Essays in Development Economics

HARVARD UNIVERSITY
Graduate School of Arts and Sciences

DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the

Department of Economics

have examined a dissertation entitled

Essays in Development Economics

presented by Moya Chin

candidate for the degree of Doctor of Philosophy and hereby certify that it is worthy of acceptance.

Signature

Typed name: Prof. Nathan Nunn

Signature

Typed name: Prof. Melissa Dell

Signature

Typed name: Prof. Vincent Pons

Signature

Typed name: Prof. Gautam Rao

Date: May 11, 2020

Essays in Development Economics

A dissertation presented

by

Moya Chin

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

May 2020

Dissertation Advisor:
Nathan Nunn

Author:
Moya Chin

Essays in Development Economics

Abstract

This dissertation explores the political and economic determinants of why households in developing countries have different economic outcomes. The first chapter documents how different electoral rules can lead to different patterns of voting and public goods provision. This chapter uses data from Brazilian municipal elections between $1996-2016$ to show that two-round systems, compared to single-round systems, lead politicians to appeal more broadly to the electorate and to provide public goods more broadly. The second chapter quantifies the effect of the anti-sweatshop movement in the 1990s on female workers in textile factories in Indonesia. The third chapter studies the implementation of a proxy-means testing system in Colombia. This chapter uses machine learning techniques to evaluate the performance of the government's targeting system and documents why households may manipulate their eligibility for social programs.

Contents

List of Tables

List of Figures

Acknowledgments

PhD dissertations only have one author, but the number of people who made this possible cannot fit on one page. If I somehow missed you, I apologize but thank you.

I would not have gone to graduate school were it not for the opportunities and mentorship from my professors at MIT. Thank you for giving me a chance, even when graduate school was not a decision I made until my senior year.

This dissertation would not have been possible without the encouragement and support from my advisors: Nathan Nunn, Melissa Dell, Gautam Rao, and Vincent Pons. There was so much I learned from you, but you also gave me the confidence to be proud of my work.

Thank you to all of the friends who have helped me through graduate school. Thank you to the lunch group, for making lunches less lonely: Nir Hak, Natasha Sarin, Omar Barbiero, Edoardo Acabbi, Paulo Costa, Talia Gillis, Alejandro Lagomarsino, Tzachi Raz, and Jonathan Roth. Thank you to my other econ friends, for helping me stay sane: Lisa Abraham, Sunny Lee, Dev Patel, Adrienne Sabety, and so many others. Thank you to my wonderful lifelong friends outside of the department, for only occasionally asking me when I was going to get a job.

To my mother, father, and brother: thank you for your support and for making me the curious person I am today.

And lastly, thank you to my partner, Arturo Gonzalez, without whom I would not be here today. Thank you for your love and for being my biggest cheerleader.

Introduction

This dissertation explores the political and economic determinants of why households in developing countries have different economic outcomes. The first chapter documents how different electoral rules can lead to different patterns of voting and public goods provision. Electoral rules determine how voters' preferences are aggregated and translated into political representation, and their design can lead to the election of representatives who represent broader or narrower constituencies. This chapter examines how single- and two-round elections in Brazil affect municipal mayoral races using a regression discontinuity design. Two-round elections use two rounds of voting to elect a winner, ensuring that the eventual winner must obtain at least 50% of the vote. Theoretically, this may provide incentives for candidates to secure a broader base of support. Consistent with this, I show that in two-round systems, candidates represent a more geographically diverse group of voters, more resources are allocated to public schools, and there is less variance in resources allocated to public schools across the municipality. I find evidence suggesting that these effects are driven by strategic responses of candidates, rather than differential entry of candidates into races. The findings suggest that two-round systems can lead candidates to secure broader voter bases and subsequently exhibit less political favoritism when implementing policy.

The second chapter quantifies the effect of the anti-sweatshop movement on the textile, footwear, and apparel industry in Indonesia and the gendered impacts on female workers. In the 1990s, companies such as Nike, Adidas, and Reebok became targets of anti-sweatshop activism and implemented reforms to improve labor standards in their overseas factories. I use a difference-in-differences strategy, by comparing differences among workers before and after the anti-sweatshop movement in either targeted sectors or in locations with targeted firms. I show that the movement succeeded in improving labor standards and shifted the educational composition of workers, suggesting that the movement affected the entry and exit of workers into the sector. I find some, but limited effects on marriage and fertility outcomes for female workers. Female workers are less likely to be married and to have had children, but I do not find effects on contraception use nor school attendance rates of their children. These results suggest that labor reforms in industries with predominantly female labor forces can generate better work environments but that these improvements lead to modest gains in female autonomy and empowerment.

