Mechanism Design and Marginal Distributions

Mechanism Design for Social Good 2018 workshop

Robert Manduca - Harvard Sociology

June 22, 2018

Marginal distributions and allocation processes

- Social outcomes are determined by two types of input:
  - Marginal distributions – what is the set of possible outcomes available
  - Allocation processes – who gets assigned to which positions
- Sociologists have historically been most interested in allocation processes:
  - Who lives in which neighborhoods?
  - Who gets hired for which jobs?
  - What determines where a child goes to school?

Today's big social challenges will not be solved solely by better allocation of existing positions

- Many of our most pressing social problems are not fully solvable through better allocation of existing social positions alone. In this talk I will give two examples from my research:
  - Upward income mobility
  - Racial economic equality
- Then I will describe ways mechanism design might be harnessed to create a better marginal distribution of outcomes, rather than trying to more optimally allocate the inadequate set we have today

Allocation and distribution 1: declining upward income mobility

- Upward mobility is central to American identity
  - "A better life for your children"
  - Also key for tolerance, fairness, democracy… (Friedman 2005; Mullainathan and Shafir 2013)
- But absolute upward mobility rates have been falling for 50 years
- How do we reverse this decline?

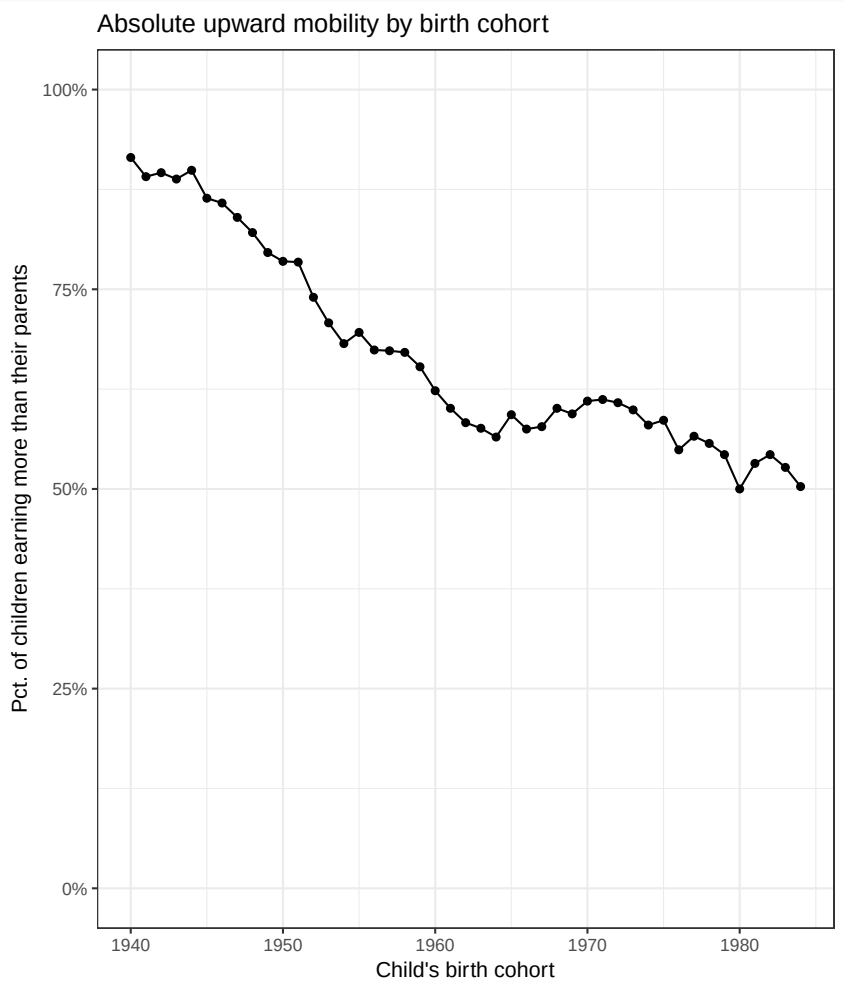

Fairer allocation processes would not restore mass upward mobility

- "Equality of opportunity" is typically defined as children's economic positions have zero correlation with their parents'
- But perfect equality of opportunity would not increase aggregate upward mobility
  - Observed mean = 46.9%
  - Mean with perfect relative mobility = 46.4%

