University of Cyprus
Department of Economics

Working Paper 04-2022

Lognormal (Re)Distribution: A Macrofounded Theory of Inequality

Jon Eguia and Dimitrios Xefteris

Department of Economics, University of Cyprus, P.O. Box 20537, 1678 Nicosia, Cyprus
Web site: http://www.ucy.ac.cy/econ/en

Lognormal (Re)Distribution: A Macrofounded Theory of Inequality[*]

Jon Eguia[†] Dimitrios Xefteris[‡]

May 24, 2022

Abstract

We study how political competition over redistribution determines income inequality under two macrofounded premises: a) the income distribution is approximately log-normal before and after any policy intervention; and b) voters' income and turnout rates are positively related. The unique equilibrium features substantial income inequality and less than maximal redistribution, even if the voters' median income is very low. Fitting our model to the US economy, we argue that either the efficiency cost of redistribution is higher than estimates in the literature, or else US's redistributive policies are optimal for an agent richer than the median voter.

Keywords: Income inequality, redistribution, lognormal distribution, macrofoundations.

JEL Codes: D72, H20, E62.

[*]We thank Dylan Moore for the excellent comments and recommendations.

[†]Michigan State University, email: eguia@msu.edu

[‡]University of Cyprus, email: xefteris.dimitrios@ucy.ac.cy

1 Introduction

Our analysis of income inequality rests on an observed empirical regularity: the distribution of income is closely approximated by a lognormal distribution. This lognormality of income is robust to the particularities of idiosyncratic institutional frameworks and policy interventions. First established by Gibrat (1931), it has been documented in the US (Battistin, Blundell and Lewbel 2009), the UK (Kalecki 1945), India (Rajaraman 1975), the USSR (Alexeev and Gaddy 1993), Japan (Souma 2001) and several other countries (Cowell 2011, p.188).[1] That is, the lognormality of the eventual income distribution seems to have deep roots that extend beyond the specific institutional/political details of each case.[2]

An implication for political competition in a representative democracy in which candidates present competing policy platforms to the electorate, is that these platforms, regardless of how they deal with specific policy domains (e.g. taxation, unemployment benefits, access to health care/education, etc.), ultimately result in alternative income distributions that are all approximately lognormal. Voters interested in maximizing their own family income would thus evaluate competing policies by the income their family would obtain under the lognormal income distribution that would arise if each policy were implemented. For office-motivated candidates, electoral competition then boils down to choosing the feasible lognormal distribution that maximizes electoral support, and announcing policies that would (approximately) generate this income distribution. Assuming that total aggregate income along the Pareto frontier is pinned down by the chosen level of income inequality, and that redistributive policies preserve citizens' order in the income distribution, we first prove that electoral competition becomes unidimensional. That is, we show that preferences over lognormal distributions are single-peaked, so that in the unique equilibrium of a two-candidate electoral competition game, income inequality maximizes the voters' median income.[3]

Under the empirically supported assumption that the voters' median income is greater than society's median income (see for instance Larcinese 2007, or Wichowsky 2012), we then establish that the equilibrium features some degree of income inequality—i.e. redistribution is not maximal—regardless of how

[1] Cowell (2011, section 4.2) provides an extensive discussion of the lognormal function as a model of the income distribution. Among notable features, if income is lognormally distributed, its Lorenz curve is symmetric, and thus any two lognormal distributions have non-intersecting Lorenz curves.

[2] The lognormal distribution is the maximum entropy distribution for a variable x, with a well-defined mean and variance of $\ln x$. That is, assuming that the agents' beliefs regarding the future income distribution after certain policy interventions are given by a lognormal distribution aligns with the Principle of Maximum Entropy (PME) in statistics (i.e. the distribution which best represents the current state of knowledge is the one with largest entropy), while the fact that the realized distribution of income turns out to be (approximately) lognormal is compatible with the PME in physics (i.e. in a closed system with fixed energy, entropy tends to maximize). Lognormal distributions of post-tax income arise also in simple models of nonlinear income taxation with additively separable and isoelastic preferences, when we employ, for instance, the tax function of the influential HSV (Heathcote, Storesletten and Violante 2017) approach (details are available from the authors upon request).

[3] The fit between the lognormal functional form and observed income distributions is not as good at the tails of the distribution; rather, the fit is at its best where the voters' median income lies, away from the tails.

low the voters' median income is. On the other hand, if the marginal efficiency cost of redistribution (the aggregate income lost if the income to the poor increases by one unit) is not too high, equilibrium redistribution is strictly positive and benefits the median voter even if her income is slightly higher than society's mean income.

These results differ from those derived from theories of redistribution with proportional taxation and lump-sum transfers (à la Meltzer and Richard 1981), which deliver maximal redistribution if the voters' median income is very low, and none at all if the voters' median income is above society's mean income. The reason for the discrepancy is that, in the traditional setting, lump-sum transfers to any given agent, including the median voter, are linear in the total amount raised and redistributed: with n agents, they are $1/n$ of the total amount raised, for any amount raised. Whereas, if the reduction of inequality takes the form of reducing the variance of the lognormal income distribution (and, empirically, it does), then the *redistributive benefits do not flow linearly in the reduction of inequality.* On the contrary, starting with a very unequal society in which income is concentrated, say, in the top 1% of the distribution, reductions in inequality first favor a large swath of the population, from the poor to the relatively affluent, including—crucially—the median voter, while subsequent reductions in inequality benefit exclusively the poor. The median voter thus favors greater income equality up to a point, namely, up to the point where more equality harms her.

Starting with a very unequal distribution of income, a continuous reduction in income inequality all the way to fully equalized income, works like a wave: for each agent in any position in the income distribution other than those at the very top, this process first raises the agent's income, and then it lowers it. As the wave moves down the income distribution, each position in this distribution attains its peak value only once, with lower positions in the income distribution attaining their peak later, after greater reductions in income inequality.