The third chapter studies the implementation of a proxy-means testing system in Colombia. I investigate the extent to which households manipulate their eligibility for a social program. Eligibility is determined by a poverty score that is calculated from answers to a household survey, the algorithm for which was released four years after the start of the program. Because proxy-means testing systems can potentially predict household poverty poorly, households have incentives to manipulate their eligibility. I document a significant discontinuity at the eligibility threshold and that one method households use to manipulate their eligibility is by having their poverty score overwritten. I then use machine-learning techniques to predict households' actual poverty level. First, I find that the government poverty score predicts household poverty poorly and performs particularly poorly for households far below the poverty line. Second, I find that *before* the public release of the algorithm, households who manipulate their score are richer than households with the same score, but that *after* the public release of the algorithm, households who manipulate their scores are poorer. These findings suggest that not all proxy-means testing systems predict household poverty well and when they do not, households can self-target by manipulating their eligibility.

Chapter 1

When Do Politicians Appeal Broadly? The Economic Consequences of Electoral Rules in Brazil

1.1 Introduction

Electoral rules determine how voters' preferences are aggregated and translated into political representation, and their design can lead to the election of representatives who represent broader or narrower constituencies. This is particularly important given the evidence that more inclusive political institutions are beneficial for long-term growth (Acemoglu and Robinson, 2008; Acemoglu et al., 2014; Bardhan and Mookherjee, 2006). This paper studies how one difference in electoral rules – namely, if elections feature a single or two rounds – affects the extent to which elected representatives appeal to a broader constituency and how this, in turn, affects the overall level and distribution of public goods.

I take advantage of a unique policy in Brazil that assigns a municipality's electoral rule based on a threshold of 200,000 registered voters. Municipalities below this threshold elect their mayor in a single-round election, and municipalities above this threshold elect their mayor in a two-round election. Using a regression discontinuity framework that compares municipalities close to the threshold, this paper provides causal empirical evidence that politicians elected under two-round systems secure geographically broader bases of support, provide more resources to public schools, and allocate these resources more equitably.

Two-round systems are the most widely used rule in democratic presidential elections: 64.0% of elections under presidential systems used the two-round system between 2000 and 2010 (Bormann and Golder, 2013 and Figure 1.1). Together, single- and two-round systems account for 86.7% of elections in this period. In a single-round system, voters vote once and the candidate with the most votes wins. In a two-round system, voters first vote and, if no candidate receives a majority, vote a second time between the top two candidates.[1] This difference generates two important distinctions. First, two-round systems require winners to attain a vote share above 50%. Second, the existence of a second round effectively limits the number of candidates (Lizzeri and Persico, 2005). Even in the first round, the top candidate effectively only needs to be concerned with the runner-up, who poses the threat of either victory in the first round or opposition in the second. Because of these distinctions, it has been argued that two-round systems incentivize candidates to secure a broader base of support and legitimize the winner's position once elected (Bouton, 2013; Bouton and Gratton, 2015).[2] The intuition behind this is that the rules imposed in a two-round system make it more difficult for politicians to win with policies that appeal to a narrow group of voters.

To the extent that politicians commit to their campaign promises, the policies they offer in order to win the election may translate into economic consequences. This paper tests the hypothesis that electoral rules affect the level of public goods provided and the manner in which they are allocated *across* the electorate. When electoral rules make it more difficult for politicians to win with a narrow constituency, this reduces their incentive to provide public goods supported by a narrow constituency once in office.

To guide the empirical analysis and explain this intuition, I propose a model of electoral competition to illustrate how two-round systems create incentivizes for politicians to appeal to broader groups of voters, provide more public goods, and allocate public goods differently. I adapt a standard probabilistic voting model where candidates offer policy proposals that (i) specify the overall size of the government budget and (ii) target government resources to specific localities within a municipality, as in Genicot et al. (2018). These policy proposals

[1] This is the case in Brazil and most countries with two-round systems. Outside of single- and two-round systems, a small number of countries use qualified two-round, electoral college, and alternative vote.

[2] While not a focus of this paper, there is also a large literature arguing that two-round systems allow voters to vote more sincerely in the first round (see Bouton et al., 2019 for a review) and to better communicate their policy preferences to candidates (Piketty, 2000).

