall that the Lorenz curve is depicted through the  
30 cumulative share of population on the x-axis and cumulative share of income on the y-axis. We therefore construct a dataset  
31 that contains the share of population (from low-income to high-income) who own a certain percentage of total income, such  
32 that we can draw a Lorenz curve using the cumulative sum of these data points.

33 While the EPI report already presented the high-income earner data in such a way, further processing had to be undertaken  
34 for the ACS data: the data were given as headcounts per income bucket, which required transformation to income shares for  
35 the Lorenz curve representation. For this transformation, we assumed that people within income buckets were distributed  
36 symmetrically around the mean of the respective bucket. For example, a uniform distribution of people within an income  
37 bucket seems plausible in that people's income is likely to be equidistantly spread between the narrow boundaries of 45 000  
38 USD and 49 999 USD per year. We could then calculate the volume of income held by the people belonging to that bucket by  
39 multiplying the number of people in the respective income bucket with the mean value of the bucket range, and then dividing  
40 this number by the aggregate income in that county, giving us the share of total income. Based on this transformation, a Lorenz  
41 curve could be constructed for each US county. To verify that our approximated Lorenz curve data are in line with the true  
42 income share percentiles of that ACS dataset (the 20<sup>th</sup>, 40<sup>th</sup>, 60<sup>th</sup>, 80<sup>th</sup> and 95<sup>th</sup> income share percentiles are provided), we  
43 evaluated deviations between our approximated Lorenz curve and true income share data from ACS. We found good agreement  
44 between the approximated Lorenz curves with the ACS income shares, which we detail in Section 3.

45 Matching the ACS and EPI datasets revealed that, on average, the EPI data implied a higher level of inequality than the  
46 ACS data. This may arise in part because the EPI data are based on actual tax records at the taxpayer level, whereas the  
47 ACS data are from a self-reported survey at the household level, the latter of which is already an aggregate that typically  
48 underestimates the inequality suggested by the according Lorenz curve (20). For both ACS and EPI data, the exact 95<sup>th</sup>  
49 percentile was available, which enabled us to perform an exact scaling, i.e., adjusting the ACS household-level data to the EPI  
50 taxpayer-level data, using this data point as a link between datasets, see section 3 detailing this procedure. We adjust to the  
51 taxpayer level because it reflects the true level of income inequality in that individuals earn income, not households as a unit  
52 itself. We further believe that the EPI data are closer to reality, as tax reports are more difficult to manipulate and do not rely  
53 on self-reports that might be inaccurate, falsely remembered, or strategically misrepresented.

54 **Merging Source Tables.** This subsection =describes the code `data_cleaning_merge_b6_nhigs.R` which was used to merge the  
55 raw data tables provided by ACS and EPI.

56 We merge Tables B6 and B4\* from <https://www.epi.org/publication/income-inequality-in-the-us/#epi-toc-20> and Tables NHGIS  
57 A and NHGIS B from <https://data2.nhgis.org/main> that are from the American Community Survey 2011-2015. Source Table  
58 NHGIS A is taken from the dataset with NHGIS code 2011\_2015\_ACSa, and the source codes of the variables are B19001,  
59 B19013, B19025. Source Table NHGIS B is taken from the dataset with NHGIS code 2011\_2015\_ACSb, and the source codes  
60 of the variables are B19080, B19081, B19082, B19083. As additional information, a file with abbreviations and full names of  
61 US states (e.g. AK = Alaska) is taken from [https://developers.google.com/public-data/docs/canonical/states\\_csv](https://developers.google.com/public-data/docs/canonical/states_csv).

62 The procedure to merge the source tables is as follows:

- 63 • Load data and exclude Puerto Rico and the District of Columbia
- 64 • Merge ACS data NHGIS A and NHGIS B by county name such that all data from the survey are in a single dataset
- 65 • Adjust county names to prepare for the match: let the B6 county names (format: "San Francisco, CA") look like NHGIS  
66 county names (format: "San Francisco County, California"). To do so, the B6 county data is split at "," to separate the  
67 county name and state name. With the additional file on state abbreviations and names, the county state abbreviations  
68 are transformed into their actual name (e.g. from CA to California). Not only does the state name abbreviation differ in  
69 the B6 from the NHGIS format; it also says "San Francisco *County*, California". Therefore, to create a new B6 column  
70 that looks like the NHGIS county name, the county name (San Francisco), the word "County", ",", and the full state  
71 name "California" are pasted into a single column such that we end up with a column in B6 of the county name format  
72 "San Francisco County, California" to match with NHGIS
  - 73 - For the special cases Census Areas or Cities: don't paste "County" after "Census Area" or "City"
  - 74 - For the special case Alaska: Alaska is not divided into counties but into cities, boroughs, or census areas. NHGIS  
75 names them as City/Borough/Census Areas, but B6 does not, so we omit everything after the first word (which is a  
76 unique determinant of the actual area) in both datasets to derive a matching name for the corresponding area in Alaska
  - 77 - For the special case Louisiana: Louisiana is not divided into counties but into Parishes, so we paste "Parish"  
78 instead of "County" after county names in Louisiana
- 79 • Transform encoding of NHGIS data from 'ISO-8859-1' to = 'UTF-8'

\*Note that B4 is relevant not for the present study but for other (future) studies that intend using this dataset.

- Use a fuzzy string matching algorithm to merge B4/B6 and NHGIS data by county name: Fuzzy string matching has to be double checked by visual inspection of the county names to ensure that only correct merges have taken place. Iterative procedure to minimize the amount of counties that have to be inspected and matched by hand: From all the imperfect matches (distance > 0), which exhibit a very similar pattern, e.g. “St.” instead of “St”, transform “St.” to “St” such that all of these cases are now perfect matches (distance = 0). Fuzzy match again and repeat procedure. When most of the common structures like “St.” -> “St” are cured, we can inspect the resulting imperfect matches for counties that we need to match by hand. For some counties, different names exist, e.g. Shannon County, South Dakota, is another name for Oglala Lakota County, South Dakota
- Write a single .csv file for the merged tables

### 3. Calculation of the Lorenz Curves

This subsection describes the code `create_lorenz_curves.R` to calculate Lorenz curve values for each county. The goal is to calculate the share of income held by shares of the population (from low-income to high-income). A quick recap of the information that the ACS and EPI source tables give us:

- Table B6: Income share held by 90<sup>th</sup>, 95<sup>th</sup> and 99<sup>th</sup> percentile of the population → no further transformation needed
- NHGIS B: Income share held by 20<sup>th</sup>, 40<sup>th</sup>, 60<sup>th</sup>, 80<sup>th</sup> and 95<sup>th</sup> percentile of the population → no further transformation needed
- NHGIS A: Aggregated income per county, people per county, count of people that fall into a certain income bucket, e.g., have an income between 45,000 USD and 49,999 USD a year (see codebook in zip file for details) → need to transform this information, procedure:
  - Assume that people are symmetrically distributed around the mean value of the income bucket range within each closed income bucket, i.e., we do not use the top income bucket > 200,000 USD.
  - Use mean value of the income bucket range multiplied by the number of people that fall into that bucket as estimate of the income held by people belonging to the corresponding income bucket.
  - Divide this number by the income aggregate for the respective county, such that we end up with the share of total income held by the income bucket
  - Divide the number of people belonging to that income bucket by the total number of people in that county to get the share of people belonging to that income bucket
- Check for consistency in the ACS dataset: Inspect whether the estimated income shares per bucket are coherent with the information on the (true) income shares held by the 20<sup>th</sup>, 40<sup>th</sup>, 60<sup>th</sup>, 80<sup>th</sup>, and 95<sup>th</sup> percentile of the population → found to be consistent; see related [Supplementary Figure 1](#).
- Merge Lorenz curve data from ACS and EPI: Table B6 systematically suggests a higher level of inequality than the ACS data. This is a well-known phenomenon (20), as the ACS is at the household level (already an aggregate, e.g., two income earners living together in a household) whereas the B6 data are at the taxpayer level). We favor B6 data to depict a more realistic picture of the true inequality and hence decided to scale the ACS data to match the B6 data at the 95<sup>th</sup> percentile:
  - We have exact information on the 95<sup>th</sup> percentile, so we can use the 95<sup>th</sup> percentile as the anchor point for scaling to account for the difference in the data induced by the fact that B6 is at the taxpayer level and NHGIS at the household level. This means that we multiply the NHGIS percentile data by the 95<sup>th</sup> percentile of the B6 data and then divide it by the 95<sup>th</sup> percentile of the NHGIS data. To ensure convexity, we use solely ACS data below the 95<sup>th</sup> percentile and solely EPI data above the 95<sup>th</sup> percentile.
  - Check for data consistency prior and post scaling: Visually, most of the scaled data are close to the non-scaled data. However, as an example of an extreme case, which also illustrates that Table B6 delivers valuable information, we can look at Teton County, WY, further described in 3.

**Systematic Evaluation of Constructed Lorenz Curves.** We have already performed a brief cross-check for data consistency of the ACS dataset; i.e., we checked whether our approximation of income shares using the income buckets is close to the few true income share percentiles provided by the ACS. Now, we check the consistency of the ACS data more systematically.

We estimated the share of total income held by each income bucket (for all closed income buckets; i.e., we omit the top income bucket > 200,000 USD) under the assumption of symmetrically distributed incomes around the mean income within each income bucket. As we have true income share percentiles for some percentiles of the population, namely, the 20<sup>th</sup>, 40<sup>th</sup>, 60<sup>th</sup>, 80<sup>th</sup>, and 95<sup>th</sup> population percentiles, we can evaluate our estimated income shares by adding the true percentiles to our estimated Lorenz curves and for their fit. Remember that empirical Lorenz curves are defined by data points that are then linearly interpolated. Hence, we also linearly interpolate between our estimated income percentiles and calculate the estimated income percentile at the 20<sup>th</sup>, 40<sup>th</sup>, 60<sup>th</sup>, and 80<sup>th</sup> population percentile for which the ACS provides exact data. This allows us to

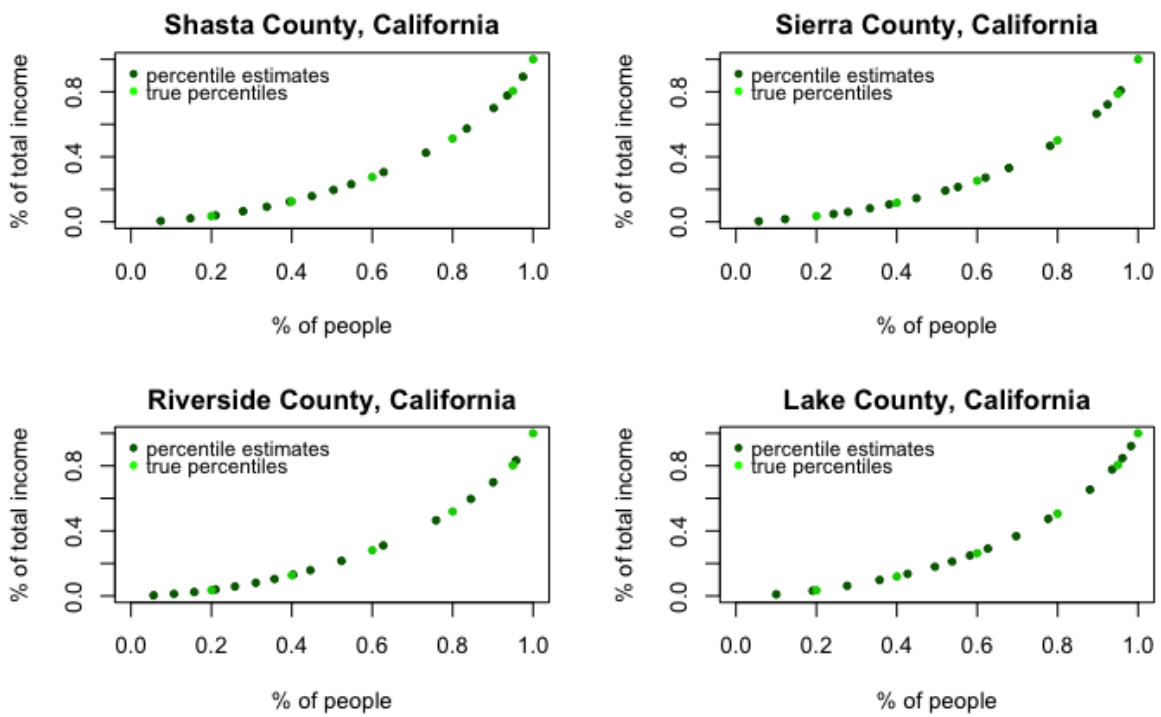
133 calculate the residual sum of squares (RSS) between the estimated income percentile at the 20<sup>th</sup>, 40<sup>th</sup>, 60<sup>th</sup> and 80<sup>th</sup> population  
134 percentiles and the true 20<sup>th</sup>, 40<sup>th</sup>, 60<sup>th</sup> and 80<sup>th</sup> percentiles.<sup>†</sup>

135 While [Supplementary Figure 1](#) already suggested that the estimated income percentiles from the income buckets seem to fit  
136 very well to the true income percentiles, we aim to quantify the fit more formally and calculate the RSS as described above. In  
137 [Supplementary Figure 2](#), we can see one clear outlier, and potentially three more. Hence, we take a closer look at the counties  
138 with the top four RSS scores, which turn out to be [1] Falls Church city, Virginia, [2] Monroe County, Alabama, [3] Allendale  
139 County, South Carolina, and [4] Holmes County, Mississippi.