Our analysis of income inequality and redistribution builds most directly on classical studies of affine taxation and redistribution, with a linear tax and a positive lump sum transfer. Itsumi (1974) finds that under some conditions on the distribution of ability, preferences over tax rates are single-peaked, and derives optimal tax rates for some parameterizations. Romer (1975) finds that the optimal tax rate is interior due to the disincentives to work introduced by taxation (as studied by Mirrless 1971). Roberts (1977) shows that as long as the ordering of agents by income is the same across any tax schedule in operation (an assumption we maintain throughout), majority rule aggregates preferences to a rational collective order even if individual preferences are not single-peaked. Meltzer and Richard (1981) show that the share of income redistributed increases in the ratio of mean income to the income of the decisive voter; and when the decisive voter has greater income than the mean, there will no redistribution. This literature solves each individual's agent labor allocation problem, and derives national aggregates (on income, taxation, redistribution and inequality) by summing or integrating across all individuals. Instead, we zoom out, and work directly with distributional aggregates such as the nation's income distribution.[4]

Overall, by letting our redistribution theory be informed by robust empirical regularities, rather than basing it on a detailed—yet, abstract—microfoundation, we gain additional insights and a better fit with observed data. In an application to the US economy, we find that either the efficiency cost of redistribution is toward the higher end of the estimates in the literature, or else the decisive voter is much richer than the median one. That is, the calibration of our model to the US shows that the redistributive policies are optimal for an agent with greater income than the median voter. This is compatible with the literature that explains political outcomes, not only by the mere number of supporters for each alternative, but, also, by the resources that agents are willing to spend to affect the outcome (see, for instance, Campante 2011; Karabarbounis 2011; Bonica, McCarty, Poole, and Rosenthal 2013; Eguia and Xefteris 2021; etc.).

Note that macro models often seek microfoundations to back their assumptions; whereas, micro models rarely rely on results from macroeconomic analyses. Our approach indicates that a more symmetric relationship between micro and macro economic theories is possible: we present a standard micro (i.e. game-theoretic) theory of electoral competition, but our assumptions regarding the voters' preferences are built, in effect, on macrofoundations (i.e. they rely on economy-wide phenomena). Our theory shows that a macrofounded game-theoretic analysis is feasible, and that it uncovers new angles to understanding the relationship between inequality and redistribution.

2 Theoretical argument

Consider a society with a unit mass of agents, indexed by their position $q \in [0, 1]$ in the income distribution, with the poorest agent denoted by $q = 0$ and the richest by $q = 1$. Assume that the initial income distribution, before any policy intervention, follows a lognormal functional form $\ln\mathrm{N}\left[\mu, \sigma_I^2\right]$, where $e^{\mu} \in \mathbb{R}$ is the median income, and $\sigma_I \in \mathbb{R}_{++}$ is the standard deviation of log-income, which we use as a measure of income inequality.

Let $\tilde{y} \in \mathbb{R}_+$ denote the income of a randomly chosen agent. Income being distributed according to a lognormal $\ln\mathrm{N}\left[\mu, \sigma_I^2\right]$ means that $\ln(\tilde{y})$ is distributed according to a normal distribution $\mathrm{N}\left[\mu, \sigma_I^2\right]$, and

[4] Our results, by describing how economic and political inequality (and in particular, the gap between the decisive voter's income and the median income) determines economic policy, speaks as well to more recent work documenting and analyzing the causes and consequences of higher income inequality around the turn of the 20th to the 21st century (Piketty and Saez 2003; Blundell, Pistaferri and Preston 2008; Card 2009; Greenwood, Guner, Kocharkov, and Santos 2014; Aguiar and Bils 2015).

that the mean income is

$$e^{\mu+\frac{\sigma_I^2}{2}}. \tag{1}$$

We normalize initial mean income to 1; under this normalization, $\mu = -\frac{\sigma_I^2}{2}$ and the distribution of income is $\ln\mathrm{N}\left[\frac{\sigma_I^2}{2}, \sigma_I^2\right]$.

Redistributive policies can reduce income inequality while preserving the lognormal functional form of the distribution of income, but such redistribution may destroy aggregate income. Formally, there exists a set of policies $\mathcal{P}$, and a cost parameter $c \in [0,1]$ such that for each $\sigma \in [0, \sigma_I]$, there is one policy in $\mathcal{P}$ that, if implemented, the distribution of income remains lognormal with income inequality reduced to σ, while the order of agents by income is preserved, but mean income decreases to

$$1 - c\frac{(\sigma_I - \sigma)}{\sigma_I}. \tag{2}$$

This means that the share of aggregate income lost by implementing policies that reduce inequality is linear in the relative decrease in income inequality, in proportion equal to the cost c of redistribution.

We assume that the policy choice is made by a democratically-elected government. Two office-motivated parties A and B compete in an election, by committing to a policy in $\mathcal{P}$ each. The set of voters V is exogenously given as a strict subset of the population, such that the position in society's income distribution of the median income among voters is in $\left(\frac{1}{2}, 1\right)$. Let q_v denote this position.[5] Voters observe the policy commitments, and then each voter votes for the party whose policy gives her a higher income (or abstains if both policies give her the same income). The party with greater vote share (or a randomly chosen one if they are tied) is elected and implements its policy.

Note that if the implemented policy generates income inequality σ, then equating (1) to (2), and solving for μ, we obtain that the final distribution of income is $\ln\mathrm{N}\left[\frac{-\sigma^2}{2} + \ln\left(1 - c\frac{\sigma_I - \sigma}{\sigma_I}\right), \sigma^2\right]$.[6]

Therefore, the policy choice boils down to choosing a final income distribution among the set

$$\left\{\ln\mathrm{N}\left[-\frac{\sigma^2}{2} + \ln\left(1 - c\frac{(\sigma_I - \sigma)}{\sigma_I}\right), \sigma^2\right]\right\}_{\sigma\in[0,\sigma_I]},$$

by choosing an income inequality level $\sigma \in [0, \sigma_I]$.

[5]Three different channels contribute to make the median income among voters greater than the median income among all residents. First, non-citizen residents without the right to vote are poorer than those with citizenship and the right to vote; second, poorer citizens register to vote at lower rates; and third, richer registered voters may turn out to vote at higher rates than poorer ones. Aggregating these three factors results in an income bias in turnout, such that richer residents vote at higher rates than poorer ones (Wichowsky 2012), and the median income among voters is greater than among all residents.

[6]Notice that both the mean of this distribution, $1 - c\frac{(\sigma_I - \sigma)}{\sigma_I}$, and its variance, $(e^{\sigma^2} - 1)(1 - c\frac{(\sigma_I - \sigma)}{\sigma_I})^2$, are increasing in σ for any $\sigma \in [0, \sigma_I]$ and any $c \in [0,1]$. Hence, refering to σ as income inequality is well-motivated.

Our solution concept is a Nash equilibrium of the two-party electoral competition game, given that each voter votes for the party whose policy would give the voter a higher income.