Figure 1.1 Electoral rules in presidential systems around the world

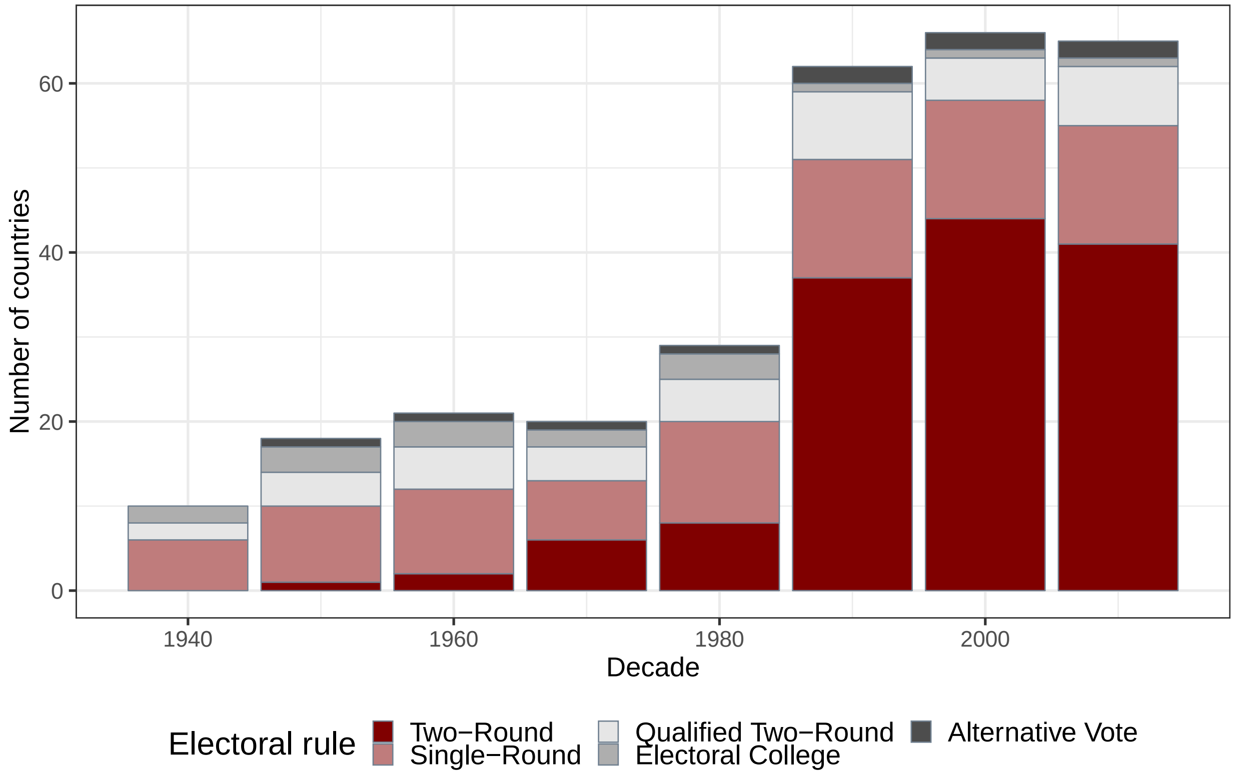

Single-round plurality (light red) and *Two-round majority* (dark red) correspond to the single- and two-round system in Brazil. Vertical axis corresponds to the number of countries which held an election within the specified time period under the specified electoral rule. Data encompasses 497 presidential elections across 71 countries with democratic regimes and presidential systems, between 1946-2016. *Source:* Bormann and Golder (2013).

are announced prior to the election and, in a two-round system, are binding between rounds. In these elections, the top two candidates must contend with a small, third candidate who is non-strategic and commits to allocating all resources to a single locality.

In the model, two-round elections lead to different outcomes, because the second round raises the marginal return for the top two candidates to allocate resources to all localities. First, in a two-round election, to win in the first round, candidates must attain not only the most votes, but must attain at least 50% of the vote. If they do not, candidates must compete in a second round. This conditionality raises the value of each vote. As a result, candidates promise a higher overall budget. Second, candidates in a single-round election face low incentives to allocate resources to the locality dominated by the third candidate. In contrast, in a two-round election, the top two candidates in the first round must look ahead to the possibility of a second round in which the third candidate will not be present. As a result, candidates promise more resources to the locality that the third candidate appeals to, even if they are unlikely to attract those votes in the first round. This increased allocation generates a force to reduce inequality in the allocation of government resources.