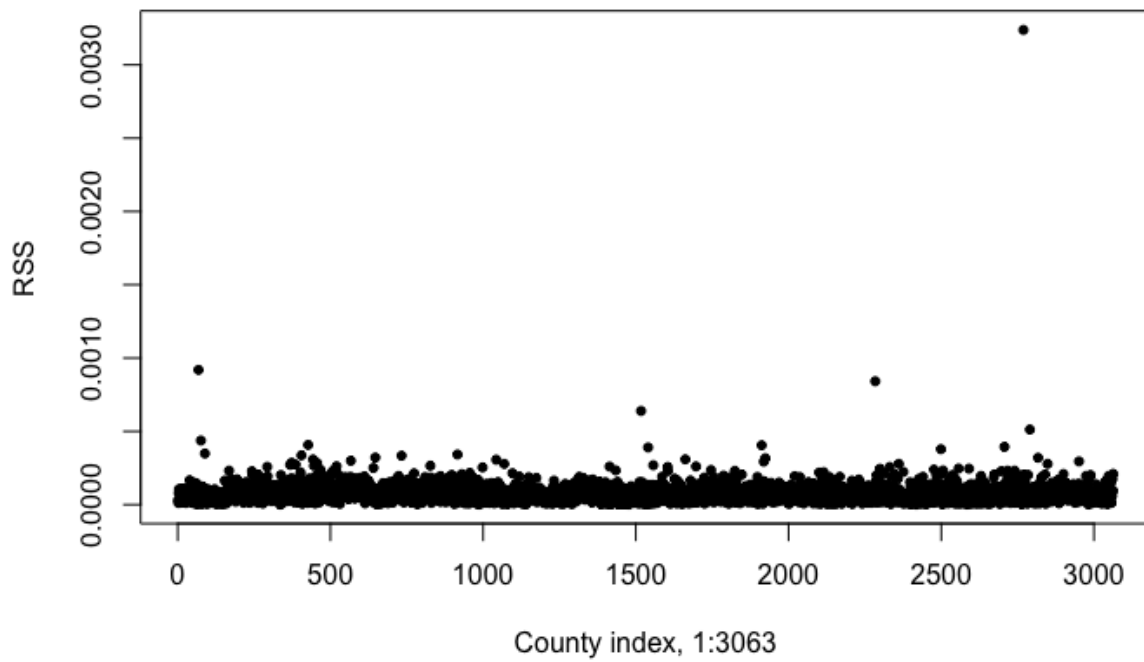
140 The Lorenz curve plots of these counties, depicted in [Supplementary Figure 3](#), reveal the following: for the county with the  
141 highest RSS score, Falls Church, we can clearly see that this high RSS score results from the fact that a significant fraction of  
142 its population falls into the top income bucket, > 200 000 USD. This forces a linear interpolation straight from a 0.73 percentile  
143 to the boundary of (1,1). We know this interpolation is not trustworthy, which is why we enrich the data at the top percentiles  
144 with EPI data and hence a comparably large deviation from the true 80<sup>th</sup> percentile should not worry us too much. For the  
145 remaining counties, the percentiles still seem to fit the Lorenz curve reasonably well. Therefore, we can conclude that there is  
146 no need to exclude any outliers from further analyses.

---

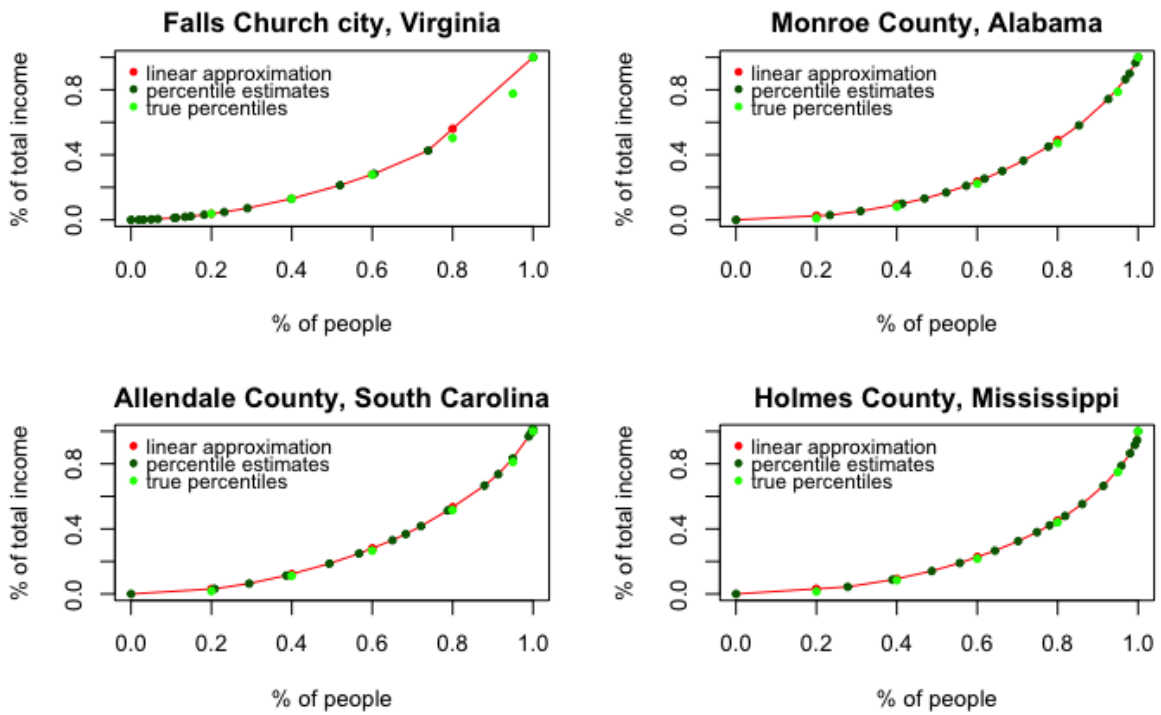
<sup>†</sup>We omit the 95<sup>th</sup> from the analyses here because we know that linear approximation is not a good approximation for top income shares, which is why we use EPI data from B6 for the 95<sup>th</sup> percentile and above.



Supplementary Figure 1. Estimated and true income percentiles for some exemplary counties



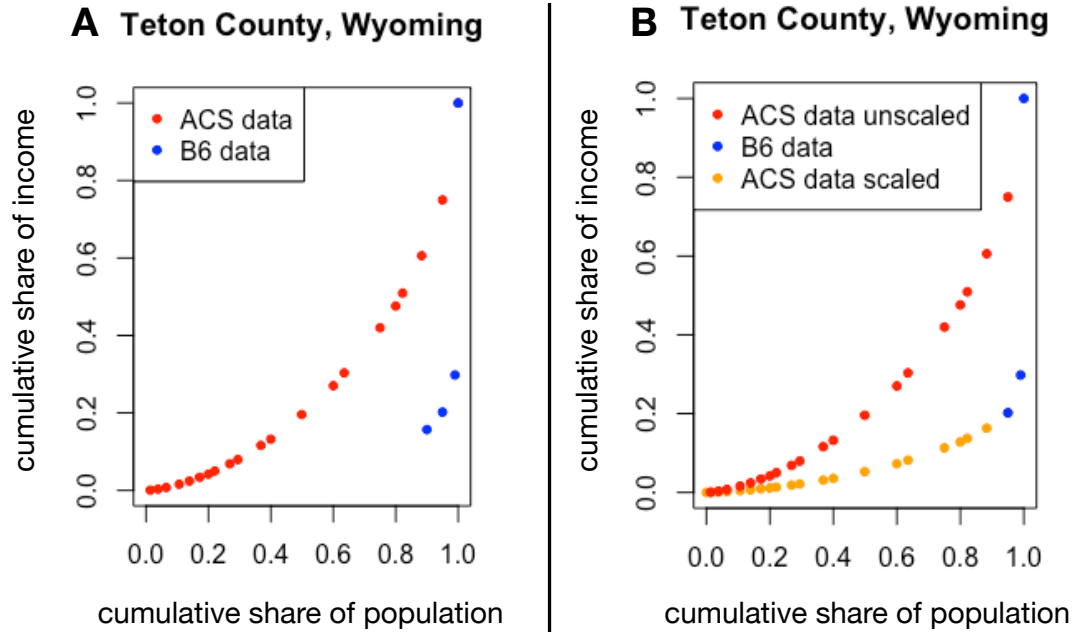
**Supplementary Figure 2.** RSS for estimated percentiles of income shares. Here, we refer to residuals as the difference between the true income share and estimated income share. Residuals are then squared and summed over all available data points. This was performed for each county out of all 3063 counties.



Supplementary Figure 3. Interpolated Lorenz curves from estimated income shares for the four counties with highest RSS scores in Supplementary Figure 2.

147 **Scaling the Data in an Exemplary County: Teton, Wyoming.** Teton, WY, is an example of a county that exhibits a special  
148 distribution of income that we could not have guessed with the ACS data alone. The data of the American Community Survey  
149 alone are fine-grained for low and medium income levels, yet the ACS data alone might lead to unrealistic approximations of  
150 the top populations' income shares, as the top income bucket > 200 000 USD is an open interval that does not provide any  
151 information on how people are distributed within that interval. Table B6, however, gives us detailed information on the income  
152 shares of the top-income percentages of the population on the taxpayer level.

153 Apparently, there are a few people living in Teton, WY, that have an income far above the threshold 200 000 USD. In  
154 [Supplementary Figure 4](#) Panel A, we can clearly see that the income share of the top 5% and top 1% percent of income  
155 earners far exceeds what we would have expected from the American Community Survey data. Now looking at the scaled data  
156 presented in Panel B of [Supplementary Figure 4](#), i.e., taking into consideration the information from the EPI dataset, we can  
157 clearly see the immense difference. This example highlights the importance of considering Table B6 as an additional data  
158 resource for the construction of close-to-reality Lorenz curves.



**Supplementary Figure 4.** Panel A provides raw Lorenz curve data from ACS and Table B6; Panel B depicts scaled data for Teton County, Wyoming.



#### 4. Maximum Likelihood Estimation (MLE) via Dirichlet Distribution

An approach to estimate Lorenz curves based on maximum likelihood estimation (MLE) was proposed by Chotikapanich et al. (2002) (21). They assume the income shares from grouped data to follow a Dirichlet distribution. Chang et al. (2018) (22) agree with this perspective and argue that the Dirichlet distribution “naturally accommodates the proportional nature of income share data and the dependence structure between the shares” (22, p. 2), which is a major advantage compared with the NLS estimation procedure (15). Chotikapanich et al. (2002) (21) demonstrate analytically that it is possible to relate desired functional forms of the Lorenz curve to the Dirichlet parameters; i.e., parameters of the Dirichlet distribution are set so that they incorporate the proposed functional form of the Lorenz curve with its parameters. The density of the Dirichlet distribution (with newly defined parameters that consist of the Lorenz curve parameters) is then used to construct the likelihood that is maximized later on. In detail, the procedure to model the Lorenz curve models with maximum likelihood estimation using the Dirichlet distribution described in (21) is as follows:

Let  $\eta_i = L(u_i; \theta)$  be the cumulative income share held by the cumulative share of the population  $u_i$ . Then,  $q = (q_1, \dots, q_M)$  with  $q_i = \eta_i - \eta_{i-1}$  are assumed to be random variables that follow a Dirichlet distribution. The probability density function of the Dirichlet distribution is given by

$$f(q|\alpha) = \frac{\Gamma(\alpha_1 + \alpha_2 + \dots + \alpha_M)}{\Gamma(\alpha_1)\Gamma(\alpha_2)\dots\Gamma(\alpha_M)} \cdot q_1^{\alpha_1-1} q_2^{\alpha_2-1} \dots q_M^{\alpha_M-1}$$

where the gamma function is defined as  $\Gamma(\alpha) = \int_0^\infty x^{\alpha-1} \exp^{-x} dx$ . The method is now to relate the parameters  $\alpha$  of the Dirichlet distribution to the functional form of the Lorenz curve that we want to estimate. This can be conveniently be done by setting

$$\alpha_i = \lambda[L(u_i; \theta) - L(u_{i-1}; \theta)]$$

where  $\lambda$  is an additional unknown parameter. Now we can write the probability density function as

$$f(q|\lambda, \theta) = \Gamma(\lambda) \prod_{i=1}^M \frac{q_i^{\lambda[L(u_i; \theta) - L(u_{i-1}; \theta)] - 1}}{\Gamma(\lambda[L(u_i; \theta) - L(u_{i-1}; \theta)])}$$

To now estimate the parameters, we simply have to maximize the log-likelihood that takes the form

$$\log[f(q|\lambda, \theta)] = \log \Gamma(\lambda) + \sum_{i=1}^M (\lambda[L(u_i; \theta) - L(u_{i-1}; \theta)] - 1) \cdot q_i - \sum_{i=1}^M \log \Gamma(\lambda[L(u_i; \theta) - L(u_{i-1}; \theta)])$$

This maximum likelihood based estimation of Lorenz curve parameters is, however, not widely used. The original study of (21) was replicated and advanced by (22) and (15), finding mixed results. In detail, (22) find that the MLE estimation via the Dirichlet distribution provides a better fit to empirical data, and (15) find that NLS provides a “better and more reliable fit compared to the maximum likelihood estimation” (15, p. 117). Moreover, (21) find that most Lorenz curve parameter estimates are not sensitive to the estimation method; i.e., they compared parameters estimated by NLS and MLE and found them yielding very similar point estimates for the parameters for most Lorenz curves proposed (but not all of them, which they attribute to estimation instability). (15) find similar point estimates of NLS and MLE as well, but report, as (21), much larger standard errors of the estimated parameters of the MLE method.