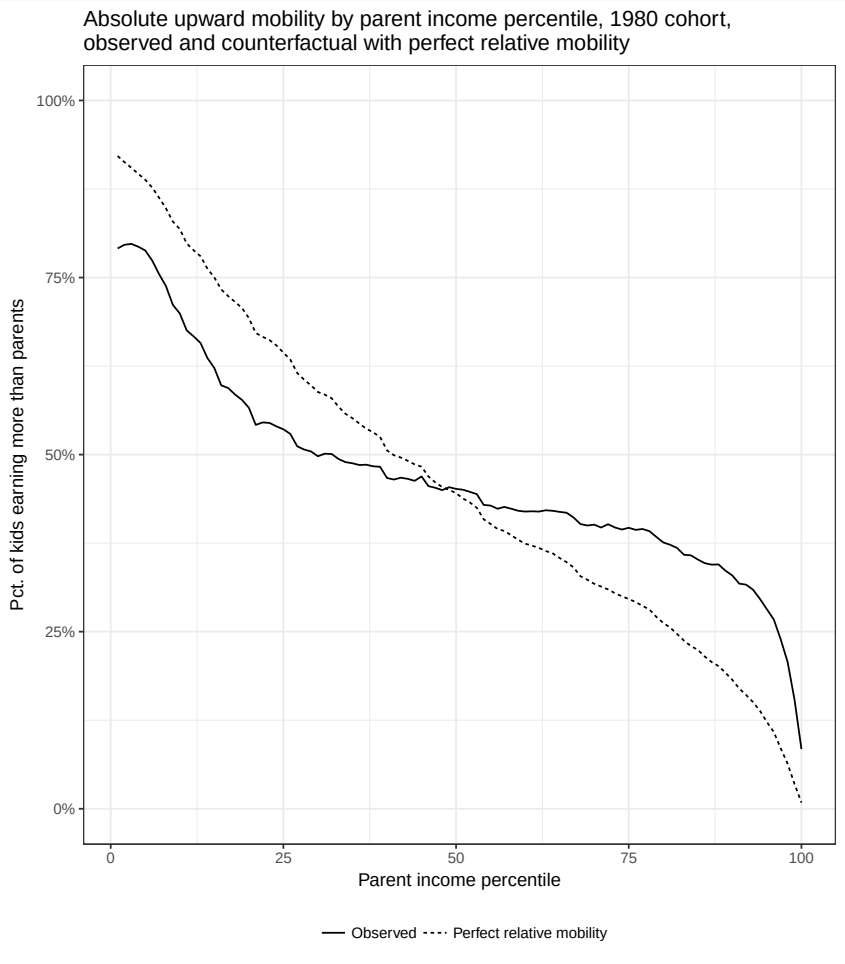

Widespread upward mobility requires a more equitable income distribution

- If fair allocation is not enough, how do we restore upward mobility?
- Returning to the income distribution of 1970 would reverse most of the mobility decline
  - Observed mean = 46.9%
  - Mean with 1970 income distribution = 78.3%