I test these predictions empirically across six municipal elections, between 1996 and 2016, in Brazil. I exploit the fact that Brazil's electoral authority assigns the electoral rule for mayor in each election based on the number of registered voters in the municipality and whether this lies above or below an arbitrary threshold. To estimate the causal impact on electoral and economic outcomes after each election, I compare municipalities just above the registered voter threshold with those just below, in a regression discontinuity design. I obtain three main empirical results.

First, candidates in two-round elections receive broader geographical support. Using each candidate's vote count at each polling station, I measure the geographic distribution of voters for specific candidates in the first round in two ways: (i) indices of voter concentration to quantify the *overall* level at which voters are geographically concentrated, and (ii) the standard deviation of candidates' vote shares to quantify a *candidate*-level measure of geographic concentration. In two-round municipalities, voters are less geographically concentrated, corresponding to a reduction that is 27.4–45.6% of the average level of concentration in single-round municipalities. The decrease in concentration only occurs among the top two candidates, indicating that the electoral rule mainly affects the candidates with a chance of winning. I show that this is not driven by the increased number of candidates in two-round elections.[3] These results suggest that two-round elections lead to greater inclusiveness: Elected candidates represent a geographically broader constituency. In turn, I find inclusiveness along another dimension: Voters are more engaged in the political process. While turnout is

[3]The effect of the electoral rule on the number of candidates is known as Duverger's Law, which states that single-round elections will lead to a two-party system, while two-round elections will lead to a multi-party system. This has been formalized, and sometimes challenged, in recent literature (Bouton, 2013; Bouton and Gratton, 2015; Callander, 2005; Cox, 1997; Fujiwara, 2011; Osborne and Slivinski, 1996).

unaffected, as turning out to vote is mandatory in Brazil, I find that the number of blank and invalid ballots is significantly lower in two-round elections.

Second, once in office, politicians elected under two-round systems provide and allocate public goods differently. I find effects on both the *level* and *distribution* of municipal resources. I look at the provision of a local public good that can be geographically targeted and is the main public good completely under the jurisdiction of the municipality: municipal education. To quantify provision, I measure the level of resources present in public schools in the municipality. I find that (i) public schools in two-round municipalities measure 0.057–0.081 percentiles higher in the national distribution of resources, and (ii) the standard deviation of resources across schools is lower. Schools with the fewest resources in the municipality benefit the most from these additional resources. When politicians secure broader bases of support, they provide more public goods and distribute these resources more evenly across the municipality.

Third, I find evidence that educational outcomes are improved in two-round municipalities. Specifically, drop-out rates are lower and literacy rates higher among cohorts of school age during the electoral term. I find limited improvements along more downstream economic indicators. Two-round municipalities have fewer low-income households, but there are no significant improvements in average income, employment, or night lights. These results suggest that electoral rules impact direct policy outcomes, but have limited effects on more downstream economic outcomes.

Turning to mechanisms, I do not find that two-round elections cause differential selection of candidates. More specifically, I do not find that different types of candidates enter the races, nor do different types of candidates win the elections. Candidates in two-round elections are not observably different in terms of demographic characteristics, place of birth, educational attainment, or previous occupation. Winners also do not differ along these observable characteristics. More candidates enter the race in two-round elections – many of these additional candidates are represented by smaller parties and have run in previous elections. However, these candidates are not more likely to win the election.

I find suggestive evidence that the effects I find are driven by the different strategic incentives candidates face in two-round elections. These strategic incentives may affect politicians' behavior once in office through two mechanisms. First, candidates may adopt different strategies during the campaign, which, to the extent that politicians fulfill campaign promises, lead to different behaviors in office. This is in line with the model's predictions. Second, politicians may behave differently in office in anticipation of the forthcoming election and what electoral rule it will follow. If the second mechanism is present, candidates facing different reelection incentives should behave differently. I rule out the second mechanism, as I do not find different treatment effects among mayors facing different reelection incentives. Instead, I find that different outcomes are due to candidates adopting different strategies during the campaign. First, in two-round elections, geographic concentration decreases between the first and second round, suggesting that candidates adopt strategies between rounds to consolidate their voter bases.[4] Second, candidates in two-round elections finance their campaigns differently: They rely less on donations from corporations. To the extent that corporations represent narrower swaths of the electorate, this suggests that candidates in two-round elections adopt strategies that appeal more broadly to individuals rather than corporations. These findings suggest that candidates in two-round elections offer different policy platforms and build more geographically diverse constituencies. When their supporters are less geographically concentrated, mayors allocate public goods in a less geographically concentrated manner.