Define the **income function** $y : \mathbb{R}_{++} \times [0,1] \times [0,1] \times [0,\sigma_I] \longrightarrow \mathbb{R}_+$ so that $y(\sigma_I, c, q, \sigma)$ denotes the income of agent q in the income distribution, in a society that starts with income inequality σ_I and cost c of redistribution, and that redistributes to end up with inequality level σ. For each $(\sigma_I, c, q) \in \mathbb{R}_+ \times [0,1] \times [0,1]$, define $\sigma(\sigma_I, c, q) \equiv \arg\max_{\sigma \in [0,\sigma_I]} y(\sigma_I, c, q, \sigma)$, so $\sigma(\sigma_I, c, q)$ is the level of income inequality that maximizes the income function y with respect to σ, given (σ_I, c, q).

Definition 1 *For any $b \in \mathbb{R}_{++}$, we say that a function $f : [0,b] \to \mathbb{R}$ is **single-peaked** if there exists a maximum $x \in [0,b]$ such that the function is strictly increasing in $[0,x]$ and strictly decreasing in $[x,b]$.*

Our first result establishes that each voter's income is single-peaked in income inequality, with the peaks ordered by income.

Lemma 1 *The income function y is single-peaked in inequality σ for any $q \in [0,1]$ and $c \in [0,1]$, with maximizer $\sigma(\sigma_I, c, q)$ increasing in q.*

Proof. For any $\sigma_I \in \mathbb{R}$, any $c \in [0,1]$ and any $\sigma \in [0,\sigma_I]$, let $F[\sigma_I, c, \sigma]$ denote the cumulative distribution function of $\ln\mathrm{N}\left[-\frac{\sigma^2}{2} + \ln\left(1 - c\frac{(\sigma_I - \sigma)}{\sigma_I}\right), \sigma^2\right]$, and let $F[\sigma_I, c, \sigma]^{-1} : [0,1] \to \mathbb{R}_+$ denote its inverse, so that $F[\sigma_I, c, \sigma]^{-1}(q)$ denotes the income of an agent in position q in the income distribution. Then $y(\sigma_I, c, q, \sigma) = F[\sigma_I, c, \sigma]^{-1}(q)$.

Observe that

$$F[\sigma_I, c, \sigma](y) = \frac{1}{2}\left(1 + \mathrm{erf}\left[\frac{2\ln y + \sigma^2 - 2\ln\left(1 - c\frac{\sigma_I - \sigma}{\sigma_I}\right)}{2\sqrt{2}\sigma}\right]\right),$$

where $\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}}\int_0^x e^{-t^2} dt$ (Gauss error function), and hence for each $q \in [0,1]$,

$$F[\sigma_I, c, \sigma]^{-1}(q) = e^{\sqrt{2}\sigma\ \mathrm{erf}^{-1}(2q-1) - \frac{\sigma^2}{2} + \ln\left(1 - c\frac{\sigma_I - \sigma}{\sigma_I}\right)}.$$

Then, note that

$$\frac{\partial F^{-1}[\sigma_I, c, \sigma](q)}{\partial \sigma} = e^{\sqrt{2}\sigma\mathrm{erf}^{-1}(2q-1) - \frac{\sigma^2}{2} + \ln\left(1 - c\frac{(\sigma_I - \sigma)}{\sigma_I}\right)}\left(\sqrt{2}\mathrm{erf}^{-1}(2q-1) - \sigma + \frac{\frac{c}{\sigma_I}}{1 - c\frac{(\sigma_I - \sigma)}{\sigma_I}}\right),$$

which is strictly positive [zero] if and only if

$$\sqrt{2}\text{erf}^{-1}(2q-1) - \sigma + \frac{c}{\sigma_I - c(\sigma_I - \sigma)} \tag{3}$$

is strictly positive [zero]. Expression (3) is strictly decreasing in σ, and thus it attains a value of zero at most at a unique value of σ, and thus $\frac{\partial \bar{F}^{-1}[\sigma_I,c,\sigma](q)}{\partial \sigma}$ is strictly positive up to this value (if it exists) and strictly negative beyond it, or else, if there is no value such that $\frac{\partial \bar{F}^{-1}[\sigma_I,c,\sigma](q)}{\partial \sigma}$ is zero, then it is strictly positive for any value of σ, or strictly negative for any value of σ. In any of the three cases, for any (σ_I, q, c), $F^{-1}[\sigma_I, c, \sigma](q) = y(\sigma_I, c, q, \sigma)$ is single-peaked with respect to σ over $[0, \sigma_I]$.

Since, for any (σ_I, c, q), $y(\sigma_I, c, q, \sigma)$ is single-peaked with respect to σ, it follows that the maximizer $\sigma(\sigma_I, c, q)$ is a function (not a multi-valued correspondence). We want to show that for any $\sigma_I \in \mathbb{R}_{++}$ and any $c \in [0, 1]$, $\sigma(\sigma_I, c, q)$ is non-decreasing in q. We noted that Expression (3) is strictly decreasing in σ. Note as well that for $\sigma = 0$, Expression (3) is non-negative for any $c \in [0, 1]$. It follows that for any $(\sigma_I, c, q) \in \mathbb{R}_{++} \times [0, 1] \times [0, 1]$, a maximizer $\sigma(\sigma_I, c, q)$ of the income function y is either a corner solution at $\sigma(\sigma_I, c, q) = \sigma_I$ (if Expression (3) is positive for $\sigma = \sigma_I$), or an interior solution equal to the value σ such that Expression (3) equals zero, otherwise. Observe that Expression (3) is strictly increasing in q. Thus, holding fixed σ_I and c, if q increases, a corner solution at $\sigma(q, c) = 0$ may turn into an interior solution, and any interior solution for σ must increase with increases in q, unless it becomes a corner solution at $\sigma(\sigma_I, c, q) = \sigma_I$. Thus, $\sigma(\sigma_I, c, q)$ is increasing in q, strictly so where $\sigma(\sigma_I, c, q) \in (0, \sigma_I)$. ■

Figure 1 illustrates Lemma 1 for initial inequality $\sigma_I = 1$, cost of redistribution $c = 0.2$, and agents in quantiles $q = 0.5$, $q = 0.6$ and $q = 0.7$ in the distribution of income. Their incomes $y(\sigma_I, c, q, \sigma)$ are all single-peaked, with respective maxima at $\sigma(1, 0.2, 0.5) = 0.24$, $\sigma(1, 0.2, 0.6) = 0.48$, and $\sigma(1, 0.2, 0.7) = 0.74$. Since voters' preferences over income inequality are single-peaked, and the individual peaks (income maxima) are ordered by the voters' position in the income distribution, it follows as a corollary that a majority of voters prefers the inequality level that maximizes the voters' median income to any other inequality level in pairwise comparisons (Black 1948), and that the political equilibrium is such that both parties propose this level of inequality that maximizes the voters' median income (Downs 1957).