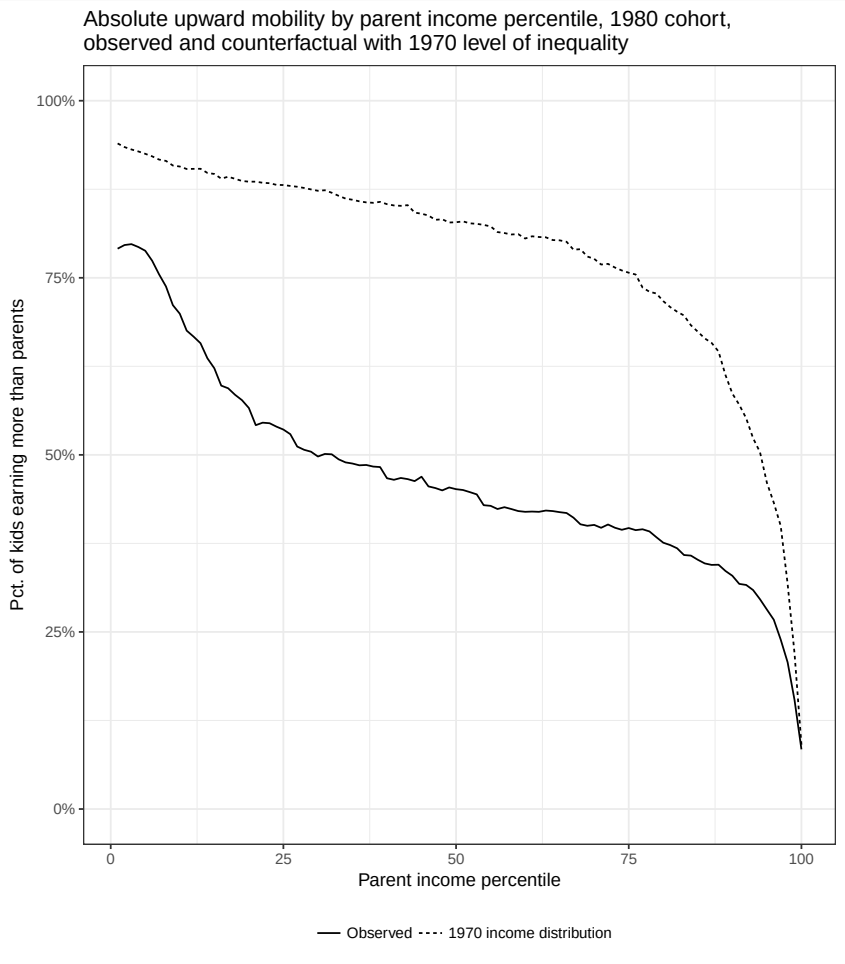

Fixing upward mobility requires changing the set of outcomes that are available, not better allocating the outcomes we have now

Allocation and distribution 2: Black-white family income disparities

- The family income gap between blacks and whites has not changed for the last 50 years
- Most explanations for the lack of racial progress emphasize continued (and well documented) racial stratification:
  - Processes that sort whites into better jobs, better schools, etc.