This study builds on a large, mostly theoretical literature studying the role of electoral rules on political incentives. These studies document how different rules impact electoral accountability (Persson et al., 2003) and personal vote-seeking (Carey and Shugart, 1993).[5]

[4]There is empirical evidence that candidates qualifying for the second round rally votes from supporters of the candidates eliminated after the first round. Pons and Tricaud (2018) find that in France, the qualification of a third candidate in the second round reduces the top two candidates' vote share, indicating that when the third candidate is not present in the second round (as is always the case in Brazil), the top two candidates capture votes from the third placed candidate's supporters.

[5]See Persson and Tabellini (2004) and Persson et al. (2004) for a review of other consequences of electoral

Importantly, these studies highlight how electoral rules incentivize politicians to appeal to broader groups of voters (Myerson, 1993), provide public goods with broader benefits (Lizzeri and Persico, 2001, 2005; Persson and Tabellini, 1999), and target public spending to specific groups of voters (Genicot et al., 2018; Milesi-Ferretti et al., 2002). With the exception of Lizzeri and Persico (2005), these studies do not compare single- to two-round systems. In Lizzeri and Persico (2005), politicians in proportional systems provide broad public goods at a higher rate than targetable goods when there is less political competition. They extend their model to argue that two-round elections should lead to higher provision of broad public goods, because the second round effectively limits the number of candidates. This paper provides a theoretical framework that models public goods provision in single- and two-round systems, generates predictions on both the overall level and the exact allocation, and tests these predictions empirically.

This paper adds to a growing empirical literature providing causal evidence on the impacts of local electoral rules. These studies, which compare proportional and single-round systems in addition to single- and two-round systems, have measured the impact on electoral outcomes and fiscal expenditures (Chamon et al., 2018; Cipullo, 2019; Eggers, 2013; Fujiwara, 2011; Pellicer and Wegner, 2013). Chamon et al. (2018) and Fujiwara (2011) also study the Brazilian context. Of particular interest are Bordignon et al. (2016), who compare single- and two-round systems in Italian municipalities and find that municipalities under two-round systems exhibit more policy moderation, as measured by the volatility of a municipal tax rate across elections. In contrast to Bordignon et al. (2016), I study not only an *aggregate* policy outcome – the overall level of public goods provision – but also the *allocation* of this policy across the electorate. This paper's contribution is to provide evidence that electoral rules have economic consequences, both on the level of public goods provision and how these public goods are distributed.

More broadly, this paper connects to a literature examining inequalities in the allocation

rules.

of state resources. A large literature documents the role of political factors in creating these inequalities – in particular, how politicians politically favor certain subgroups, such as those of the same ethnicity or partisanship. A key insight that emerges is that the extent to which politicians practice political favoritism is reduced when political institutions are stronger, elections are more competitive, and citizens are more broadly engaged in the electoral process (Burgess et al., 2015; Fujiwara and Wantchekon, 2013; Hodler and Raschky, 2014). Notably, Golden and Min (2013) emphasize the importance of policy responsiveness to voter preferences. Electoral rules serve as a key channel through which voter preferences are translated into policy outcomes. The final contribution of this paper is to demonstrate the role of another factor in political favoritism, the electoral rule, and the incentives it creates for politicians to broaden their appeal.

The remainder of this paper is organized as follows. Section 1.2 presents a theoretical framework for single- and two-round systems. Section 1.3 describes the context, and Section 1.4 describes the empirical strategy. Section 1.5 presents the results on electoral and economic outcomes. Section 1.6 discusses the results and mechanisms. Section 1.7 concludes.

1.2 Theoretical framework

I present a stylized model to illustrate how two-round elections create incentives for politicians to secure a broader base of support and provide public goods differently. My model of electoral competition adapts a standard probabilistic voting model (Burden, 1997; Lindbeck and Weibull, 1987; Persson and Tabellini, 2000) and follows the setup in Genicot et al. (2018) by allowing for targeting of government interventions to specific localities within a municipality. I extend this model by (i) introducing a third non-strategic candidate who appeals to a single locality, (ii) allowing candidates to exert effort to increase the municipal budget, and (iii) adapting it to the context of single- and two-round elections.