#### 5. Akaike Information Criterion (AIC) and AIC<sub>c</sub> Simulation Study

While the AIC measure of goodness-of-fit is well known as a tool for model selection in many fields of applied statistics, such as ecology (23) or astrophysics (24), it has not previously been used to systematically analyze the optimal number of parameters needed to adequately represent empirical Lorenz curves. One reason the AIC has not been used in prior literature may be the more common use of nonlinear least squares (NLS) approaches as an estimation procedure for Lorenz curves, which does not allow for the use of AIC. The NLS approach is widespread because it does not impose distributional assumptions on the data, which is a requirement for MLE. However, within the NLS framework, researchers typically rely on the residual sum of squares as a measure of goodness-of-fit. Residual sum of squares does not trade-off fit for model complexity, which commonly results in the most complicated Lorenz model as the winner. For our research question—determining how many parameters are necessary to capture relevant information—we therefore focus on the MLE/AIC framework in order to balance complexity and model fit. As mentioned in the paper, we use the small-sample bias adjusted version of the criterion, namely AIC<sub>c</sub>.

Our key question we want to answer with our simulation study is: Will AIC<sub>c</sub> suggest that we use the correct model? To answer this question, we will simulate Lorenz curve data points according to a certain model. Based on these data points, we will estimate the parameters of all 17 models we analyzed in the previous chapters and then let AIC<sub>c</sub> choose the best model. If AIC<sub>c</sub> actually picks the correct model that was used for data generation sufficiently often, the reliability of AIC<sub>c</sub> as a criterion for model selection is supported for our setting. However, if AIC<sub>c</sub> fails to pick the correct model, we have to question our previous results and take them with a (big) grain (rock) of salt. We will vary the sample size, i.e., the number of data points used for model estimation, to get a clearer picture of where our setting stands with respect to the extent to which we trust in AIC<sub>c</sub> picking the correct model. Only then we can judge whether AIC<sub>c</sub> can be used as an indicator of the number of parameters needed to describe income-inequality Lorenz curves.

210 Through an  $AIC_c$ -based ranking and Borda voting procedure, we found the Ortega Lorenz curve model (2 parameters),  
211 the GB2 Lorenz curve model (3 parameters), and the Wang Lorenz curve model (5 parameters) to be among the most  
212 suitable models. To verify that our judgment, especially between those three most promising models, is trustworthy, we will  
213 focus on those three models for income share generation. In detail, we will run three simulations, with the only difference  
214 being the model used to generate the income shares. One might wonder why we run the simulation not only with one  
215 exemplary income-generating model but with three models. The reason is that we then can cross-compare results between the  
216 income-generating routines. For example, we could detect whether a certain model is preferred by  $AIC_c$  regardless of the true  
217 data-generating process. In other words,  $AIC_c$  might always choose the same model.

218 **Simulation Setup.** For ease of comprehensibility, we will describe the simulation procedure in a numbered list. The structure of  
219 the simulation study is as follows:

- 220 1. Generate a vector that imitates population shares:  $\pi = (0, \pi_1, \dots, \pi_n, 1)$  with  $\pi_i \sim \text{Unif}(0,1)$ .
- 221 2. Generate a vector of cumulative income shares  $\eta = L(\pi, \theta)$ , where  $L(\theta)$  is a known Lorenz curve model of either type  
222 Ortega, GB2, or Wang with known parameters  $\theta^\ddagger$  and population shares  $\pi$  that were generated in the previous step. For  
223 each Lorenz curve model used for income-share generation, we run a separate simulation.
- 224 3. Use MLE to fit all 17 Lorenz curve models<sup>§</sup> to the data generated above and store the model name with minimum  $AIC_c$   
225 value.
- 226 4. Evaluate whether  $AIC_c$  has chosen the model that was used to generate the cumulative income shares or not.
- 227 5. Repeat this procedure for  $sim = 1\,000$  population share vectors generated. Then vary the length of the population share  
228 vector and apply the same procedure.
- 229 6. Evaluate the percentage of instances where  $AIC_c$  was able to detect the model that was used for income-share generation  
230 for each vector length and each of the the Lorenz curve models that are used to generate income.

231 **Simulation Results.** Results show that we observed a high true-model detection rate even for small sample sizes, see Tables  
232 [Supplementary Table 2](#), [Supplementary Table 3](#), [Supplementary Table 4](#), and Figures [Supplementary Figure 5](#), [Supplementary](#)  
233 [Figure 6](#), [Supplementary Figure 7](#). For our sample size range of 19-23 data points—and assuming that the two-parameter  
234 Ortega truly was the Lorenz curve generating model—the true discovery rate would be  $\geq 0.97$  (lower bound of 95% confidence  
235 interval), see [Supplementary Table 4](#) and [Supplementary Figure 7](#)). This result provides additional confidence in the reliability  
236 of  $AIC_c$  given our specific setting.

---

<sup>‡</sup>To find reasonable parameters, we used the mean value across the US county parameter estimates.

<sup>§</sup>See Table 1 in the paper

**Supplementary Table 2. Rate of bias corrected AIC picking the true data-generating model for varying sample sizes out of 1 000 simulation runs. A sample size of 102 means we have 100 data points generated between 0 and 1, plus 0 and 1 as boundary values. Lower and upper bounds correspond to the 95% confidence interval, based on a binomial test. True model: GB2**

sample size	rate	lower bound	upper bound
6	0.766	0.738	0.792
7	0.830	0.805	0.853
8	0.879	0.857	0.899
9	0.881	0.859	0.900
10	0.894	0.873	0.912
11	0.909	0.889	0.926
12	0.915	0.896	0.932
13	0.923	0.905	0.939
14	0.911	0.892	0.928
15	0.912	0.893	0.929
16	0.921	0.903	0.937
17	0.911	0.892	0.928
18	0.909	0.889	0.926
19	0.921	0.903	0.937
20	0.917	0.898	0.933
21	0.921	0.903	0.937
22	0.914	0.895	0.931
23	0.920	0.901	0.936
24	0.910	0.891	0.927
25	0.923	0.905	0.939
26	0.921	0.903	0.937
27	0.925	0.907	0.941
32	0.928	0.910	0.943
42	0.940	0.923	0.954
52	0.952	0.937	0.964
77	0.974	0.962	0.983
102	0.974	0.962	0.983
127	0.981	0.970	0.989
152	0.992	0.984	0.997
177	0.991	0.983	0.996
202	0.978	0.967	0.986


**Supplementary Table 3. Rate of bias corrected AIC picking the true data-generating model for varying sample sizes out of 1000 simulation runs. A sample size of 102 means we have 100 data points generated between 0 and 1, plus 0 and 1 as boundary values. Lower and upper bounds correspond to the 95% confidence interval, based on a binomial test. True model: Wang**

sample size	rate	lower bound	upper bound
6	0.000	0.000	0.004
7	0.000	0.000	0.004
8	0.169	0.146	0.194
9	0.408	0.377	0.439
10	0.558	0.527	0.589
11	0.648	0.617	0.678
12	0.683	0.653	0.712
13	0.709	0.680	0.737
14	0.727	0.698	0.754
15	0.739	0.711	0.766
16	0.728	0.699	0.755
17	0.755	0.727	0.781
18	0.768	0.741	0.794
19	0.739	0.711	0.766
20	0.765	0.737	0.791
21	0.780	0.753	0.805
22	0.775	0.748	0.801
23	0.774	0.747	0.800
24	0.781	0.754	0.806
25	0.802	0.776	0.826
26	0.765	0.737	0.791
27	0.776	0.749	0.801
32	0.787	0.760	0.812
42	0.825	0.800	0.848
52	0.842	0.818	0.864
77	0.878	0.856	0.898
102	0.923	0.905	0.939
127	0.933	0.916	0.948
152	0.959	0.945	0.970
177	0.969	0.956	0.979
202	0.983	0.973	0.990

**Supplementary Table 4. Rate of bias corrected AIC picking the true data generating model for varying sample sizes out of 1 000 simulation runs. A sample size of 102 means we have 100 data points generated between 0 and 1, plus 0 and 1 as boundary values. Lower and upper bounds correspond to the 95% confidence interval, based on a binomial test. True model: Ortega**


sample size	rate	lower bound	upper bound
6	0.986	0.977	0.992
7	0.991	0.983	0.996
8	0.995	0.988	0.998
9	0.992	0.984	0.997
10	0.984	0.974	0.991
11	0.984	0.974	0.991
12	0.984	0.974	0.991
13	0.980	0.969	0.988
14	0.978	0.967	0.986
15	0.986	0.977	0.992
16	0.983	0.973	0.990
17	0.983	0.973	0.990
18	0.985	0.975	0.992
19	0.985	0.975	0.992
20	0.989	0.980	0.994
21	0.981	0.970	0.989
22	0.985	0.975	0.992
23	0.986	0.977	0.992
24	0.987	0.978	0.993
25	0.981	0.970	0.989
26	0.972	0.960	0.981
27	0.969	0.956	0.979
32	0.965	0.952	0.976
42	0.972	0.960	0.981
52	0.980	0.969	0.988
77	0.986	0.977	0.992
102	0.984	0.974	0.991
127	0.981	0.970	0.989
152	0.994	0.987	0.998
177	0.991	0.983	0.996
202	0.997	0.991	0.999

### Model choice of bias corrected AIC for varying sample sizes




**Supplementary Figure 5.** Simulation results for GB2 being the true income share generating model out of a selection of 17 possible models. Point estimates of the percentage of correct model detection are reported together with confidence bounds of the 95% confidence interval.

### Model choice of bias corrected AIC for varying sample sizes



**Supplementary Figure 6.** Simulation results for Wang being the true income share generating model out of a selection of 17 possible models. Point estimates of the percentage of correct model detection are reported together with confidence bounds of the 95% confidence interval.

**Model choice of bias corrected AIC for varying sample sizes**



**Supplementary Figure 7.** Simulation results for Ortega being the true income share generating model out of a selection of 17 possible models. Point estimates of the percentage of correct model detection are reported together with confidence bounds of the 95% confidence interval.



## 6. Voting

For interested readers, we recommend the literature of the *Handbook of Social Choice and Welfare* (25), which describes the voting procedures in more depth. This section is based on this handbook as well and aims to present voting procedures relevant for our study in a comprehensive way.

According to Arrow's impossibility theorem, there exists no single best voting procedure across the board (26). As a result, researchers have to choose the voting procedure that best fits the problem at hand. We suggest that the Borda count is particularly well suited for our context as it provides insight into which fitted model has good performance across all counties instead of a great fit in some counties but an inferior fit in other counties. We note that others arrive at a different conclusion and prefer a different voting procedure; in that case, we encourage interested readers to use our comprehensive voting results given in the subsection below.

Relying on the principle 'the winner takes all,' *plurality voting* is a simple and intuitive voting procedure. Each individual has one vote, and the candidate receiving the most votes wins. Of course, this reveals only a fraction of the voters' preferences, namely their top choice, but it neglects any remaining preference ordering behind the top choice. In our case, plurality voting corresponds to evaluating which Lorenz curve model was ranked first the most.

A procedure that does not only take the first choice into consideration but performs pairwise comparisons between options is the so-called *Condorcet procedure*. In detail, each option is compared with any other option, and a winner between those options is determined. A quick example illustrates the procedure: Imagine that there are three possible options, A, B, and C, to choose from. Individual 1 has the preference ordering  $A > B > C$ <sup>4</sup> while the preference of individual 2 is  $B > C > A$ . To aggregate the preferences of both individuals, we can now compare how often an option was ranked ahead of another option. In this case, option A was preferred over B once (by individual 1), B was preferred over C twice (by individual 1 and 2), and C was preferred over A once (by individual 2); so in this case, the winner of the Condorcet procedure is option B. As we have an  $AIC_c$ -based ranking between Lorenz curve models for each county, we can perform such pairwise comparisons across counties. Note that the dominance matrix introduced above depicts these pairwise comparisons, i.e., displays how often a certain model was preferred over the remaining Lorenz curve models.