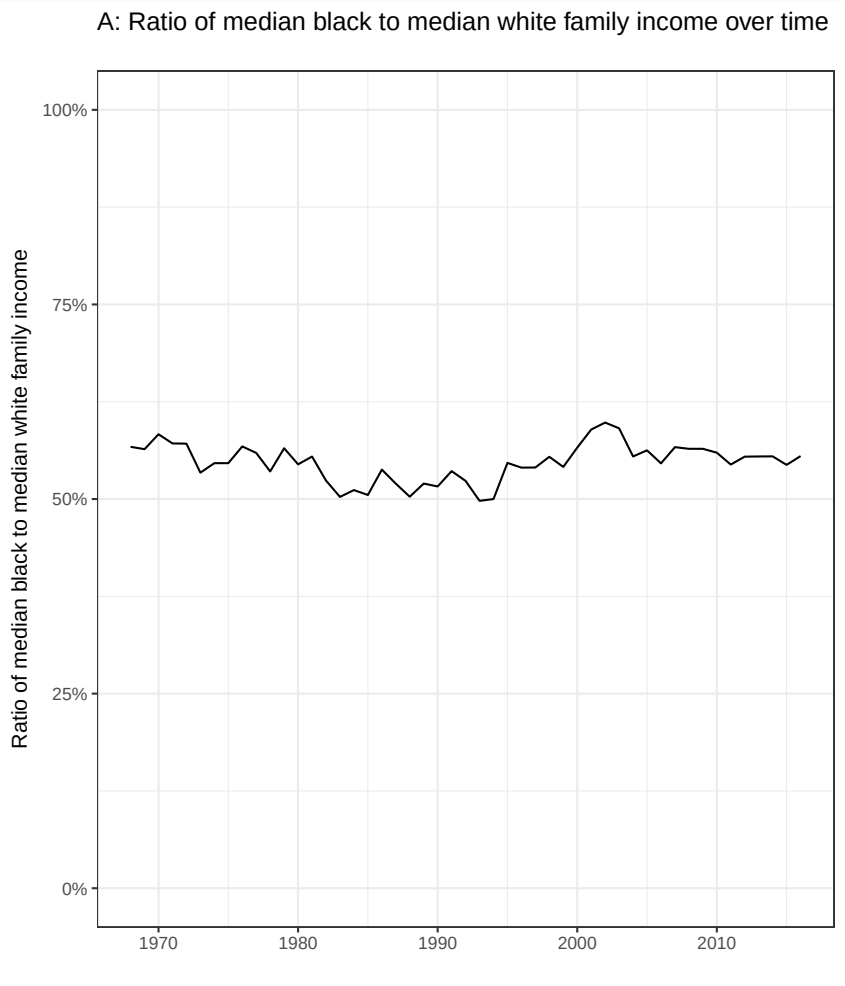

US racial income stratification has decreased substantially since the 1960s

- In 1968, the median black American was at the 25th percentile of US family income
- In 2016 they were at the 35th percentile
- The black-white gap rank has shrunk by 28%

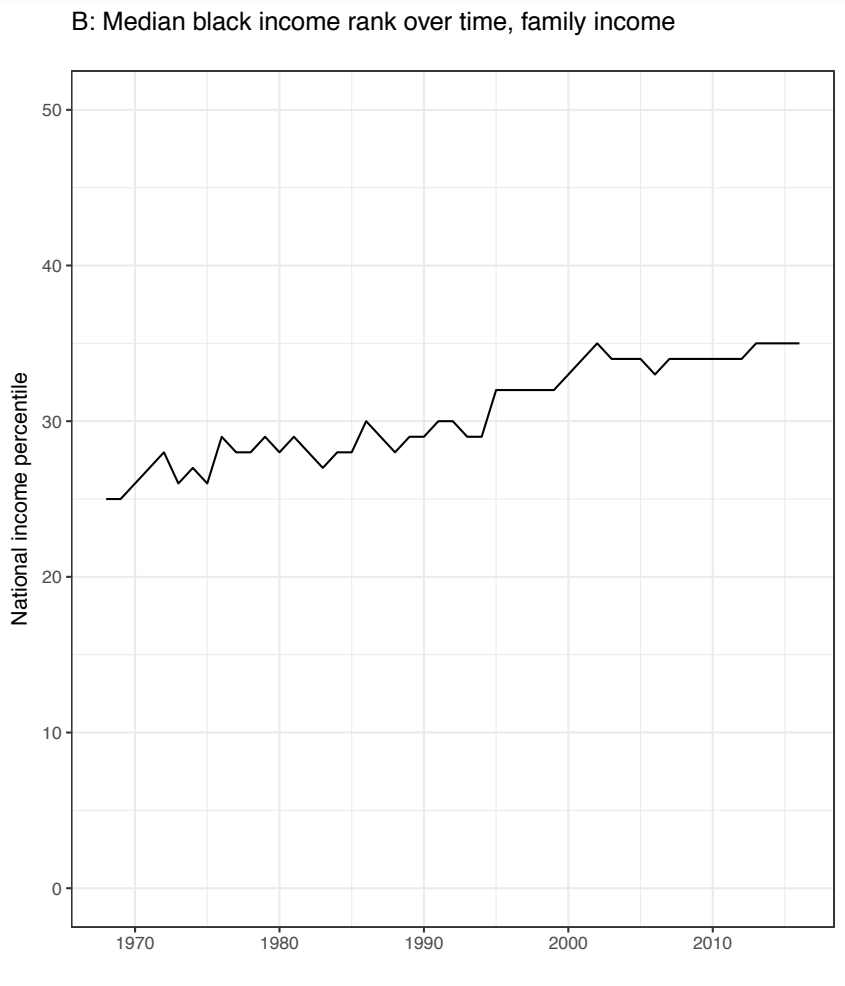

Reductions in racial stratification were negated by rising economy-wide income inequality