The intuition is that two-round elections perform two functions. First, the winner must attain a vote share above 50%. This makes it difficult for politicians who appeal to a minority of the electorate to gain enough votes to win the election. It also raises the marginal value of *every* vote and results in candidates exerting more effort to increase the overall government budget. Second, the existence of a second round effectively limits the number of candidates. Candidates who expect to gather enough first round votes to qualify for the second round but less than the 50% required to win in the first round must behave as though a second round will occur where only one other candidate will stand. This incentivizes candidates to offer policy proposals that appeal to all localities, even those that are likely to vote for another candidate in the first round. While this effect runs counter to a literature documenting the positive impacts of political competition, in a single-round election, higher electoral competition incentivizes candidates to appeal to narrower groups and ignore other voters.[6]

1.2.1 The environment

Consider an election with three politicians and J localities within a municipality. Politicians are indexed by $c \in \{A, B, C\}$, and localities are indexed by $j \in \{1, 2, \cdots, J\}$ where $J \geq 3$. Each locality has a continuum of voters of mass $1/J$.

Prior to election day, politicians simultaneously announce a platform that describes (i) the total government budget, G^c, and (ii) the amount of the government budget to be allocated to each locality, $\mathbf{q}^c = (q_1^c, q_2^c, \ldots, q_J^c)$, where $q_j^c \geq 0$. The politician's budget constraint is:

$$\sum_{j=1}^{J} q_j^c \leq G^c$$

Since each locality has the same number of voters, each voter receives the same fraction of the government budget allocated to their locality. I assume without loss of generality that

[6] A theoretical literature argues that higher electoral competition has the negative effect of incentivizing candidates to focus on narrower groups. Myerson (1993) shows that candidates offer more unequal campaign promises when electoral competition increases in elections using rank-scoring rules. Lizzeri and Persico (2005) extend this model to other electoral rules and introduce a public good, finding a negative effect of political competition on the equality of campaign promises and public good provision.

voters care about the total amount allocated to their locality, q_j^c. In promising a certain budget, politicians face a cost that is quadratic in the size of the budget:

$$\frac{1}{2}\kappa\left(G^c\right)^2$$

for a constant κ. Platforms are binding for politicians between rounds and after the election.[7]

To make solving a three-candidate model more tractable, I assume that the third candidate C is a non-strategic candidate with the following platform:[8]

$$\mathbf{q}^C = \left(0, 0, \ldots, 0, G^C\right)$$

There are several ways to interpret candidate C. In my model, because voters are partitioned into geographic localities, C is a candidate whose supporters are all located in the same geographic area. However, my model easily translates into other interpretations of candidate C. For example, C may be a candidate whose supporters share a common trait and vote for her due to descriptive representation. These traits could include geography, but also other dimensions such as age or race. Another possibility is that C is a single-issue candidate who attracts voters who only care about that issue. In any of these cases, C should be viewed as a small candidate and A and B as front-runners in relation to C. This also empirically matches elections in Brazil. Third placed candidates receive on average 11.9% of the vote, and the vote spread between the second and third placed candidate is on average 23.8%.

Following Genicot et al. (2018), voters in locality j have preferences over government spending. They obtain utility $u_j(q_j)$ from government spending q_j, where $u_j(q_j)$ is strictly increasing and concave in q_j. In addition to the policy component of voters' preferences,

[7] This is not unrealistic, as the time between rounds is often short compared to the length of the campaign. In Brazil, the second round is three weeks after the first. Intuitively, my model translates into contexts where this assumption is relaxed as long as there is some continuity between the two rounds – first, if voters' second round vote depends on a candidate's policy proposal in both rounds; and second, if candidates can change their policy proposals between rounds but are constrained in the extent to which their proposals can change.

[8] I assume that G^C is the highest offer in locality J, ie. that candidate A and B's equilibrium allocations to J are smaller than G^C. See Appendix A.1.9.

there is an individual shock v_i and a municipality shock δ toward candidate A, which are independently and uniformly distributed:

$$v_i \sim U\left[-\frac{1}{2\psi}, \frac{1}{2\psi}\right] \qquad \delta \sim U\left[-\frac{1}{2\gamma}, \frac{1}{2\gamma}\right]$$

The individual shock, v_i, captures idiosyncratic voter preferences towards candidate A. The municipality shock, δ, captures any political dimensions that swing voters in the municipality as a whole toward candidate A, such as economic shocks, and is independent across rounds.