Our environment features three primitive parameters: the initial income inequality $\sigma_I \in \mathbb{R}_{++}$; the position (quantile) $q_v \in \left(\frac{1}{2}, 1\right)$ of the median voter in the income distribution; and the cost of redistribution $c \in [0, 1]$. For each party $J \in \{A, B\}$, let $\sigma_J : \mathbb{R}_{++} \times \left(\frac{1}{2}, 1\right) \times [0, 1] \to [0, \sigma_I]$ denote the equilibrium income inequality proposed by party J, and let $\sigma^* : \mathbb{R}_{++} \times \left(\frac{1}{2}, 1\right) \times [0, 1] \to [0, \sigma_I]$ denote the equilibrium outcome inequality, as a function of parameters. Recall that $\sigma(\sigma_I, c, q)$ denotes the income inequality that

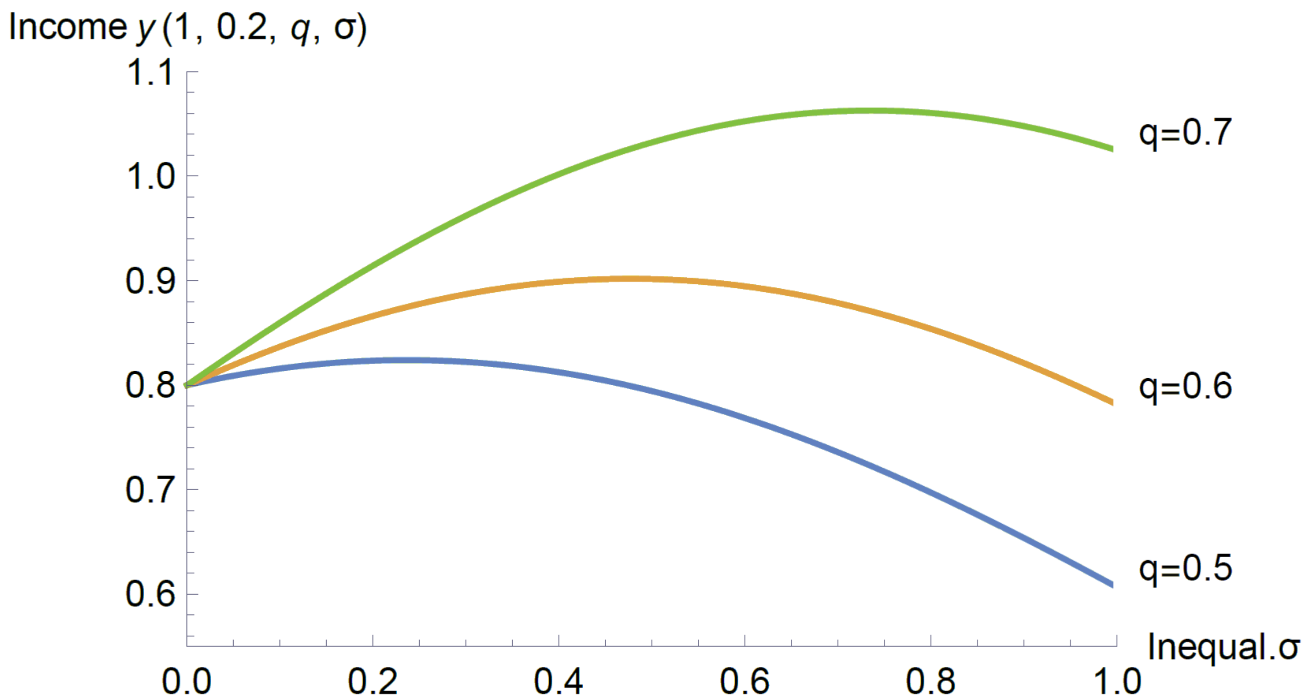

Figure 1: Income y is single-peaked in inequality σ for each q (shown for $\sigma_I = 1, c = 0.2$).

maximizes the income of an agent at position q in the income distribution. We can then formalize the corollary mentioned in the previous paragraph thus.

Corollary 2 ***(Median Voter Theorem)*** *For any $\sigma_I \in \mathbb{R}_{++}$, $q_v \in \left(\frac{1}{2}, 1\right)$ and $c \in [0,1]$, the equilibrium is unique and is such that $\sigma_A(\sigma_I, q_v, c) = \sigma_B(\sigma_I, q_v, c) = \sigma^*(\sigma_I, q_v, c) = \sigma(\sigma_I, c, q_v)$.*

In light of this corollary, we describe subsequent results in terms of the implemented policy σ^*, with the implication that they apply as well to the parties' proposals, and to the median voters' income-maximizing inequality level.

It is instructive to consider the results with linear taxation and lump-sum redistribution (Meltzer and Richard 1981) as a benchmark. Because redistributive policies typically reduce aggregate income, define "redistribution" as the aggregate increase in income among agents whose income increases. Noting that the initial and final outcome income of an agent in position q in the income distribution are respectively $y(\sigma_I, c, q, \sigma_I)$ and $y(\sigma_I, c, q, \sigma)$, redistribution thus defined is equal to

$$\int_0^1 \max\{y(\sigma_I, c, q, \sigma) - y(\sigma_I, c, q, \sigma_I), 0\}\mathrm{d}q.$$

This is the shaded area in Figure 2. Under Meltzer and Richard (1981) redistributive scheme, tax rates