- As the rank gap was closing, overall inequality was rising
  - 1968 income shares:
    - Richest 1%: 12.0%
    - Poorest 50%: 19.2%
  - 2014 income shares:
    - Richest 1%: 19.0%
    - Poorest 50%: 10.3%
- These shifts reduced the payoff for African Americans climbing the income ladder

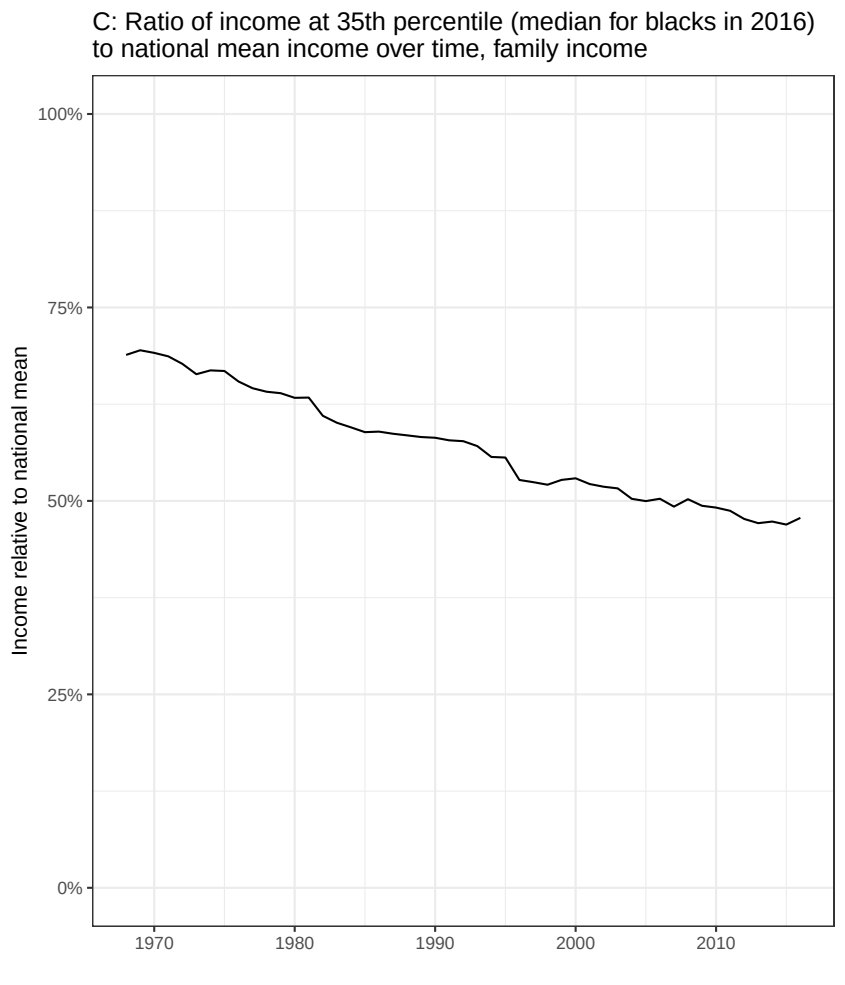

Allocation processes have become less racially stratified, but that improvement was entirely undone by changes to the marginal distribution

Many social problems require marginal shifts

- In both examples, improvements to allocation processes could not overcome worsening marginal distributions
- Similar dynamics exist in many important issue areas:
  - Education
  - Health care
  - Housing
- Mechanism design can help improve marginal distributions

Mechanism Design 4 Better Marginals

- Our current marginal distributions result in part from **poor social decision making**. Across many domains, current policy differs markedly from majority opinion:
  - Support for single player health care: 63%
  - Support for marijuana legalization: 60%
  - Support for limiting political spending by individuals: 77%
- Aggregating from individual preferences to social choices is hard, and leaves room for manipulation by well-organized or wealthy interests
- Mechanism design can help!