However, the Condorcet procedure can result in circular preferences and compares the options only in a pairwise fashion. A voting procedure that fully takes into account the ranking of the options is the so-called *Borda count*. This procedure scores the different options according to their ranks. In detail, if there are  $n$  options to choose from, the option ranking first receives  $n$  points, the option ranking second  $n - 1$  points, . . . , the least favored option receives 0 points. The points received are summed for all individuals, and the option receiving the most points wins the Borda count. Thus, options with a consistently high ranking have a greater chance to win than options that are brilliant for some individuals but heavily undesirable for others. This is exactly the behavior we desire for our Lorenz curve model comparison: we want to detect the model that overall achieves good performance across counties. Therefore, the Borda count is the most relevant voting procedure for our purpose.

**Voting results.** It is important to again emphasize that the Borda count winner is not the only choice one could make. Other Lorenz curve models winning other voting procedures might be legitimate models as well. The crucial point is that one has to decide which aspects to focus on. By design, different voting mechanisms will lead to different model winners, as they—purposely—emphasize different aspects. Where researchers want to emphasize other aspects, another Lorenz curve model might be more useful. As Arrow's impossibility theorem states, the aggregation of preferences cannot be performed using a single best selection procedure but with different procedures for different kinds of problems and suitable outcomes. For our setting, we find the Borda count procedure superior. However, we do not want to discourage researchers from concluding that other Lorenz curve models might be superior if faced with a different scenario. We therefore provide various voting results below.

In our application, the results are as follows: in plurality voting, the Wang Lorenz curve model wins; applying the Condorcet procedure, the winner is the GB2 Lorenz curve model; and the Borda count winner is the Ortega Lorenz curve model. As the Borda count voting procedure depends on the goodness-of-fit criterion used to judge the models, we cross-check whether those results are driven by  $AIC_c$  or whether they are robust to the use of another information criterion. Therefore, we rerun the Borda voting procedure using the Bayesian information criterion (BIC) as indicator to rank the models. The BIC is defined as

$$BIC = -2 \cdot \ell(\hat{\theta}) + 2p \cdot \ln(n)$$

Voting results are similar to the  $AIC_c$ -based Borda count; see 6. This result shows that these three models (Wang, GB2, and Ortega) are the most promising.

<sup>4</sup>In words: Individual 1 prefers option A over B over C, so individual 1 ranks A first, B second, and C third.

	ABDALLA_HASSAN	CHOTIKAPANICH	DAGUM	GAMMA	GB1	GB2	GENERALIZED_GAMMA	KARWANI_PODDER	LOGNORMAL	ORTEGA	PARETO	RASCHE	RHODE	SARABIA	SINGH_MADDALA	WANG	WEIBULL
ABDALLA_HASSAN	0	3038	1243	3039	2585	161	2800	3037	2894	69	3053	2156	3055	2192	1642	1344	3046
CHOTIKAPANICH	18	0	17	9	1	18	0	2992	18	17	796	17	99	19	17	20	13
DAGUM	1813	3039	0	3042	2594	1154	2840	3038	2984	773	3051	2838	3055	2079	2084	1361	3046
GAMMA	17	3047	14	0	65	15	25	3046	32	12	1919	11	3000	38	11	41	2633
GB1	471	3055	462	2991	0	439	359	3033	467	477	2965	443	3055	603	439	550	3023
GB2	2895	3038	1902	3041	2617	0	2836	3037	2938	1619	3054	2401	3055	2587	2255	1586	3045
GENERALIZED_GAMMA	256	3056	216	3031	2697	220	0	3055	233	237	2808	186	3055	493	179	412	3044
KARWANI_PODDER	19	64	18	10	3	19	1	0	18	18	753	18	97	19	18	21	14
LOGNORMAL	162	3038	72	3024	2589	118	2823	3038	0	101	2738	45	3055	507	44	360	3035
ORTEGA	2987	3039	2283	3044	2579	1437	2819	3038	2955	0	3055	2684	3055	2442	2599	1533	3048
PARETO	3	2260	5	1137	91	2	248	2303	318	1	0	4	1887	4	0	4	1255
RASCHE	900	3039	218	3045	2613	655	2870	3038	3011	372	3052	0	3055	1682	446	1085	3050
RHODE	1	2957	1	56	1	1	1	2959	1	1	1169	1	0	2	1	4	140
SARABIA	864	3037	977	3018	2453	469	2563	3037	2549	614	3052	1374	3054	0	1079	1102	3030
SINGH_MADDALA	1414	3039	972	3045	2617	801	2877	3038	3012	457	3056	2610	3055	1977	0	1256	3048
WANG	1712	3036	1695	3015	2506	1470	2644	3035	2696	1523	3052	1971	3052	1954	1800	0	3026
WEIBULL	10	3043	10	423	33	11	12	3042	21	8	1801	6	2916	26	8	30	0

Supplementary Figure 8. Condorcet matrix on a county level. Count of how often models in the rows achieve a higher  $AIC_c$  rank than models in the columns, out of all 3 056 counties.

**Supplementary Table 5. Plurality voting results. In each county, the Lorenz curve model with the lowest  $AIC_c$  value gets one vote. The model with the highest number of total votes wins.**

Num. of Parameters	Model	Votes
5	Wang	998
2	Ortega	546
2	Dagum	399
3	GB2	364
3	GB1	355
4	Sarabia	153
2	Generalized Gamma	80
2	Rasche	70
2	Singh-Maddala	53
1	Lognormal	28
1	Gamma	6
1	Weibull	2
3	Abdalla-Hassan	1
1	Kakwani-Podder	1

**Supplementary Table 6. Borda count result using  $AIC_c$  as information criterion. In each county, the Lorenz curve models were scored using the Borda count procedure. The model with the highest Borda score wins.**

Num. of Parameters	Model	Borda Score
2	Ortega	42597
3	GB2	41906
2	Dagum	38791
5	Wang	38187
2	Singh-Maddala	36274
3	Abdalla-Hassan	35354
4	Sarabia	32272
2	Rasche	32131
1	Lognormal	24749
2	Generalized Gamma	23178
3	GB1	22852
1	Gamma	13926
1	Weibull	11400
1	Pareto	9522
1	Rhode	7296
1	Chotikapanich	4071
1	Kakwani-Podder	1110

**Supplementary Table 7. Borda count result using BIC as information criterion. In each county, the Lorenz curve models were scored using the Borda count procedure. The model with the highest Borda score wins.**

Num. of Parameters	Model	Borda Score
2	Ortega	42595
3	GB2	41760
5	Wang	38861
2	Dagum	38806
2	Singh-Maddala	36297
3	Abdalla-Hassan	35153
4	Sarabia	32208
2	Rasche	32109
1	Lognormal	24830
2	Generalized Gamma	23084
3	GB1	22779
1	Gamma	13931
1	Weibull	11420
1	Pareto	9594
1	Rhode	7310
1	Chotikapanich	3909
1	Kakwani-Podder	970

286 **7. Analysis of BIC differences**


In order to rule out that the choice of information criterion ( $AIC_c$ ) influenced the results of our analysis, we reran the  $\Delta$  analysis while using the Bayesian information criterion (BIC). The differences in BIC are defined in analogy to the  $AIC_c$  differences ( $\Delta$ ) as

$$\text{BIC difference} = BIC_i - BIC_j \tag{1}$$


287 For BIC, the analysis of differences is also applied in the literature, yet with a slightly differing usage of wording and  
288 boundaries. While the interpretation of the differences is the same for both differences in  $AIC_c$  and BIC (namely, the larger the  
289 difference between the values, the less support there is for the competing model's ability to provide as good an approximation of  
290 the data as the other one), the boundaries are shifted. (27) sets the boundaries of BIC differences as described in 7. Respecting  
291 those boundaries, we arrive at similar histograms as with the analysis of  $AIC_c$  differences; see Figures [Supplementary Figure 9](#),  
292 [Supplementary Figure 10](#), and [Supplementary Figure 11](#). Hence, we conclude that the superiority of Ortega compared with  
single-parameter models is irrespective of the chosen information criterion.

BIC difference	Evidence
0-2	weak
2-6	positive
6-10	strong
>10	very strong

293




**Supplementary Figure 9.** Histogram of BIC differences between the one-parameter lognormal model  $i$  and the two-parameter Ortega and  $j$ .



Supplementary Figure 10. Histogram of BIC differences between the three-parameter GB2 model  $i$  and the two-parameter Ortega and  $j$ .





Supplementary Figure 11. Histogram of BIC differences between the five-parameter Wang model  $i$  and the two-parameter Ortega and  $j$ .

## 294 8. $\Delta$ -AIC analysis of Ortega vs. GB2 and Ortega vs. Wang model

295 Using the Borda count voting procedure, we have determined the two-parameter Ortega Lorenz curve to be the winning model.  
296 However, the GB2 model using three parameters tightly comes second in the Borda count, and the Wang five-parameter model  
297 also performs well and wins the majority voting procedure. So do the three- and five-parameter models potentially provide  
298 substantially more information for some counties than a two-parameter model? To investigate this question, we calculated the  
299  $AIC_c$  differences between Ortega and GB2 as well as Ortega and the Wang model.

300 We draw on prior literature, namely the guidelines given by Burnham and Anderson (28), to set up an evaluation strategy  
301 tied to the specific problem at hand of investigating the extent to which a certain model fits the data better than other models.

302 Burnham and Anderson (28) acknowledge that an interpretation of absolute AIC values, and hence a comparison between  
303 competing models, is hindered because of arbitrary constants. Instead, (28) propose using differences in AIC values,  $\Delta_i =$   
304  $AIC_i - AIC_{min}$ , that represent the information loss experienced when using model  $i$  rather than the best model which exhibits  
305 the minimum AIC value  $AIC_{min}$ . The severity of information loss can be characterized by defining intervals for  $\Delta_i$  values, with  
306 larger values representing a higher amount of information loss. Burnham and Anderson (28) provide some rules of thumb:  
307 Models  $i$  with  $\Delta_{i,j} \leq 2$  have substantial support; for  $4 \leq \Delta_{i,j} \leq 7$  considerably less support and for  $\Delta_{i,j} > 10$  no support for  
308 being the best approximating model in the candidate set. In other words, the higher the  $\Delta_i$  value, the less support there is  
309 for the hypothesis that the two models of comparison provide an equally well characterization of the empirical data. This  
310 information can then be used to evaluate the strength-of-evidence in favor of the minimum AIC model (28), i.e., to get a sense  
311 of whether the minimum AIC model is substantially better.


312 For the setting of a Lorenz curve comparison as outlined in this paper, we generalize the evaluation strategy of (28) and  
313 fine-tune the interpretation in order to provide a more intuitive understanding. First, let us note that we work with the  
314 small-sample bias corrected version of AIC values ( $AIC_c$  values), which does not affect the evaluation strategy, but changes the  
315 name of the strategy to evaluating  $AIC_c$  differences instead of AIC differences. Second, we do not necessarily compare the  
316 model of interest to the minimum  $AIC_c$  model in the respective US county, but fixed models, e.g., Ortega versus lognormal  
317 model. Hence, instead of  $\Delta_i = AIC_i - AIC_{min}$ , we introduce a more general version  $\Delta_{i,j} := AIC_i - AIC_j$ . To enhance ease of  
318 interpretation, we do not take on the perspective of (28) that focus on characterising the support of various models in being the  
319 best approximation of the data, but propose a slightly different perspective on the values: Starting off with the interpretation  
320 of (28) that  $\Delta_{i,j}$  represents the information loss experienced when using model  $i$  rather than model  $j$ , we frame the  $\Delta_{i,j}$  values  
321 directly as strength-of-evidence in favor of model  $j$ . This means that higher values of  $\Delta_{i,j}$  provide evidence in favor of model  
322  $j$  capturing the information given by the empirical data more aptly. With this general setup of  $\Delta_{i,j}$  values, we might now  
323 encounter the situation of negative values in  $AIC_c$  differences, which is not possible with the AIC difference values defined in  
324 (28) as they set model  $j$  to the model with minimum AIC value. However, negative values of  $AIC_c$  values simply correspond to  
325 the case where  $i$  and  $j$  are reversed, hence gathering evidence for model  $i$  or, in other words, evidence for counter model  $j$ .  
326 Finally, we are forced to redefine the value intervals: (28) leave out interpretation guidelines for  $\Delta_i$  in the intervals  $[2, 4]$  and  
327  $[7, 10]$ , and we therefore extend their intervals in a conservative manner.