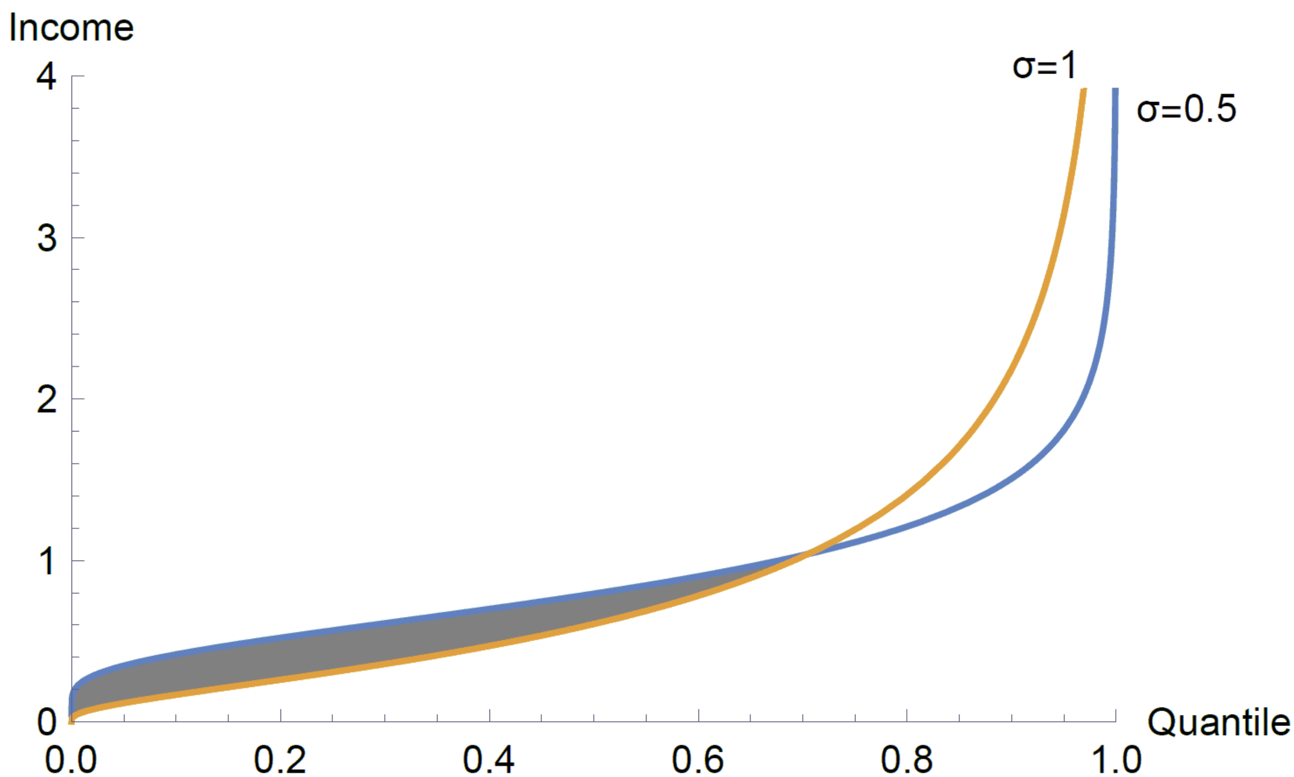

Figure 2: Income by quantile, $\sigma_I = 1$ and $c = 0.2$.

and redistribution:

i) are continuous and weakly increasing in the ratio of voters' initial median income to society's initial mean income (which we normalize to one);

ii) are maximal if this income ratio is sufficiently low; and

iii) are zero (minimal) if the voters' initial median income is greater than one.

In our environment, the voters' median income is determined by the voters' median position in the income distribution, denoted q_v, and by the initial income inequality σ_I. Let $\bar{q}$ denote the initial position in the income distribution of a voter with mean income; formally, $\bar{q} = \ln\mathrm{N}\left[\frac{\sigma_I^2}{2}, \sigma_I^2\right](1) \in \left(\frac{1}{2}, 1\right)$. Then, for any initial income inequality, the voters' median income is greater than one if the voters' median position in the income distribution is high enough (i.e. if $q_v \in (\bar{q}, 1)$), and for any position q_v, the voters' median income is arbitrarily low if inequality σ_I is sufficiently high.

Our results agree with Meltzer and Richard's (1981) first result: the redistribution and the reduction in inequality are continuous in parameters. However, we reject the bang-bang nature of the other two predictions: redistribution is never maximal because maximal redistribution is never in the interest of a median voter who is richer than the median agent; on the other hand, for a relevant range of parameters with voters' median income above society's mean income, the equilibrium policy features some redistribution and reduction in inequality, which benefits the median voter. These are our main results.

Proposition 3 *There exists $\bar{c} > 0$ such that for every $c \in [0, \bar{c}]$, the equilibrium level of income inequality $\sigma^*(\sigma_I, q_v, c)$ satisfies the following properties:*

i) It is continuous and increasing (strictly if at an interior solution) in the position q_v of the voters' median income in the distribution of income, and in the cost c of redistribution;

ii) It is strictly greater than zero for any parameter values, and redistribution is less than maximal; and

iii) For any $q_v \in \left(\frac{1}{2}, 1\right)$, there exists a non-empty interval $(a(c), b(c)) \subset \mathbb{R}_{++}$ such that if $\sigma_I \in (a(c), b(c))$, then the voters' initial median income is greater than society's mean income, and yet the equilibrium features a reduction in income inequality from σ_I to $\sigma^*(\sigma_I, q_v, c) < \sigma_I$, and strictly positive redistribution.

Proof. From Corollary 2 we obtain that $\sigma^*(\theta) = \sigma(\sigma_I, c, q_v)$. From the proof of Lemma 1, we obtain that

$$\sigma(\sigma_I, c, q_v) = \begin{cases} \sigma_I & \text{if } \sqrt{2}\text{erf}^{-1}(2q_v - 1) - \sigma_I + \frac{c}{\sigma_I} \geq 0, \\ \sqrt{2}\text{erf}^{-1}(2q_v - 1) + \frac{c}{(1-c)\sigma_I + c\sigma} & \text{otherwise,} \end{cases} \tag{4}$$

where in the second case, σ is left as an implicit function (we can solve for σ explicitly, if desired).

From $\sigma_I > 0$ and $q_v > \frac{1}{2}$ (so $\text{erf}^{-1}(2q_v - 1) > 0$), we obtain $\sigma^*(\sigma_I, q_v, c) > 0$ as stated in property ii).

From the continuity of the error function erf and its inverse we obtain that a sufficiently small change in q_v leads to an arbitrarily small change in $\sqrt{2}\text{erf}^{-1}(2q_v - 1)$, which, if the solution is interior and remains interior, in order to continue satisfying the equation for an interior equilibrium, implies an equally small change in $\sigma - \frac{c}{(1-c)\sigma_I + c\sigma}$, which implies an arbitrarily small change in σ. Continuity of $\sigma(\sigma_I, c, q_v)$ with respect to q_v as the solution transits from an interior solution to the corner one follows from

$$\lim_{\sigma \to \sigma_I} \left(\sqrt{2}\text{erf}^{-1}(2q_v - 1) + \frac{c}{(1-c)\sigma_I + c\sigma} \right) = \sqrt{2}\text{erf}^{-1}(2q_v - 1) + c.$$

Thus, σ^* is continuous in q_v.