MD4BM 1: Participatory budgeting

Preference Elicitation For Participatory Budgeting

GERDUS BENADE, SWAPRAVA NATH, and ARIEL D. PROCACCIA, Carnegie Mellon University
NISARG SHAH, Harvard University

Participatory budgeting enables the allocation of public funds by collecting and aggregating individual preferences; it has already had a sizable real-world impact. But making the most of this new paradigm requires a rethinking of some of the basics of computational social choice, including the very way in which individuals express their preferences. We analytically compare four preference elicitation methods — knapsack votes, rankings by value or value for money, and threshold approval votes — through the lens of *implicit utilitarian voting*, and find that threshold approval votes are qualitatively superior. This conclusion is supported by experiments using data from real participatory budgeting elections.

CCS Concepts: • **Computing methodologies** → **Multi-agent systems**; • **Applied computing** → **Economics**; • **Theory of computation** → *Approximation algorithms analysis*;

ACM Reference format:
Gerdus Benade, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. 2017. Preference Elicitation For Participatory Budgeting. *J. ACM* 1, 1, Article 1 (January 2017), 27 pages.
DOI:

MD4BM 2: Voting systems

THE WALL STREET JOURNAL.

California Primaries Could Shut Parties Out of Key November Midterms

Outcome of all-party gubernatorial primary may sway turnout for key races in November

The Democrats have a "blue wave" of momentum building for the 2018 midterms, thanks to a motivated base, success in special elections and a low approval rating for President Trump. Will that be enough to take back the House and the Senate?

By Janet Hook and Reid J. Epstein

Updated June 6, 2018 12:01 a.m. ET

Menu | Subscribe | Subscriber Log in

ELECTION 2018 >

Posted **Yesterday at 1:12 PM** | Updated **June 13**

Passage of Maine's ranked-choice ballot question seen as 'turning point' in election history

Supporters cheered their referendum victory on Wednesday and said it has the chance to reform politics.

BY **ERIC RUSSELL** STAFF WRITER

MD4BM 3: Easier group decision-making

Getting large groups of people to reach consensus on a course of action is hard, and a major barrier to political organizing and cooperative ownership. Internet platforms with good mechanism design can help overcome this

MD4BM 4: Determining social preferences

Rules for Choosing Societal Tradeoffs

Vincent Conitzer and **Rupert Freeman** and **Markus Brill** and **Yuqian Li**
Department of Computer Science
Duke University
Durham, NC 27708, USA
{conitzer,rupert,brill,yuqian}@cs.duke.edu

Abstract

We study the societal tradeoffs problem, where a set of voters each submit their ideal tradeoff value between each pair of activities (e.g., "using a gallon of gasoline is as bad as creating 2 bags of landfill trash"), and these are then aggregated into the societal tradeoff vector using a rule. We introduce the family of distance-based rules and show that these can be justified as maximum likelihood estimators of the truth. Within this family, we single out the logarithmic distance-based rule as especially appealing based on a social-choice-theoretic axiomatization. We give an efficient algorithm for executing this rule as well as an approximate hill climbing algorithm, and evaluate these experimentally.