328 In summary, our strength-of-evidence classification in terms of  $AIC_c$  differences is as follows: We find inconclusive evidence  
329 on whether model  $j$ , e.g., the Ortega model, is superior in modeling relevant information compared to model  $i$ , e.g., the  
330 lognormal model, if the  $AIC_c$  difference  $\Delta_{i,j}$  is in  $[-4, 4]$ , some evidence that model  $j$  is superior if  $\Delta_{i,j} \in [4, 10]$  and decisive  
331 evidence that model  $j$  is superior to model  $i$  if  $\Delta_{i,j} > 10$ . If  $\Delta_{i,j} \in [-4, -10]$ , we find some evidence *against* model  $j$ 's  
332 superiority, and decisive evidence *against* model  $j$ 's superiority for  $\Delta_{i,j} < -10$ . With histograms of  $AIC_c$  differences ( $\Delta_{i,j}$ ), we  
333 can see how often, i.e., in how many US counties, we find supporting evidence for whether one model indeed provides more  
334 substantial information about the data.


335 As a recap, for the comparison between an Ortega two-parameter model and the single-parameter lognormal model, we find  
336 a clear picture in support of the two-parameter model; see Figure 2 in the main text.

337 Now evaluating Ortega versus GB2, we see a much more inconclusive picture; see [Supplementary Figure 12](#). For most of the  
338 counties, there is inconclusive evidence; i.e., there is substantial support that both models perform similarly well in modeling  
339 the information given in the empirical data. This indicates that the three- and two-parameter models are somewhat comparable.  
340 Given this information, it is debatable which model to prefer, but as Ortega is the simpler model, we clearly favor it over GB2.

341 In comparing the Ortega model and the five-parameter Wang model, we get a more distinct histogram; see [Supplementary](#)  
342 [Figure 13](#). On the one hand, we can clearly see that for many counties, we have evidence that the five-parameter model  
343 captures relevant information better than the two-parameter Ortega model. On the other hand, we find counter-evidence in  
344 many counties as well: i.e., that the two-parameter model performs that task better. This result is unsurprising given which  
345 aspects the various voting procedures emphasize: the Borda count values good performance across counties (Ortega won),  
346 whereas majority voting honors how often a model performs best in a county (Wang won). That is, the Wang five-parameter  
347 model is excellent many times but also inferior many times compared with the two-parameter Ortega model. As we seek a  
348 model that performs well across all US counties, we prefer Ortega for that purpose.



**Supplementary Figure 12.** Histogram of  $AIC_c$  differences ( $\Delta_{i,j}$ ) between the three-parameter GB2 model  $i$  and the two-parameter Ortega and  $j$ .



**Supplementary Figure 13.** Histogram of  $AIC_c$  differences ( $\Delta_{i,j}$ ) between the five-parameter Wang model  $i$  and the two-parameter Ortega and  $j$ .

## 349 9. Nonlinear Least Squares (NLS)

In terms of Lorenz curves, we are dealing with functions that are nonlinear in their parameters, which is why we call the framework in this case nonlinear least squares (NLS). The NLS approach is a widely used method for estimating the parameters of functional forms of the Lorenz curve, e.g., in (8, 15, 29–32). The objective we are trying to minimize is the sum of squared residuals. We recognize the estimation task as

$$\min_{\theta} \sum_{i=1}^K (L(u_i, \theta) - \eta_i)^2 \quad [2]$$

350 where  $\theta$  is the parameter vector of the Lorenz curve model and  $\eta_i$  the cumulative empirical income share observed for the  
351 cumulative population share  $u_i$ .

352 Using the NLS procedure, we get consistent estimates. However, they are not efficient, as least squares estimation in  
353 the Lorenz curve setting exhibits auto-correlated and heteroskedastic residuals (5, 8). Krause (2014) used the approach of  
354 minimizing the MSE in their recent study and mentions that other procedures to gain efficiency, e.g., proposed by (10), hardly  
355 change results given their setting.

356 A main disadvantage of NLS stems from ignoring the proportional nature of the data (33) and “overlook[ing] the fact that  
357 the sum of the income shares is, by definition, equal to one” (8, p. 11). Both features of the data are neglected by NLS and  
358 hence fruitful opportunities in using this special structure of the data are missed.

359 Apart from that, the NLS estimation method is still widely used for estimating Lorenz curves and does not provide efficient,  
360 but more importantly, consistent, estimates.

361 NLS estimates for each county are provided for the present study and will be evaluated as a robustness check.

## 362 10. Comparison of MLE and NLS Estimates

363 We explore whether potential estimation method artifacts account for our results by comparing the estimated parameters  
364 for the 17 Lorenz curve models using both NLS and MLE. We find similar point estimates for most model parameters. The  
365 median relative difference between the MLE and NLS estimates across counties is depicted in [Supplementary Table 8](#) below.  
366 An exception is the GB1 model, for which differences were large: for the generalized gamma and GB1 Lorenz curve model,  
367 the differences between NLS and MLE estimates were large, e.g., 84.1659 for the second GB1 parameter. This observation is  
368 not surprising, as those two Lorenz curve models exhibited severe estimation instabilities, which we take as indicating their  
369 unsuitability as a basis for deriving inequality measures. For this reason, we classify the GB1 model as unsuitable and exclude  
370 it from further analysis.

371 The remaining models exhibit small relative differences between both estimation methods. For example, the median relative  
372 difference between MLE and NLS point estimates of the Ortega parameters was 0.0234 for Ortega parameter  $\alpha$  and 0.0165 for  
373 Ortega parameter  $\beta$ . Hence, we have no reason to believe that the estimation technique has a systematic influence on the  
374 model parameters estimated.

375 We refer to the relative difference as given by

$$\text{relative difference} = \frac{|\hat{\theta}_{MLE} - \hat{\theta}_{NLS}|}{|\hat{\theta}_{NLS}|}$$

377 The median of the relative difference of parameter estimates across all  $N = 3\,056$  US counties included in our study is given  
378 in [Supplementary Table 8](#).

**Supplementary Table 8. Median relative difference between MLE and NLS estimates across all counties.**

Model	Param. 1	Param. 2	Param. 3	Param. 4	Param. 5
Abdalla-Hassan	0.0315	0.9410	0.0168	-	-
Chotikapanich	0.1498	-	-	-	-
Dagum	0.0605	0.0186	-	-	-
Gamma	0.2135	-	-	-	-
GB1	83.3003	84.1659	0.8906	-	-
GB2	0.2329	0.1884	0.1548	-	-
Generalized Gamma	80.8858	0.8900	-	-	-
Kakwani-Podder	0.2040	-	-	-	-
Lognormal	0.0165	-	-	-	-
Ortega	0.0234	0.0165	-	-	-
Pareto	0.0751	-	-	-	-
Rasche	0.0218	0.0145	-	-	-
Rhode	0.0250	-	-	-	-
Sarabia	0.3857	0.0292	0.0890	0.0680	-
Singh-Maddala	0.0718	0.0342	-	-	-
Wang	0.2122	0.1616	0.0894	0.4698	0.9285
Weibull	0.0810	-	-	-	-

## 11. Relationship between Ortega parameters and Pareto index

Sarabia et al. (1999) (34) introduced a general method to build ordered families of Lorenz curves, noting that one of the Pareto Lorenz curve families coincides with the Ortega Lorenz curve. We draw on this work in advancing the correspondence between the Pareto distribution parameter and one of the Ortega parameters.

To derive the relationship between Ortega parameter  $\beta$  and the Pareto index, let us first introduce some definitions. The Ortega Lorenz curve is given by (12):

$$L_{Ortega}(u) = u^\alpha \cdot (1 - (1 - u)^\beta) \quad [3]$$

where  $\alpha \leq 0, 0 < \beta \leq 1$ .

The cumulative distribution function of the classical Pareto distribution is given by

$$F(x) = 1 - \left(\frac{\sigma}{x}\right)^a \quad [4]$$

where  $\sigma, a > 0$ . Following this notation, we can recognize  $\sigma$  as a scale parameter and  $a$  as a shape parameter. The Pareto index equals the shape parameter of the classical Pareto distribution (e.g., used in (35)). Being consistent with our notation, we can therefore define

$$\text{Pareto index} := a \quad [5]$$

To show that there is a relationship between  $\beta$  and  $a$ , it is useful to calculate the Lorenz curve for the classical Pareto distribution first. The general definition of a Lorenz curve is given by (36):

$$L(u) = \mu^{-1} \int_0^u F^{-1}(t) dt \quad [6]$$

where  $\mu$  is the finite mean and  $F^{-1}(t)$  the inverse of the cumulative distribution function. For the classical Pareto case with  $\mu = \frac{a\sigma}{a-1}$  and  $F^{-1}(t) = \sigma(1-t)^{-\frac{1}{a}}$ , we get

$$L_{Pareto}(u) = \frac{a-1}{a\sigma} \int_0^u \sigma(1-t)^{-\frac{1}{a}} dt \quad [7]$$

$$= \frac{a-1}{a\sigma} \left[ \frac{-\sigma}{1-\frac{1}{a}} \cdot (1-t)^{1-\frac{1}{a}} \right]_0^u \quad [8]$$

$$= \frac{a-1}{a\sigma} \left[ \left( \frac{-\sigma}{1-\frac{1}{a}} \cdot (1-u)^{1-\frac{1}{a}} \right) - \left( \frac{-\sigma}{1-\frac{1}{a}} \right) \right] \quad [9]$$

$$= \left(1 - \frac{1}{a}\right) \cdot \left[ \frac{-1}{1-\frac{1}{a}} (1-u)^{1-\frac{1}{a}} + \frac{1}{1-\frac{1}{a}} \right] \quad [10]$$

$$= 1 - (1-u)^{1-\frac{1}{a}} \quad [11]$$

We can see that the Pareto Lorenz curve depends on the Pareto index  $a$  only. If we are able to relate the Pareto Lorenz curve to the Ortega Lorenz curve and demonstrate that the Pareto index is linked to one of the two Pareto parameters only, we know that we can transform that parameter into the Pareto index. (34) actually introduced a family of Lorenz curves that helps explain the relationship between the Pareto Lorenz curve and the Ortega Lorenz curve. In detail, their second theorem states:

**Theorem 2 ((34))** *Let  $L(p)$  be a Lorenz curve and consider the transformation  $L_\alpha(p) = p^\alpha \cdot L(p)$ , where  $\alpha \leq 0$ . Then, if  $\alpha \geq 1$ ,  $L_\alpha(p)$  is a Lorenz curve too. In addition, if  $0 \leq \alpha < 1$  and  $L'''(p) \geq 0$ ,  $L_\alpha(p)$  is also a Lorenz curve.*

(34) further show that the condition  $L'''(p)$  is satisfied for the Pareto Lorenz curve such that for  $\alpha \geq 0$ , we can transform the Pareto Lorenz curve using theorem 2, which yields

$$L_\alpha(u) = u^\alpha \cdot L_{Pareto}(u) \quad [12]$$

$$= u^\alpha \cdot \left(1 - (1-u)^{1-\frac{1}{a}}\right) \quad [13]$$

Now looking at the Ortega Lorenz curve as defined in 3, we can clearly see that the Ortega Lorenz curve is nothing other than the Pareto Lorenz curve, extended by a newly introduced parameter  $\alpha$  through the use of theorem 2 and a redefined parameter

$$\beta := 1 - \frac{1}{a} \quad [14]$$

In other words, we can see the Ortega Lorenz curve as an extension to the Pareto Lorenz curve. Having established this close link between the two Lorenz curves, we can think of Ortega parameter  $\beta$  as being in close relation to the Pareto index  $a$ , using the relationship defined in 14. If the true income distribution were to follow a Pareto distribution, Ortega parameter  $\alpha$  would be zero and the Ortega parameter  $\beta$  would be an exact monotonic transformation of the Pareto index. However, in cases where the true income distribution was not generated by a Pareto distribution, of course, the additional estimation of Ortega parameter  $\alpha$  might capture aspects that are also correlated to  $\beta$ , such that the exact monotonic transformation given in 14

411 is rather an approximate relationship, depending on the data. Although this is a weaker statement, it is still useful for our  
412 purpose: we want to know which aspects of the income distribution the Ortega parameters capture. We know that the lower  
413 the Pareto index, the larger the proportion of very-high-income people. And we derived above that the higher the Pareto  
414 index associated with the income distribution, the higher the Ortega parameter  $\beta$ . Having demonstrated the close relationship  
415 between  $\beta$  and the Pareto index  $a$  in the above section, we see this as evidence of  $a$  capturing the occurrence of very top  
416 incomes. We therefore conclude that Ortega  $\beta$  has the following interpretation: the lower the Pareto index, the larger the  
417 proportion of very-high-income people. We therefore propose it as a measure of top-concentrated income inequality.