To establish continuity of $\sigma(\sigma_I, c, q_v)$ with respect to c, first note that if $c' - c > 0$ is sufficiently small, and the solutions $\sigma(\sigma_I, c, q_v)$ and $\sigma(\sigma_I, c', q_v)$ are both interior, then $c' - c$ is sufficiently small implies $\left| \frac{c}{(1-c)\sigma_I + c\sigma(\sigma_I, c', q_v)} - \frac{c'}{(1-c')\sigma_I + c'\sigma(\sigma_I, c', q_v)} \right|$ is arbitrarily small, and thus $|\sigma(\sigma_I, c, q_v) - \sigma(\sigma_I, c', q_v)|$ is arbitrarily small as well. Further, as noted above, if the solution transitions from an interior solution to a corner solution, it does so continuously, converging to value σ_I as it approaches the corner solution, and staying at σ_I in the corner solution. Therefore, $\sigma(\sigma_I, c, q_v)$ is continuous in c.

Note that since the Gauss error function erf and its inverse are monotonically increasing, $\sqrt{2}\text{erf}^{-1}(2q_v - 1)$ strictly increases in q_v. Thus, if for the initial value of q_v, $\sigma(\sigma_I, c, q_v) = \sigma_I$, then this holds as well for any higher value of q_v. Further, if $\sigma(\sigma_I, c, q_v)$ is at the intermediate value defined by the implicit function, then note that a strict increase in $\sqrt{2}\text{erf}^{-1}(2q_v - 1)$ means that to continue to satisfy the implicit equation, $\sigma - \frac{c}{(1-c)\sigma_I + c\sigma}$ must strictly increase as well, which implies that σ must strictly increase, so $\sigma(\sigma_I, c, q_v)$ is strictly increasing in q_v for any interior solution, and increasing for any solution.

Since $\frac{c}{\sigma_I}$ and $\sqrt{2}\text{erf}^{-1}(2q_v - 1) - \sigma_I + \frac{c}{\sigma_I}$ are strictly increasing in c, if for the initial value of c, $\sigma(\sigma_I, c, q_v) = \sigma_I$, then this holds as well for any higher value of c. To show that $\sigma(\sigma_I, c, q_v)$ is strictly increasing in c for any interior solution, notice that in any interior solution, Expression (3) must be equal to zero, so $-\sigma + \frac{c}{\sigma_I - c(\sigma_I - \sigma)}$ must be constant with the increase in c. Since $\frac{c}{\sigma_I - c(\sigma_I - \sigma)}$ strictly increases in c and strictly decreases in σ, in order for it to stay constant with a strict increase in c, σ must strictly increase as well. Therefore, $\sigma(\sigma_I, c, q_v)$ is increasing in c, strictly for any interior solution, completing our proof of all the claims stated in property i).

Further, note that if $c = 0$, redistribution is maximal if and only if income is fully equalized, i.e. if $\sigma = 0$. Since $\sigma^* > 0$, it does not feature maximal redistribution. By continuity of σ^* with respect to c, and continuity of redistribution with respect to σ^* and c, we obtain that there exists $\bar{c} > 0$ such that redistribution is maximal if and only if income is fully equalized, income fails to be fully equalized, and redistribution is less than maximal for any $c \in [0, \bar{c}]$, establishing property ii).

For property iii), from the solution for $\sigma(\sigma_I, c, q_v))$ in Expression (4), for any $q_v \in \left(\frac{1}{2}, 1\right)$, and for any $c \in [0, 1]$,

$$\sigma(\sigma_I, c, q_v)) < \sigma_I \iff \sigma_I > \frac{\sqrt{2}\text{erf}^{-1}(2q_v - 1) + \sqrt{2\left(\text{erf}^{-1}(2q_v - 1)\right)^2 + 4c}}{2}.$$

Further, note that

$$(y \text{ evaluated at } (\sigma_I, c, q_v, \sigma_I) \text{ is strictly greater than one}) \iff \sigma_I \in \left(0, 2\sqrt{2}\text{erf}^{-1}(2q_v - 1)\right).$$

Together, this pair of necessary and sufficient conditions imply that for any $c \in \left[0, 4\left(\text{erf}^{-1}(2q_v - 1)\right)^2\right)$, and for any $\sigma_I \in \left(\frac{\sqrt{2}\text{erf}^{-1}(2q_v-1) + \sqrt{2\left(\text{erf}^{-1}(2q_v-1)\right)^2 + 4c}}{2}, 2\sqrt{2}\text{erf}^{-1}(2q_v - 1)\right)$, the initial income of agent q_v is greater than mean income (one), and yet agent q_v prefers strictly positive redistribution and a reduction in inequality to $\sigma(\sigma_I, c, q_v)) < \sigma_I$, which increases her income to $y(\sigma_I, c, q_v, \sigma(\sigma_I, c, q_v)) > y(\sigma_I, c, q_v, \sigma_I) > 1$.

■

Perhaps the most striking result is that for a range of initial income inequality levels, purely selfish voters with greater than mean income support redistributive policies that reduce inequality. This is so because for such range of income inequality, redistribution first flows from the very rich at the right tail, to most other agents, from the poorest to the somewhat rich, including the median voter.[7]

We have stated our comparative statics results focusing only on the most interesting cases regarding the admissible costs of redistribution (i.e. c close to zero). Property i) holds for general costs of redistribution, and, intuitively, if c is very large then the median voter does not want to redistribute income as the reduction in the average income dominates the gains from the transfers she receives from the richer agents.

3 Application: Income inequality in the United States

We next fit the theory to the United States' economy and its political environment. Our theory has three parameters: (σ_I, q_v, c), respectively the initial income inequality, the quantile of the voters' median income in society's distribution of income, and the cost of redistribution expressed as the linear parameter of loss of aggregate income in the reduction of inequality.

The quantile q_v can be estimated using Table 1 on turnout by income bracket, from the 2018 US Current Population Survey. Let $\hat{q}_v$ denote the estimated value of q_v. Interpolating linearly within the \$75,000 to \$99,000 income bracket, we estimate the voters' median family income is \$87,000, at quantile $\hat{q_v} = 0.59$.

According to 2018 data from the Congressional Budget Office (CBO), means-tested transfers to the bottom three quantiles of the income distribution amounted to 5.2% of the US economy, and the ratio of the mean family income to the median family income in the economy was 1.17. We seek to recover a pair of parameters (σ_I, c) such that if the (counterfactual) distribution of US income in the absence of such transfers were a lognormal $\ln\mathrm{N}\left[\frac{\sigma_I^2}{2}, \sigma_I^2\right]$, and if the cost of reducing income inequality were proportional to c, then according to our theory, total transfers would amount to 5.2% of the economy, and the ratio of mean to median income after the transfers would be 1.17, as in the data. Further, we wish that the marginal efficiency cost of redistribution (MECR) in the equilibrium resulting from initial conditions $(\sigma_I, \hat{q_v}, c)$ be consistent with the estimates for this cost in the literature. This marginal efficiency cost of

[7] Public policies that benefit households moderately richer than the median include those that fund services disproportionately used by these households and less so by poorer ones, such as higher education, commercial air travel, or publicly-funded culture and fine arts.