Introduction

There are many actions that we take in life that are generally agreed to have some negative effects on society. For example, consider actions with environmental downsides, such as using gasoline, creating landfill trash, and clearing forest, to name a few. Which of these is worse? To answer this, clearly one would first need to know *how much* gasoline is used, etc. This then suggests the following type of question: how many bags of trash are as bad as using one gallon of gasoline? Knowing the answer to this question could be useful to policy makers as well as to socially minded individuals or companies who are looking to reduce their environmental footprint in the most efficient way. However, since the environmental effects of these actions are different, it seems unvoter expresses for each pair of activities her ideal tradeoff value between those two. For example, a voter may feel that a gallon of gasoline corresponds to two bags of trash.

From a social-choice-theoretic viewpoint, when aggregating numbers, one submitted per voter, choosing the *median* is particularly compelling. When preferences are single-peaked, this results in choosing the Condorcet winner, and the corresponding voting rule is group-strategyproof. However, Conitzer, Brill, and Freeman (2015) pointed out that simply taking the median for each pair of activities can result in the aggregate tradeoffs being *inconsistent*, in the sense that the chosen tradeoff between a and c is not equal to the product of the tradeoff between a and b and the tradeoff between b and c. See the example in Figure 1, where a voter's tradeoffs are represented by a graph with its edges labeled with tradeoff values (e.g., voter 1 believes a gallon of gasoline is as bad as 2 bags of trash). This paradox is reminiscent of judgment aggregation paradoxes where taking majority on all individual issues results in a logically inconsistent aggregate judgment (Kornhauser and Sager 1993).

So what are we to do? We insist that the aggregate tradeoffs be consistent; if not, then it is not clear how to use them to guide decisions involving three or more activities. That means we must judiciously deviate from the median in some cases, but presumably we want to deviate as little as possible. The topic of this paper is how to make this precise.

We introduce a class of rules for this context that we call *distance-based rules*. We prove that these rules choose the median when there are only two activities and can be inter-

Liquid Democracy: An Algorithmic Perspective

Anson Kahng
Computer Science Department
Carnegie Mellon University

Simon Mackenzie
Computer Science Department
Carnegie Mellon University

Ariel D. Procaccia
Computer Science Department
Carnegie Mellon University

Abstract

We study liquid democracy, a collective decision making paradigm that allows voters to transitively delegate their votes, through an algorithmic lens. In our model, there are two alternatives, one correct and one incorrect, and we are interested in the probability that the majority opinion is correct. Our main question is whether there exist delegation mechanisms that are *guaranteed* to outperform direct voting, in the sense of being always at least as likely, and sometimes more likely, to make a correct decision. Even though we assume that voters can only delegate their votes to better-informed voters, we show that *local* delegation mechanisms, which only take the local neighborhood of each voter as input (and, arguably, capture the spirit of liquid democracy), cannot provide the foregoing guarantee. By contrast, we design a non-local delegation mechanism that does provably outperform direct voting under mild assumptions about voters.

1 Introduction

"Even if it were possible for every citizen to learn everything they could possibly know about every political issue, people who did this would be able to do little else, and massive amounts of time would be wasted in duplicated effort. Or, if every citizen voted but most people did not take the time to learn about the issues, the results would be highly random and/or highly sensitive to overly simplistic public relations campaigns."

By contrast, under liquid democracy, voters who did not invest an effort to learn about the issue at hand (presumably, most voters) would ideally delegate their votes to well-informed voters. This should intuitively lead to collective decisions that are less random, and more likely to be correct, than those that would be made under direct democracy.

Our goal is to rigorously investigate the intuition that liquid democracy "outperforms" direct democracy from an algorithmic viewpoint. Indeed, we are interested in *delegation mechanisms*, which decide how votes should be delegated based on how relatively informed voters are, and possibly even based on the structure of an underlying social network

What these approaches have in common

- Focus is on making better social decisions so that we can get the set of options that we want, instead of trying to choose among the bad options that we have
- This involves both:
  - Technical challenges – determining the optimal social choice given disparate preferences
  - Organizational challenges – overcoming coordination problems and distortionary power centers
- Mechanism design can contribute to overcoming both sets of challenges

Thank you!