## 418 12. Interpreting the Ortega Lorenz curve

419 **Visual inspection of Ortega parameters.** To visually inspect how a change in parameters affects the Ortega Lorenz curve,  
420 we simulate Ortega Lorenz curves while varying  $\alpha$  and  $\gamma$ . The R code `simulation_ortega_lorenz_curves.R` replicates this  
421 simulation and is available in the GitHub repository we provide for this paper (see [www.measuringinequality.com](http://www.measuringinequality.com)). In detail,  
422 first we plot the Ortega Lorenz curves varying  $\alpha$  between 0.01 and 1.5 while keeping  $\gamma$  fixed at 0.5 (for  $\alpha$ , the side constraint is  
423  $\geq 0$ ; the upper limit 1.5 is chosen as an extension to the empirical values that valued 1.23 at max). In our empirical estimation  
424 of US county-level Ortega Lorenz curves, for  $\alpha$  a typical value was 0.5 and for  $\gamma$  0.5, which is why we fix the respective values  
425 at that level. Then, we plot Ortega Lorenz curves with  $\alpha = 0.5$  and vary  $\gamma$  between 0.01 and 0.99 (side constraint  $0 \leq \gamma < 1$ ).

426 Our simulation results are generally in line with prior theory, i.e., that Ortega parameter  $\gamma$  is associated with top-concentrated  
427 inequality. The asymmetry line in Figure 3 in the main text of the paper, Panel B, facilitates comprehension whereby we  
428 observe a disproportionate change in the Lorenz curve on the right side (i.e., at higher incomes). Note that we observe  
429 top-concentrated inequality arising when there is a step increase in the Lorenz curve shortly ahead of the cumulative share of  
430 population reaching 100%. Further, our observations are in accordance with the direction of change we expected through the  
431 relationship between  $\gamma$  and the Pareto index, i.e., a higher value of  $\gamma$  indicating a higher level of top-concentrated inequality.

432 In sum, our simulation study suggests that  $\alpha$  is a reflection of bottom-concentrated inequality whereas  $\gamma$  is a reflection of  
433 top-concentrated inequality.

434 When varying  $\alpha$  while keeping  $\gamma$  a fixed constant, we can see that an increase in  $\alpha$  stretches the left side of the Lorenz curve  
435 toward the x-axis (i.e., at lower incomes). The higher  $\alpha$ , the more this is the case, as seen in Figure 3A in the main text. This  
436 effect can again be acknowledged when adding the asymmetry line to the plot, which helps in identifying the disproportionate  
437 change in the curves. With a more intense change on the left side, one can conclude that  $\alpha$  captures specificities on the left tail  
438 of the income distribution.<sup>||</sup> Therefore, we conclude that  $\alpha$  is a measure of bottom-concentrated inequality.

439 **Determining the relationship between Ortega parameters and other measures of inequality.** To further investigate the interpre-  
440 tation of the Ortega parameters, we relate them to income ratios, as they are more intuitive and used in some prior research to  
441 measure inequality. First, we explore the dependency between Ortega parameters and common percentile measures (95/50 and  
442 50/10 ratios). Then, we move on to evaluate which percentile ratios might reflect the information captured by the Ortega  
443 parameters more precisely.

444 A common measure of top-concentrated income inequality is the fraction of income held by the 95<sup>th</sup> percentile divided by  
445 the median income share (also known as a 95/50 ratio), whereas bottom-concentrated income is often measured using a 50/10  
446 ratio; see (37–39). We have argued that Ortega parameter  $\gamma$  is related to top-concentrated inequality and should increase with  
447 higher levels of inequality. The 95/50 ratio also aims at capturing the phenomenon of top-concentrated income inequality, which  
448 is why we suspect the quantities to be highly positively correlated. We also hypothesized that Ortega parameter  $\alpha$  is related to  
449 bottom-concentrated income inequality and should increase with higher levels of inequality. Another measure that aims at  
450 capturing bottom inequality is the 50/10 ratio, i.e., the income share held by the lower 50% of the population divided by the  
451 income share held by the lower 10% of the population. We suspect that both quantities, i.e.,  $\alpha$  and the 50/10 ratio, should be  
452 highly positively correlated because they should measure the same underlying phenomenon (bottom-concentrated inequality).

453 To test whether our suggested correlational dependencies hold true, we first simulate Ortega Lorenz curves with varying  
454 parameters  $\alpha$  and  $\gamma$ , then calculate the income percentile ratios 95/50 and 50/10 for those Lorenz curves, and consequently  
455 analyze the correlation between Ortega and percentile ratio quantities. In detail, we simulate a total of 10 000 Ortega Lorenz  
456 curves with varying parameter values. We vary  $\alpha$  from 0.01 to 1 with a step size of 0.01 and  $\gamma$  from 0 to 0.99 with the same  
457 step size of 0.01. Subsequently, we calculate partial correlations between the quantities. Doing so, we control for all other  
458 variables included in this analysis; i.e., we correlate  $\alpha$  with the 50/10 ratio controlling for  $\gamma$  and the 95/50 ratio.

459 Our results, depicted in [Supplementary Table 9](#), show that  $\alpha$  indeed highly correlates with the bottom-concentration ratio  
460 50/10 while  $\gamma$  highly correlates with the top-concentration ratio 95/50. However, it is worth pointing out that this correlational  
461 dependency only becomes apparent when focusing on the full parameter space of  $\gamma$  ( $0 \leq \gamma < 1$ ) while limiting the parameter  
462 space of  $\alpha$  for the same range as  $\gamma$ . For the empirical US county-level Lorenz curves, we encountered a parameter range of 0.12  
463 to 1.23 for  $\alpha$  and 0.3 to 0.93 for beta. In this range of parameters, the correlation between  $\alpha$  and the 50/10 ratio, and  $\gamma$  and  
464 the 95/50 ratio, gets distorted, which indicates high sensitivity of the correlational structure regarding the parameter range.

465 This gives us reason to believe that those ratios might not reflect the type of top- and bottom-concentrated inequality that  
466 is measured by the Ortega parameters. Revising Figure 3 in the main text, we can see  $\gamma$  affecting rather the very top of the  
467 distribution. Exploring the dependency structure percentile ratios and the Ortega parameters, it indeed becomes clear that

<sup>||</sup> A high level of bottom-concentrated inequality can be recognized from the Lorenz curve if the curve is rather flat near the bottom percentiles but exhibits a sharp increase before reaching the median population.



468 Ortega  $\gamma$  is instead measuring inequality in the very top percentiles and that  $\alpha$  captures a broader range of the distribution. We  
 469 find the correlational dependency between the 99/90 ratio with  $\gamma$  and 90/10 ratio with  $\alpha$  very robust to the parameter range.  
 470 Also, the strength of correlational dependency is more distinct; see [Supplementary Table 10](#), which depicts the correlations  
 471 within the same parameter range used for Lorenz curve generation as in [Supplementary Table 9](#).


472 We therefore conclude that our suggested interpretation of the Ortega parameters should not directly be linked to current  
 473 measures of top- and bottom-concentrated inequality, i.e., the 95/50 and 50/10 ratios, but to measures of inequality at the very  
 474 top (99/90 ratio) and most of the remainder of the distribution (90/10 ratio).

**Supplementary Table 9. Partial correlations between Ortega parameters and percentile ratios, controlling for all other quantities; e.g., the partial correlation between  $\gamma$  and the 95/50 ratio is 0.940 after controlling for  $\alpha$  and the 50/10 ratio.**

	50/10 ratio	95/50 ratio
Ortega $\alpha$	0.786	0.137
Ortega $\gamma$	-0.259	0.940

**Supplementary Table 10. Partial correlations between Ortega parameters and percentile ratios, controlling for all other quantities; e.g., the partial correlation between  $\gamma$  and the 99/90 ratio is 0.9088 after controlling for  $\alpha$  and the 90/10 ratio.**

	90/10 ratio	99/90 ratio
Ortega $\alpha$	0.9081	-0.0408
Ortega $\gamma$	-0.0620	0.9088



**Supplementary Figure 14.** Panel A illustrates two very different Lorenz curves exhibiting the same 90/50 percentile ratio. In Panel B we can notice that when fixing both the 90/50 and the 50/10 percentile ratios into a similar range, the resulting Lorenz curves must have a similar shape. This indicates that (at least) two parameters should be provided to limit the potential volatility of the resulting Lorenz curves.

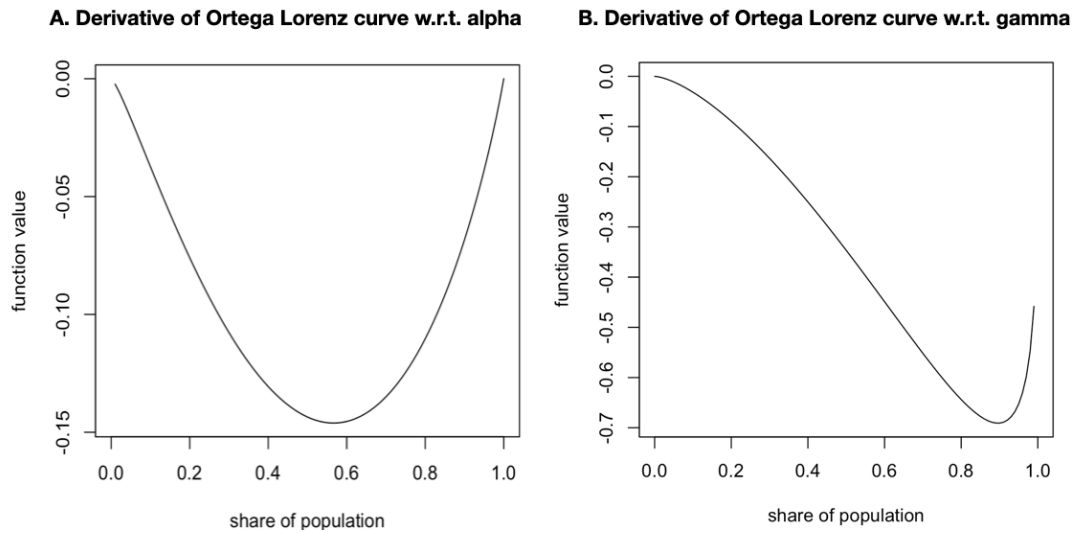
475 **Analytical investigation of the Ortega Lorenz curve: Derivatives.** A natural way to investigate how a function is affected by its  
 476 parameters is to inspect the (partial) derivatives. For the Ortega Lorenz curve, the partial derivatives with respect to  $\alpha$  and  $\gamma$   
 477 are

$$\frac{\delta}{\delta\alpha} \left( u^\alpha (1 - (1 - u)^{1-\gamma}) \right) = (u^\alpha (1 - (1 - u)^{1-\gamma}) \log(u)) \quad [15]$$

$$\frac{\delta}{\delta\gamma} \left( u^\alpha (1 - (1 - u)^{1-\gamma}) \right) = (u^\alpha (1 - u)^{1-\gamma} \log(1 - u)) \quad [16]$$

478 From this it is not immediately obvious how the Ortega Lorenz curve is affected by the parameters. However, we can  
 479 note that both derivatives are  $\leq 0$  within the allowed parameter space. What we are especially interested in is whether  
 480 the interpretation of the parameters suggested by the simulation study ( $\alpha$  more intensely emphasizing bottom-concentrated

481 inequality and  $\gamma$  highlighting top-concentrated inequality) can be seen analytically as well. To test this, we take a closer look  
 482 at the rate of change, i.e., the partial derivatives, at certain regions along the x-axis. In other words, if the Lorenz curve  
 483 function is more intensely affected by a parameter in a certain region of the population, we could conclude that this parameter  
 484 is more sensitive to this area of the population: e.g., the top or bottom. [Supplementary Figure 15](#) visualizes the derivatives of  
 485 the Ortega Lorenz curve with respect to  $\alpha$  and  $\gamma$  along the x-axis (i.e., cumulative share of population) while keeping the  
 486 parameters themselves fixed at  $\alpha = 0.5, \gamma = 0.5$ , just as when simulating Ortega Lorenz curves in the above section. Note that  
 487 we need to evaluate the absolute values of rate of change for the respective parameters, i.e., the absolute values of the partial  
 488 derivatives. From [Supplementary Figure 15](#), we can clearly see that a variation in  $\alpha$  most intensely affects the Lorenz curve  
 489 around the middle of the population (the absolute value of the derivative with respect to  $\alpha$  is largest around the percentiles  $\sim$   
 490 0.45-0.65). In contrast, a variation in  $\gamma$  has the highest rate of change within the top percentile of the population (the absolute  
 491 value of the derivative with respect to  $\gamma$  is largest around the top percentiles  $\sim$  0.80-0.95).