Income	Population	Cum. % of pop.	Turnout	Cum. % of vote
Under \$10,000	4122	2.9%	35.4%	1.6%
\$10,000 to \$14,999	3893	5.7%	39.6%	3.4%
\$15,000 to \$19,999	3153	7.9%	41.0%	4.7%
\$20,000 to \$29,999	10,395	15.2%	40.8%	9.4%
\$30,000 to \$39,999	13,493	24.8%	49.1%	16.8%
\$40,000 to \$49,999	10,045	31.9%	54.5%	23.2%
\$50,000 to \$74,999	27,937	51.6%	57.1%	42.4%
\$75,000 to \$99,999	19,796	65.6%	63.2%	57.7%
\$100,000 to \$149,999	24,241	82.8%	67.3%	78.1%
\$150,000 and over	24,396	100%	71.8%	100%

Table 1: Turnout by family income bracket, CPS 2018.

redistribution is the loss of aggregate income per dollar increase in poor agents' income. Its estimates vary, from \$0.50cts to \$1.00cts (Hendren 2020), or \$0.50cts to \$1.30cts (Ballard 1988), or \$1.29cts to \$5.84 (Browning and Johnson 1984). With two free parameters (σ_I, c), our equilibrium must match three variables, two exactly pinned down, and one estimated much less precisely.

We find that for an initial economy with parameter values $\sigma_I = 0.73$, $q_v = \hat{q_v} = 0.59$ and $c = 0.23$, the predicted inequality level is $\sigma^* = 0.56$, which, given that for any $\sigma \in \mathbb{R}_{++}$, the mean/median ratio of $\ln\mathrm{N}\left[\frac{\sigma^2}{2}, \sigma^2\right]$ is equal to $e^{\frac{\sigma 2}{2}}$, implies that the equilibrium mean/median ratio is 1.17, as in the data. We illustrate the fit of the US distribution of income after tax and transfers to our predicted lognormal distribution in Figure 3, where the bars represent the after tax and transfers income by percentile brackets as reported to the CBO (with the top decile zoomed in on the graph on the right) and the line is our predicted theoretical lognormal income distribution.

Further, redistribution (defined as the increase in income for agents whose income increase as inequality reduces from σ_I to σ^*) amounts to 4.48% of the final economy, which, given Hendren's (2020) estimate that it takes \$1.15cts of transfers to the poor to increase their income by \$1, implies total transfers of 5.2% of the economy, as in the CBO data. According to our theory, this redistribution causes a loss of aggregate income (through decreases in income to the rich) of 5.6% of the final economy, with an implied average efficiency cost of redistribution of 1.26, and a marginal efficiency cost of \$1.32cts, (coincidentally) almost at the narrow intersection between Browning and Johnson's (1984) and Ballard's (1988) estimates, and above Hendren's (2020) estimates.

The lingering doubts about the true value of the marginal cost of redistribution suggest two ways to interpret our theory. We explore both. The first view is that the true value is at the intersection between

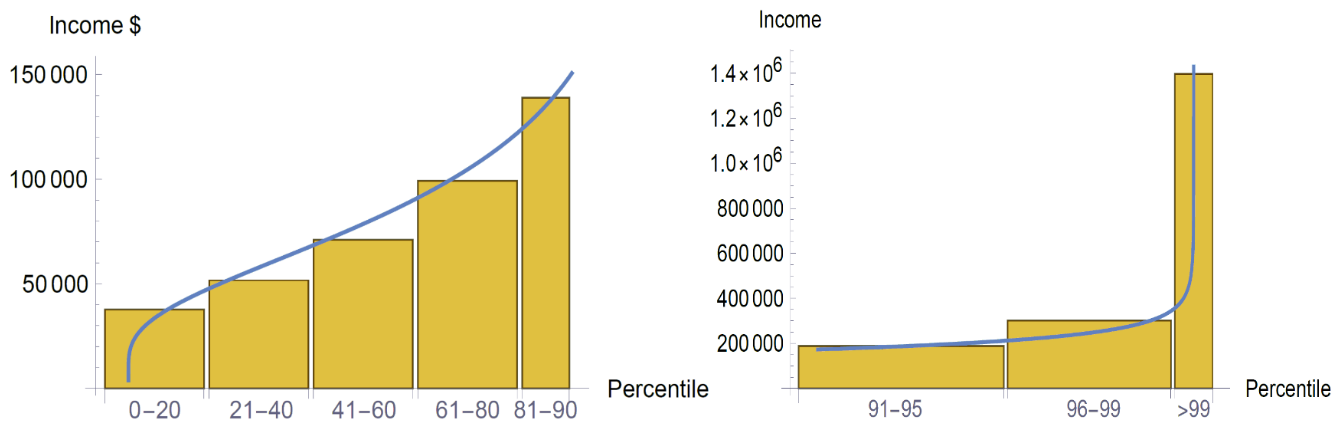

Figure 3: US Income distribution, $\sigma^* = 0.56$ and $c = 0.23$.

Ballard's (1988) and Browning and Johnson's (1984) estimates, while Hendren's (2020) estimate is too low. If so, for the proposed parametrization $(\sigma_I, q_v = \hat{q_v}, c) = (0.73, 0.59, 0.23)$, our model fits the data and the estimates of the four politico-economic variables that we track (mean/median income ratio; total mean-tested transfers; marginal efficiency cost of redistribution; and position of the voters' median income in the society's income distribution) near perfectly, and we can use the theory to derive counterfactual predictions, such as the following: halving the difference between turnout by the rich and poor noted in Table 1 (which would bring this turnout gap to European standards and would lower the position of the voters' median income to 0.545), would reduce inequality to $\sigma^* = 0.45$, would increase total mean-tested transfers to 8.3% of the total economy (an increase of approximately 600 billion dollars a year), and it would lower aggregate income by 3.4%.

The second interpretation is that Ballard's (1988) and Hendren's (2020) lower estimates of the marginal efficiency cost of redistribution are correct, and our theory errs in predicting too much redistribution for any given cost of redistribution, so that it needs to input an unrealistically high cost parameter in order to output the correct amount of redistributive transfers. Where could our theory have gone wrong? Our theory rests on two fundamental premises: First, the income distribution is lognormal, after redistribution just as well as before redistribution, so choosing redistributive policies amounts to choosing lognormal income distributions. This premise is empirically validated, as illustrated graphically by Figure 3. The second premise, inherited from Downs (1957) and Meltzer and Richard (1981), is that the equilibrium is determined by the preferences of a decisive agent, assumed to be the median voter. The assumption that the median voter determines the policy outcome is not empirically supported; to the contrary, policy is relative more responsive to the interests of rich voters than the interest of poor ones (Karabarbounis 2011), so the decisive agent whose welfare is maximized by the policy choice is richer than the median voter. How much richer? Let us formalize and try to answer this question.