Voters cast a ballot for the politician who offers them the highest payoff. In localities $j \in \{1, \ldots, J-1\}$, ie. where candidate C has not allocated resources, this amounts to voting for either A or B.[9] In locality J, ie. where candidate C is dominant, voters randomize between voting for C with probability $1-\alpha$ and for either A or B with probability α, depending on whether A or B offers the higher payoff, where $0 < \alpha < 1$.[10]

Thus, in general, voters will vote for A if and only if:

$$u_j(q_j^B) \leq u_j(q_j^A) + v_i + \delta \tag{1.1}$$

In localities $j \in \{1, \ldots, J-1\}$, all voters for whom this is true vote for A.[11] In locality J, a fraction α of voters for whom this is true vote for A. I assume that the marginal utility of locality J relative to that of the other localities, $u'_J(q)/u'_j(q)$, is not too large.[12]

The electoral rules follow those in Brazil. In a single-round system, the candidate with

[9] Since $u_j(\cdot)$ is strictly increasing, candidates A and B will invest a non-zero amount in these localities, and voters will always vote for either A or B.

[10] Assuming that $\alpha > 0$ performs two functions. First, it guarantees that candidate C gains strictly less than 1/3 in vote share and so never has the most votes. Candidate C will also never have the most votes if $J > 3$. Second, it guarantees a non-zero first order condition for locality J in the single-round election, which allows a direct comparison between single- and two-round elections. This assumption can be relaxed and will yield the same predictions; see Appendix A.1.10.

[11] For localities $j \in \{1, \ldots, J-1\}$, it is also a requirement that $u_j(q_j^A) + v_i + \delta \geq u_j(q_j^C)$. Because $u_j(q_j^B) \geq u_j(0) = u_j(q_j^C)$, condition (1.1) is sufficient for voters to vote for A in these localities.

[12] In other words, for all $j \in \{1, \ldots, J-1\}$, I assume that $u'_J(q)/u'_j(q) < \dfrac{\psi J + (1-\alpha)\gamma}{\alpha\psi J + (1-\alpha)\gamma}$, for all positive q.

the most votes wins. In a two-round system, if no candidate attains more than 50% of the vote in the first round, a second round occurs with the top two candidates.

1.2.2 Preliminaries

Let π_t^c be candidate c's total vote share in the municipality in round $t \in \{1, 2\}$. Let $\Delta u_j^{cd} \equiv u_j(q_j^c) - u_j(q_j^d)$ be the difference in utility in locality j between candidate c and d's offers.

Vote shares with three candidates.– This section applies to single-round elections and the first round of two-round elections.

As in Genicot et al. (2018), I assume that (i) there are voters to be swung in every locality and (ii) all localities are contestable.[13] The probability that candidate A attains a vote share above θ is given by:

$$Pr\left(\pi_1^A \geq \theta\right) = \frac{1}{2} + \frac{\gamma}{\psi}\left[\frac{1}{2} - \left(\frac{J}{J-1+\alpha}\right)\theta + \left(\frac{\psi}{J-1+\alpha}\right)\left(\sum_{j=1}^{J-1} \Delta u_j^{AB} + \alpha \Delta u_J^{AB}\right)\right] \quad (1.2)$$

For a detailed derivation and expressions for other candidates, see Appendix A.1.2.

With three candidates, I assume that candidate C always receives the lowest vote share: candidate C never wins a single-round election nor makes it to the second round in a two-round election (see Appendix A.1.5). This simplifies candidate A and B's maximization problem, and shuts down the channel where candidate C poses a different threat to electoral defeat.

Vote shares with two candidates.– This section applies to the second round of two-round elections (as mentioned, candidate C never makes it to the second round).

[13](i) states that $0 < \pi_{j1}^A < 1$. See Appendix A.1.1. (ii) states that $0 < Pr\left(\pi_1^A \geq \theta\right) < 1$ and $0 < Pr\left(\pi_1^B \geq \theta\right) < 1$. See Appendix A.1.4.

The probability that candidate A attains a vote share above θ is given by:

$$Pr\left(\pi_1^A \geq \theta\right) = \frac{1}{2} + \frac{\gamma}{\psi}\left(\frac{1}{2} - \theta + \frac{\psi}{J}\sum_{j=1}^{J}\Delta u_j^{AB}\right) \tag{1.3}$$

For a detailed derivation and expressions for candidate B, see Appendix A.1.3.

Candidates' maximization problem.– For both candidate A and B, her payoff is 1 if she wins the election and 0 otherwise, minus the effort cost incurred during the campaign. Candidates maximize their expected payoff, so this amounts to maximizing the probability of winning minus the effort cost.