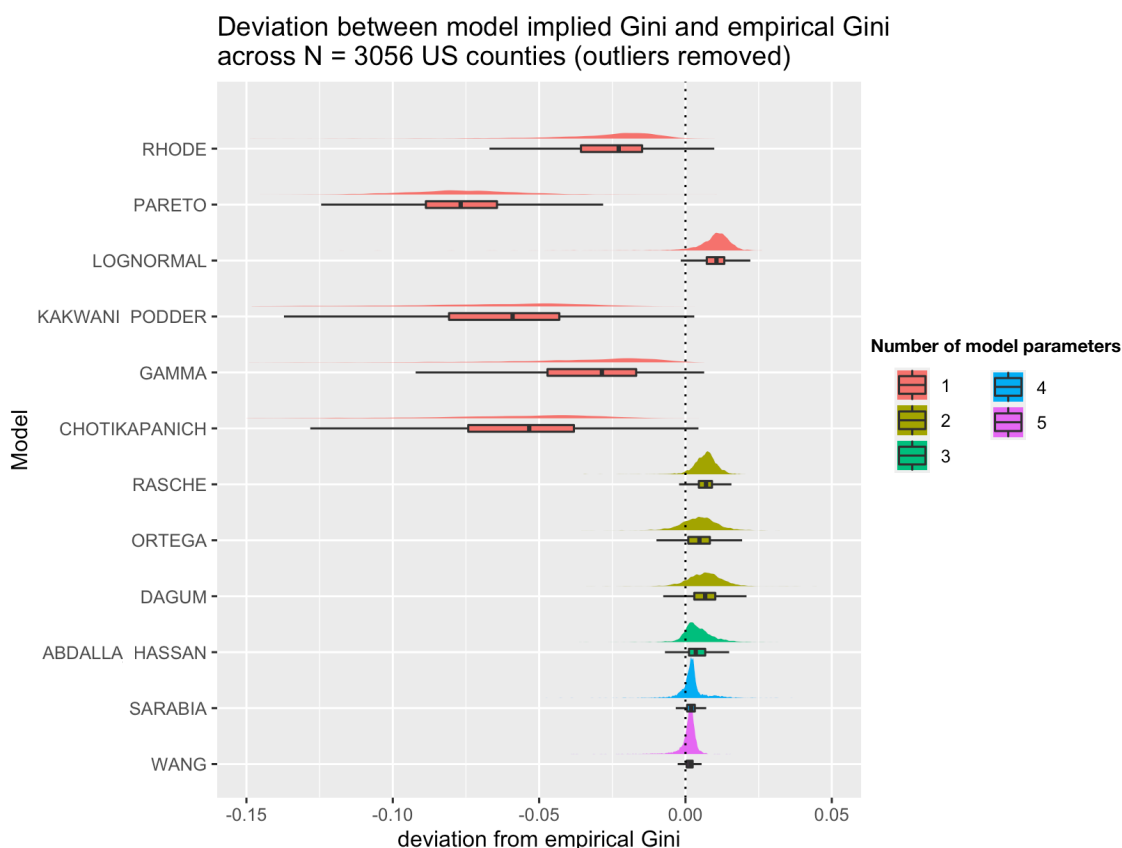


**Supplementary Figure 15.** Value of the derivatives of the Ortega Lorenz curve function  $L(u) = u^{0.5}(1 - (1 - u)^{0.5})$  across the cumulative share of population.

492 **13. Approximating the empirical Gini coefficient**

493 To assess how well the different models approximate the main distributional statistics related to inequality, we compare the  
 494 Gini coefficients implied by the model parameters with those Gini coefficients calculated nonparametrically on the US county  
 495 data. The nonparametric Gini coefficients are calculated using the given data points of the empirical income distribution with  
 496 linear interpolation, whereas the Gini coefficients implied by the models utilize integral calculus\*\* for determining the area  
 497 between the Lorenz curve and the line of perfect inequality.

498 These analyses, visualized in [Supplementary Figure 16](#), reveal that when taking into account the number of parameters  
 499 included in the model—ideally as few as possible—we can see that the Ortega model provides a reasonable trade-off between  
 500 deviation from the nonparametric Gini and the number of parameters needed. Most notably, one-parameter models (red  
 501 distributions in the figure) substantially deviate from the ideal average deviation of zero, while two-parameter models (brown)  
 502 are a major improvement. Across the two-parameter models, the Ortega model is the one closest to the deviation of zero  
 503 (dotted line) with a substantial number of data points (see boxplot touching the dotted line). While with more parameters  
 504 (green, blue, and purple boxplots), precision further increases, the improvements are much smaller than those between one-  
 505 and two-parameter models. This analysis demonstrates that using more than one parameter improves the approximation of  
 506 empirical distributional statistics such as the Gini coefficient, and that further improvement in precision with more parameters  
 507 is possible but is much smaller.



**Supplementary Figure 16.** Comparison across various parametric Lorenz curve models in approximating the empirical (nonparametric) Gini coefficient. Note that in order to prevent a masking effect of severe outliers, we omitted them in the plot. The boxes depict the 25<sup>th</sup>, 50<sup>th</sup> and 75<sup>th</sup> percentiles of the deviations from the empirical Gini. The whiskers extend from the hinge to the smallest value at most (or largest value and no further, respectively) 1.5 times the inter-quartile range of the hinge. Minimum and maximum values as well as the center of the distributions are visualized by plotting the actual distribution of deviations above the boxes.

\*\* For the Lorenz curve models based on the generalized beta distribution (GB1, GB2), we faced difficulties in calculating the integrals necessary for parametric Lorenz curve derivation, which is why these models are missing in our analysis.

508 **14. Exploratory correlational study**

509 In our exploratory correlational study, for which we provide results below, we correlate 100 variables from policy-relevant fields  
 510 to inequality measures. Our source of data is the ACS Survey 2011-2015, from which we pulled relevant source tables directly  
 511 from <https://data2.nhgis.org/main>, and the data from (40) and (41) are publicly available at <https://opportunityinsights.org>. Code to  
 512 replicate the study, as well as detailed information on the data used—i.e., a codebook—is available at [www.measuringinequality.com](http://www.measuringinequality.com).  
 513 com.

514 We propose the use of both the Ortega parameters simultaneously (i.e., in a regression setting, researchers should include  
 515 both Ortega parameters as independent variables within the regression model equation), which is why we calculate partial  
 516 Pearson correlations between covariates and Ortega parameters. For the Gini coefficient, simple Pearson correlations are  
 517 sufficient, as this is a single-parameter inequality measurement approach. We use the Gini coefficient provided by the ACS  
 518 dataset. One might argue that we should have used the Gini index implied by the empirical Lorenz curves we used in the Ortega  
 519 parameter estimation. However, the US Census Bureau, which conducts the ACS, has more fine-grained data (inaccessible to  
 520 the public) available to calculate the Gini index for each county highly accurately, which makes their Gini indices more reliable.


521 In [Supplementary Table 11](#), we provide an overview of potential outcomes and the frequency of their occurrence across our  
 522 analysis. Case ID 1 can be interpreted as Ortega’s ability to disentangle (probably counteracting) effects related to inequality  
 523 present in different parts of the income distribution, and case ID 2 might also shed light on a specific region of the income  
 524 distribution being correlated to policy outcomes. For case ID 3, i.e., that neither Gini nor Ortega parameters show significant  
 525 correlations, we have a coherent suggestion from both inequality measures that there is no association between inequality and  
 526 the correlated variable. We also find coherent guidance on whether inequality is associated with a variable for case IDs 4 and 5.  
 527 However, these cases show that use of the Ortega parameter might refine the insights we can obtain: while the Gini only reveals  
 528 that there is an association between overall inequality and the variable, using the Ortega parameters, we can differentiate which  
 529 part of the income distribution drives the significant correlation, including the magnitude. For case ID 6, i.e., that Gini is  
 530 significant but none of the Ortega parameters are, the interpretation of such cases is rather puzzling. A potential interpretation  
 531 is that in such cases, the association between inequality and the variable is driven by a feature of inequality that is captured  
 532 through the Gini coefficient measuring overall inequality but is not explained by the concentration of income in different parts  
 533 of the income distribution.

**Supplementary Table 11. Cases occurring in our exploratory study correlating 100 covariates with the Gini index and calculating partial correlations between covariates and Ortega parameters.**

Case ID	Correlation with Gini coefficient $\neq 0$	Correlation with ... Ortega parameters $\neq 0$	Number of occurrences
1	no	2	12
2	no	1	21
3	no	0	8
4	yes	1	25
5	yes	2	34
6	yes	0	0
			100 = total number of covariates


### Correlation and CI for Inequality Measures with a Variety of Aspects

Confidence level: 0.9995 (using a Bonferroni Correction)




**Supplementary Figure 17.** Pearson correlations between inequality measures and county-level covariates. The plot shows Pearson correlations with the Gini index and partial Pearson correlations with the Ortega parameters, i.e., the correlation between one Ortega parameter and the covariate while controlling for the other Ortega parameter across  $N = 3\,049$  US counties. Pearson correlation point estimates are visualized within confidence bounds of the Bonferroni corrected confidence interval.

Case 1: Gini = 0 and both Ortega parameters are  $\neq 0$   
 Confidence level: 0.9995 (using a Bonferroni Correction)




**Supplementary Figure 18.** The plot shows Pearson correlations for instances of case ID 1 (see [Supplementary Table 11](#)) with the Gini index and partial Pearson correlations with the Ortega parameters, i.e., the correlation between one Ortega parameter and the covariate while controlling for the other Ortega parameter across  $N = 3\,049$  US counties. Pearson correlation point estimates are visualized within confidence bounds of the Bonferroni corrected confidence interval.

Case 2: Gini = 0, and exactly one Ortega  $\neq 0$   
 Confidence level: 0.9995 (using a Bonferroni Correction)



**Supplementary Figure 19.** The plot shows Pearson correlations for instances of case ID 2 (see [Supplementary Table 11](#)) with the Gini index and partial Pearson correlations with the Ortega parameters, i.e., the correlation between one Ortega parameter and the covariate while controlling for the other Ortega parameter across  $N = 3\,049$  US counties. Pearson correlation point estimates are visualized within confidence bounds of the Bonferroni corrected confidence interval.


Case 3: Gini = 0, Ortega\_1 = 0, Ortega\_2 = 0  
 Confidence level: 0.9995 (using a Bonferroni Correction)



**Supplementary Figure 20.** The plot shows Pearson correlations for instances of case ID 3 (see [Supplementary Table 11](#)) with the Gini index and partial Pearson correlations with the Ortega parameters, i.e., the correlation between one Ortega parameter and the covariate while controlling for the other Ortega parameter across  $N = 3\,049$  US counties. Pearson correlation point estimates are visualized within confidence bounds of the Bonferroni corrected confidence interval.




Case 4: Gini  $\neq 0$ , and exactly one Ortega  $\neq 0$   
 Confidence level: 0.9995 (using a Bonferroni Correction)



**Supplementary Figure 21.** The plot shows Pearson correlations for instances of case ID 4 (see [Supplementary Table 11](#)) with the Gini index and partial Pearson correlations with the Ortega parameters, i.e., the correlation between one Ortega parameter and the covariate while controlling for the other Ortega parameter across  $N = 3\,049$  US counties. Pearson correlation point estimates are visualized within confidence bounds of the Bonferroni corrected confidence interval.

Case 5: Gini != 0, and both Ortega are !=0  
 Confidence level: 0.9995 (using a Bonferroni Correction)



**Supplementary Figure 22.** The plot shows Pearson correlations for instances of case ID 5 (see [Supplementary Table 11](#)) with the Gini index and partial Pearson correlations with the Ortega parameters, i.e., the correlation between one Ortega parameter and the covariate while controlling for the other Ortega parameter across  $N = 3\,049$  US counties. Pearson correlation point estimates are visualized within confidence bounds of the Bonferroni corrected confidence interval.

## 15. Simulation Study: Minimum Dataset Requirements

We introduce and evaluate **three key criteria** that datasets for inequality estimation need to possess in order for us to include them in this systematic “tournament-style” comparison to identify the best-fitting inequality measure given empirical income distributions. We find that such datasets need to contain **(1)** at least 15 or more data points per Lorenz curve; **(2)** at least two data points on top income shares above the 90th percentile of the income distribution; and **(3)** at least 60 Lorenz curves—and ideally, many more. We conducted numerous simulation studies, outlined in this section, to estimate these requirements.

In the simulation study on data granularity in the SI, Section 10, we found that for a sufficient granularity (15+ data points), and in the absence of noise, the MLE procedure will detect the correct model in almost every case if it was generated by an Ortega model (>98% of cases; see Supplementary Table 4). However, empirical observations contain observational noise. Is the AICc procedure for a given granularity of, say, 20 data points—in the presence of observational noise—still able to detect Ortega? In this case, the number of Lorenz curves available becomes crucial; i.e., if the number of Lorenz curves is too small, the reduced certainty in detecting Ortega via AICc for each Lorenz curve could lead to a false overall conclusion. But how many Lorenz curves are necessary to reduce uncertainty to reasonable amounts?