rmanduca@g.harvard.edu

How marginal distributions shape mobility

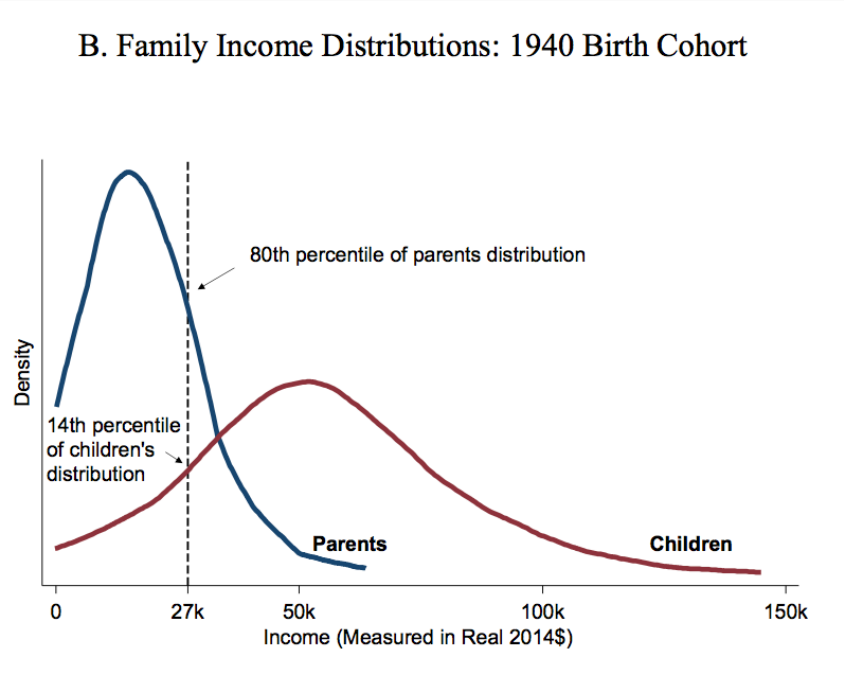

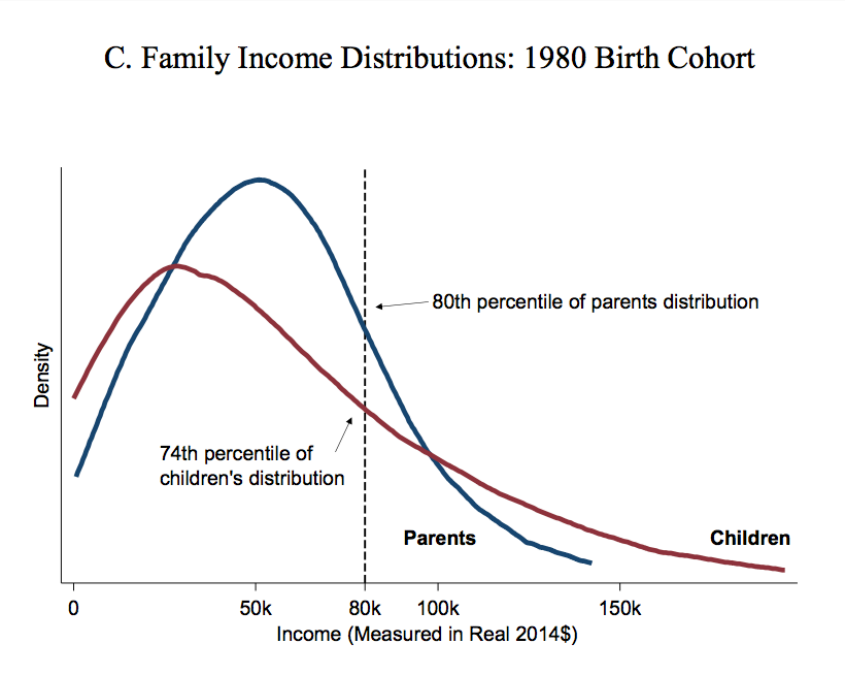

How marginal distributions shape mobility