Let q_d denote the quantile in the income distribution whose income is maximized by the equilibrium inequality. In our benchmark theory, we obtained that $q_d = q_v$, i.e. that the median voter is the decisive agent (Corollary 2). Without belaboring one of many possible micro-foundations for this departure, we now drop the Median Voter Theorem and assume instead that the equilibrium level of inequality is $\sigma(\sigma_I, c, q_d)$ (the maximizer for an agent at position q_d in the income distribution), and we recover an estimate $\hat{q}_d$ of q_d by imposing that our model fit the mean/median income ratio (1.17), the mean-tested transfers (5.2% of the economy) estimated by the CBO, and the lower marginal efficiency cost of redistribution as estimated by Hendren (2020) and Ballard (1988). Here is what we find: for $\sigma_I = 0.71$, $c = 0.15$ and $q_d = 0.635$, the equilibrium inequality σ^* is 0.56, which implies a mean/median income ratio 1.17 and mean-tested transfers of 5.2% of the economy, matching the data; and the marginal efficiency cost of redistribution is 0.73cts, near the center of Hendren's (2020) interval estimate, and Ballard's (1988) point prediction of 0.80cts. So, if we believe this lower estimate of the cost of redistribution, and our theory, then we can infer that the US macro-economic policies are optimally catering to voters in the 64th percentile in the income distribution. These voters benefit from the reduction in inequality and in equilibrium, their income raises above both the final mean society income and above the initial mean society income.

In sum, we find that the U.S. income distribution is approximately log-normal, with greater income inequality than what the median voter would want given the central estimates of the efficiency cost of redistribution. We infer that either the actual efficiency cost of redistribution is higher than those estimates, or the redistributive policies implemented, and the resulting levels of income inequality, are optimal for a voter who is substantially richer than the median.

References

[1] Aguiar, Mark, and Mark Bils (2015). "Has consumption inequality mirrored income inequality?." *American Economic Review* 105.9: 2725–56.

[2] Alexeev, Michael V., and Clifford G. Gaddy (1993). "Income Distribution in the USSR in the 1980s." *Review of Income and Wealth* 39.1: 23-36.

[3] Ballard, Charles L. (1988). "The marginal efficiency cost of redistribution." *American Economic Review* (1988): 78.5: 1019–1033.

[4] Battistin, Erich, Richard Blundell, and Arthur Lewbel (2009). "Why is consumption more log normal than income? Gibrat's law revisited." *Journal of Political Economy* 117.6: 1140–1154.

[5] Black, Duncan (1948). "On the rationale of group decision-making." *Journal of Political Economy* 56.1 (1948): 23–34.

[6] Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008). "Consumption inequality and partial insurance." *American Economic Review* 98.5: 1887–1921.

[7] Bonica, Adam, Nolan McCarty, Keith T. Poole, and Howard H. Rosenthal. (2013). "Why hasn't democracy slowed rising inequality?" *Journal of Economic Perspectives*, 27(3), 103-24.

[8] Browning, Edgar K., and William R. Johnson (1984). "The trade-off between equality and efficiency." *Journal of Political Economy* 92.2: 175–203.

[9] Campante, Filipe R. (2011). "Redistribution in a model of voting and campaign contributions." *Journal of Public Economics*, 95(7-8), 646-656.

[10] Card, David (2009). "Immigration and inequality." *American Economic Review* 99.2: 1–21.

[11] Congressional Budget Office (2021). "The distribution of household income, 2018." Report, August 4, 2021, available at https://www.cbo.gov/publication/57061.

[12] Cowell, Frank (2011). *Measuring Inequality.* Oxford University Press.

[13] Downs, Anthony (1957). *An Economic Theory of Democracy.* New York: Harper and Row.

[14] Eguia, Jon X., and Dimitrios Xefteris (2021). "Implementation by vote-buying mechanisms." *American Economic Review*, 111.9: 2811–28.

[15] Gibrat, Robert (1931). *Les Inégalites Économiques.* Paris: Librairie du Recueil Sirey.

[16] Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos (2014). “Marry your like: Assortative mating and income inequality.” *American Economic Review* 104.5: 348–53.

[17] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2017). ”Optimal tax progressivity: An analytical framework.” *The Quarterly Journal of Economics*, 132.4: 1693–1754.

[18] Hendren, Nathaniel (2020). “Measuring economic efficiency using inverse-optimum weights.” *Journal of Public Economics* 187: 104198.

[19] Itsumi, Yoshitaka (1974). “Distributional effects of linear income tax schedules.” *Review of Economic Studies* 41.3: 371–381

[20] Kalecki, Michael (1945). “On the Gibrat Distribution.” *Econometrica* 13.2: 161–70.

[21] Karabarbounis, Loukas (2011). “One Dollar, One Vote.” *Economic Journal* 121: 621–651.

[22] Larcinese, Valentino (2007). “Voting over redistribution and the size of the welfare state: the role of turnout.” *Political Studies*, 55(3), 568-585.

[23] Meltzer, Allan H., and Scott F. Richard (1981). “A rational theory of the size of government.” *Journal of Political Economy* 89.5: 914–927.

[24] Mirrlees, James A. (1971). “An exploration in the theory of optimum income taxation.” *Review of Economic Studies* 38: 175–208.

[25] Piketty, Thomas, and Emmanuel Saez (2003). ”Income inequality in the United States, 1913–1998.” *Quarterly Journal of Economics* 118.1: 1–41.

[26] Rajaraman, Indira (1975). ”Poverty, inequality and economic growth: Rural Punjab, 1960/61-1970/1971.” *Journal of Development Studies* 11.4: 278–290.

[27] Roberts, Kevin W. S. (1977). “Voting over income tax schedules.” *Journal of Public Economics* 8: 329–40.

[28] Romer, Thomas (1975). “Individual welfare, majority voting, and the properties of a linear income tax.” *Journal of Public Economics* 4: 163–185.

[29] Souma, Wataru (2001). “Universal structure of the personal income distribution.” *Fractals* 9.04: 463–470.

[30] Wichowsky, Amber (2012). "Competition, party dollars, and income bias in voter turnout, 1980–2008." *Journal of Politics* 74.2: 446-459.