We quantify uncertainty in deciding the correct model for a given number of Lorenz curves ( $N$ ) by considering each of the  $N$  Lorenz curves as independent draws from some Ortega Lorenz curve. Mathematically speaking, we can see AICc’s chance of success for detecting Ortega in each of the  $N$  Lorenz curves in terms of a Bernoulli distributed variable, i.e., AICc either detects Ortega (success = 1) or not (no success = 0). From this perspective, we can interpret the Bernoulli parameter  $p$  (probability of success) as the expected percentage of Ortega detections. For  $N$  Lorenz curves, we would expect to detect  $p \cdot N$  Lorenz curves as Ortega. Note that for simplicity, we assume the researcher decides for Ortega if it is detected in the majority of cases; hence we require  $p > 0.5$ .

The crucial point of  $N$  is that the percentage of Ortega detections, which corresponds to the maximum likelihood estimate of Bernoulli parameter  $p$ , will approximate the true value of  $p$  more accurately with increasing  $N$ : variation in estimated  $p$  across sample sizes  $N$  is the actual quantity we are interested in when quantifying the uncertainty of determining the correct model overall. We can derive the variance of this estimator analytically; i.e.,



$$\text{Var}(\hat{p}) = \frac{p(1-p)}{N} \quad [17]$$

For the simulation, we vary the number of  $N$  Lorenz curves to be generated from some underlying Ortega Lorenz curve model, allowing for each of the  $N$  samples to exhibit different Ortega parameters, and a small normally distributed random noise term (mean = 0, sd = 0.002) to reflect observational noise. We then use our MLE procedure to fit various Lorenz curve models, let AICc determine the optimum model, and divide the number of detected Ortega models by  $N$  to get an estimate for  $p$ . Repeating this procedure 10 000 times gives us an estimate for the empirical standard deviation of estimated  $p$ , i.e., the standard deviation in the percentage of correctly classified Lorenz curves.

Our results show that with increased sample size  $N$ , the standard deviation of the percentage of correct model detections decreases; critically, we show that at least 60 Lorenz curves are necessary to ensure that the share of correctly classified Lorenz curves is above 50%; see [Supplementary Figure 23](#). When fewer than 60 Lorenz curves are available, the identification of the correct model is below 50%, reflecting the challenges of using datasets that contain fewer Lorenz curves, in line with criterion #3.

In this simulation setup, we can further analyze the effects of sparse top-income data. In the base setting, we use equidistant population data shares with fixed granularity level (20 data points including population levels 0 and 1), i.e., a case where we have as much information on top-income shares as on any other parts of the income distribution. We compare this with a case where we have sparser information on top-income shares: we use the same granularity of 20 data points, but now these data points are shifted on the x-axis of the Lorenz curve toward the bottom of the income distribution, resulting in a lack of information on the top income percentiles. For example, if 1 out of the 20 data points is above the 90th percentile, this means that we have information on the bottom 90% of income earners and the 95th percentile, whereas in the case of 3 out of 20 data points being above the 90th percentile, we would have information on the bottom 90% of income earners and the 92.5th, 95th, and 97.5th percentiles. We see a considerable increase in the average percentage of true model detection as more information on top income earners is available; see [Supplementary Figure 24](#). When fewer than two data points on top-income earners above the 90th percentile are available, the share of correctly identified models again drops below 50%, in line with criterion #2. Note that the number of Lorenz curves becomes irrelevant in this case: a higher number of Lorenz curves that do not contain top-income information do not improve our selection of the overall best-fitting model, given that  $p = 0.4 < 0.5$  even when the estimated  $p$  converges with a large  $N$ . This analysis additionally reveals that our three criteria can not be treated separately but must be considered jointly.

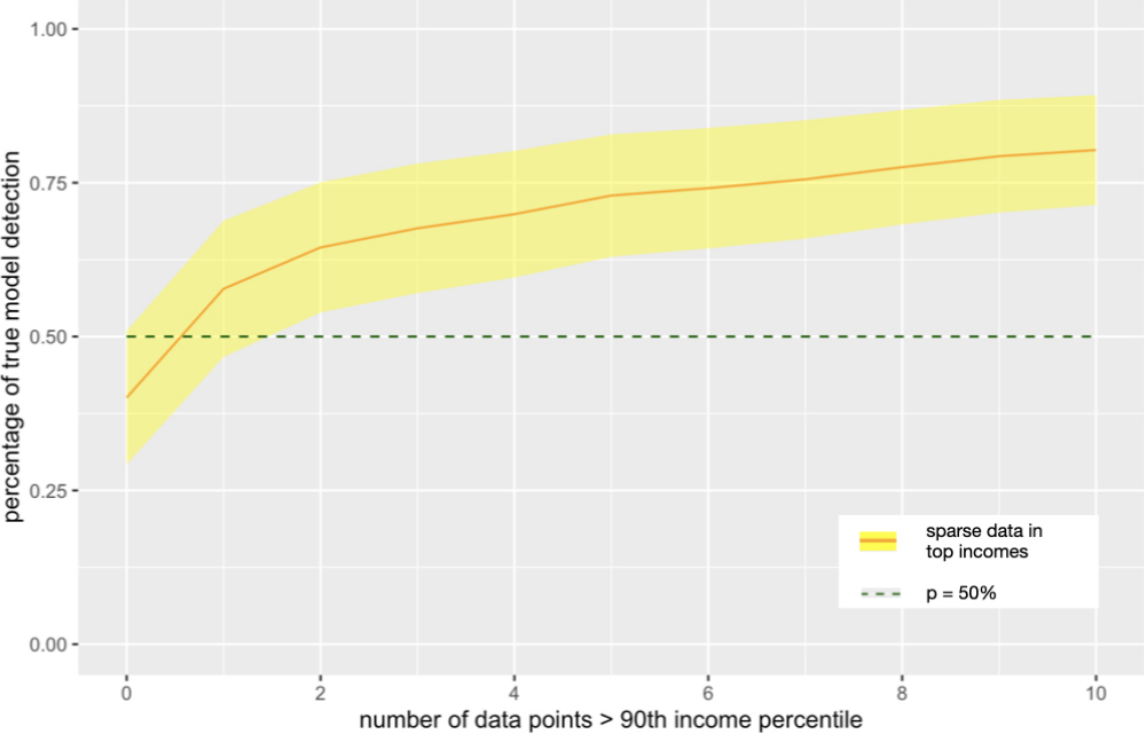
Uncertainty in true model detection for varying sample sizes with observational noise  
Average model detection rate and its standard deviation across simulation runs



Supplementary Figure 23. Uncertainty in true model detection: Variation in sample size

Average model detection rate for varying information density in top incomes

Fixed data granularity (20 data points), N = 20 Lorenz curves, average and standard deviation across 10,000 simulation runs



Supplementary Figure 24. Uncertainty in true model detection: Variation in information density within top incomes

580 **References**

- 581 1. J Fellman, Income inequality measures. *Theor. Econ. Lett.* **08**, 557–574 (2018).  
582 2. N Kakwani, On a class of poverty measures. *Econometrica* **48**, 437–446 (1980).  
583 3. R Basmann, K Hayes, D Slottje, J Johnson, A general functional form for approximating the lorenz curve. *J. Econom.* **43**,  
584 77 – 90 (1990).  
585 4. BC Arnold, JM Sarabia, *Majorization and the Lorenz Order with Applications in Applied Mathematics and Economics*.  
586 (Springer International Publishing), 1 edition, (2018).  
587 5. M Krause, Parametric lorenz curves and the modality of the income density function. *Rev. Income Wealth* **60**, 905–929  
588 (2014).  
589 6. C Kleiber, S Kotz, A characterization of income distributions in terms of generalized gini coefficients. *Soc. Choice Welf.*  
590 **19**, 789–794 (2002).  
591 7. C Dagum, Wealth distribution models: Analysis and applications. *Statistica* **3** (2006).  
592 8. V Jorda, JM Sarabia, M Jantti, Estimation of income inequality from grouped data. arxiv working paper 1808.09831  
593 (2018).  
594 9. J McDonald, Some generalized functions for the size distribution of income. *Econometrica* **52**, 647–63 (1984).  
595 10. NC Kakwani, N Podder, On the estimation of lorenz curves from grouped observations. *Int. Econ. Rev.* **14**, 278–292  
596 (1973).  
597 11. RH Rasche, J Gaffney, AYC Koo, N Obst, Functional forms for estimating the lorenz curve. *Econometrica* **48**, 1061–1062  
598 (1980).  
599 12. P Ortega, G Martín, A Fernández, M Ladoux, A García, A new functional form for estimating lorenz curves. *Rev. Income*  
600 *Wealth* **37**, 447–452 (1991).  
601 13. D Chotikapanich, A comparison of alternative functional forms for the lorenz curve. *Econ. Lett.* **41**, 129 – 138 (1993).  
602 14. JM Sarabia, E Castillo, DJ Slottje, An ordered family of lorenz curves. *J. Econom.* **91**, 43 – 60 (1999).  
603 15. IM Abdalla, MY Hassan, Maximum likelihood estimation of lorenz curves using alternative parametric model. *Metodoloski*  
604 *Zvezki* **1**, 109–118 (2004).  
605 16. N Rohde, An alternative functional form for estimating the lorenz curve. *Econ. Lett.* **105**, 61 – 63 (2009).  
606 17. Z Wang, YK Ng, R Smyth, A general method for creating lorenz curves. *Rev. Income Wealth* **57**, 561–582 (2011).  
607 18. U.S. Census Bureau, American community survey, 2011–2015: 5-year period estimates in <https://data2.nhgis.org/main>.  
608 (United States of America), (2016).  
609 19. E Sommeiller, M Price, E Wazeter, Income inequality in the us by state, metropolitan area, and county, Technical report  
610 (2016).  
611 20. FA Farris, The gini index and measures of inequality. *The Am. Math. Mon.* **117**, 851–864 (2010).  
612 21. D Chotikapanich, WE Griffiths, Estimating lorenz curves using a dirichlet distribution. *J. Bus. & Econ. Stat.* **20**, 290–295  
613 (2002).  
614 22. AC Chang, P Li, SM Martin, Comparing cross-country estimates of lorenz curves using a dirichlet distribution across  
615 estimators and datasets. *J. Appl. Econom.* **33**, 473–478 (2018).  
616 23. K Aho, D Derryberry, T Peterson, Model selection for ecologists: the worldviews of aic and bic. *Ecology* **95**, 631–636  
617 (2014).  
618 24. AR Liddle, Information criteria for astrophysical model selection. *Mon. Notices Royal Astron. Soc. Lett.* **377**, L74–L78  
619 (2007).  
620 25. SJ Brams, PC Fishburn, Chapter 4 voting procedures in *Handbook of Social Choice and Welfare*, Handbook of Social  
621 Choice and Welfare. (Elsevier) Vol. 1, pp. 173 – 236 (2002).  
622 26. KJ Arrow, A difficulty in the concept of social welfare. *J. political economy* **58**, 328–346 (1950).  
623 27. AE Raftery, Bayesian model selection in social research. *Sociol. Methodol.* **25**, 111–163 (1995).  
624 28. KP Burnham, DR Anderson, Multimodel inference: Understanding aic and bic in model selection. *Sociol. Methods & Res.*  
625 **33**, 261–304 (2004).  
626 29. E Belz, Estimating Inequality Measures from Quantile Data, (Center for Research in Economics and Management (CREM),  
627 University of Rennes 1, University of Caen and CNRS), Economics Working Paper Archive (University of Rennes 1 &  
628 University of Caen) 2019-09 (2019).  
629 30. S Md, C Saeki, A new functional form for estimating lorenz curves. *J. Bus. Econ. Res.* **1** (2011).  
630 31. KS Cheong, An empirical comparison of alternative functional forms for the lorenz curve. *Appl. Econ. Lett.* **9**, 171–176  
631 (2002).  
632 32. JM Sarabia, V Jorda, C Trueba, The lamé class of lorenz curves. *Commun. Stat. - Theory Methods* **46**, 5311–5326 (2017).  
633 33. D Chotikapanich, W Griffiths, Averaging Lorenz curves. *The J. Econ. Inequal.* **3**, 1–19 (2005).  
634 34. JM Sarabia, E Castillo, D Slottje, An ordered family of lorenz curves. *J. Econom.* **91**, 43–60 (1999).  
635 35. BC Arnold, *Pareto Distribution*. (American Cancer Society), pp. 1–10 (2015).  
636 36. JL Gastwirth, A general definition of the lorenz curve. *Econometrica* **39**, 1037–1039 (1971).  
637 37. RJ Gordon, I Dew-Becker, Selected Issues in the Rise of Income Inequality. *Brookings Pap. on Econ. Activity* **38**, 169–192  
638 (2007).  
639 38. S Voitchovsky, Does the profile of income inequality matter for economic growth? *J. Econ. growth* **10**, 273–296 (2005).  
640 39. SF Reardon, K Bischoff, Income inequality and income segregation. *Am. journal sociology* **116**, 1092–1153 (2011).

- 641 40. R. Chetty, et al., The association between income and life expectancy in the united states, 2001-2014. *Jama* **315**, 1750–1766  
642 (2016).
- 643 41. R. Chetty, N. Hendren, The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates. *The Q. J.*  
644 *Econ.* **133**, 1163–1228 (2018).