1.2.3 Equilibrium strategies in a single-round election

In a single-round election, candidate C attains a vote share of $\frac{1-\alpha}{J}$, so the probability of winning is the probability of attaining a vote share above $\frac{1}{2}\left(1 - \frac{1-\alpha}{J}\right)$. Using equation (1.2), for candidate A, this is equivalent to solving the following maximization problem:

$$\max_{G^A, \mathbf{q}^A = \left(q_1^A, \ldots, q_J^A\right)} \frac{1}{2} + \left(\frac{\gamma}{J-1+\alpha}\right)\left(\sum_{j=1}^{J-1}\Delta u_j^{AB} + \alpha\Delta u_J^{AB}\right) - \frac{1}{2}\kappa\left(G^A\right)^2 \qquad \text{s.t.} \sum_j q_j^A \leq G^A$$

In a single-round system, in equilibrium, candidates will promise less resources to locality J in comparison to the other localities (see Appendix A.1.6):

Prediction 1.1. *In a single-round election, for all $j \in \{1, \ldots, J-1\}$, we have that $q_j^A > q_J^A$.*

Prediction 1.1 results from the fact that, for a given level of government spending, candidates' marginal return to allocating resources to locality J is lower than in other localities, leading to less resources promised to locality J in equilibrium.

1.2.4 Equilibrium strategies in a two-round election

In a two-round election, the probability of winning is the probability of attaining a vote share above $\frac{1}{2}$ in the first round or, if a second round occurs, of attaining a vote share above $\frac{1}{2}$ in the second round:

$$
\begin{aligned}
& Pr(A \text{ wins in 1st round}) + Pr(\text{second round occurs}) \cdot Pr(A \text{ wins 2nd round}) \\
& = Pr\left(\pi_1^A \geq \frac{1}{2}\right) + \left(1 - Pr\left(\pi_1^A \geq \frac{1}{2}\right) - Pr\left(\pi_1^B \geq \frac{1}{2}\right)\right) Pr\left(\pi_2^A \geq \frac{1}{2}\right)
\end{aligned}
$$

Using equations (1.2) and (1.3), this corresponds to the following maximization:

$$
\begin{aligned}
\max_{G^A, \mathbf{q}^A = \left(q_1^A, \ldots, q_J^A\right)} \quad & \left(\frac{1}{2} + \frac{\gamma}{\psi}\left[\frac{1}{2}\left(\frac{\alpha - 1}{J - 1 + \alpha}\right) + \left(\frac{\psi}{J - 1 + \alpha}\right)\left(\sum_{j=1}^{J-1} \Delta u_j^{AB} + \alpha \Delta u_J^{AB}\right)\right]\right) \\
& + \frac{\gamma}{\psi}\left(\frac{1 - \alpha}{J - 1 + \alpha}\right)\left[\frac{1}{2} + \frac{\gamma}{J}\sum_{j=1}^{J} \Delta u_j^{AB}\right] - \frac{1}{2}\kappa\left(G^A\right)^2 \\
\text{s.t.} \quad & \sum_j q_j^A \leq G^A
\end{aligned}
$$

In a two-round system, in equilibrium, candidates will promise less resources to locality J in comparison to the other localities (see Appendix A.1.7):

PREDICTION 1.2. *In a two-round election, for all $j \in \{1, \ldots, J-1\}$, we have that $q_j^A > q_J^A$.*

As in the single-round system, for a given level of government spending, candidates' marginal return to allocating resources to locality J is lower than in other localities, leading to less resources promised to locality J in equilibrium.

1.2.5 Comparing single- to two-round systems

In this section, I compare three outcomes under the single- and the two-round systems: (i) politician's allocations to localities, (ii) politician's choice of the overall budget, and (iii) overall inequality in the allocation of resources. To simplify notation, denote the equilibrium allocations and overall budget for candidate A as q_j^{1R} and G^{1R} (for single-round systems) and q_j^{2R} and G^{2R} (for two-round systems).

Prediction 1.3 states that (i) candidates promise more to locality J in a two-round system than in a single-round system and (ii) the measure of candidates' promises to the other localities in a two-round system compared to a single-round system is ambiguous.

PREDICTION 1.3. *$q_J^{1R} < q_J^{2R}$ and $q_j^{1R} \lessgtr q_j^{2R}$ for all $j \in \{1, \ldots, J-1\}$.*

For locality J – where candidate C has promised the entire government budget – in a single-round system, the fact that voters there will strongly favor candidate C leads candidates A and B to ignore those voters, as the marginal return to allocating resources there is low. In contrast, the two-round system, while not producing a completely equitable distribution, incentivizes candidates to allocate more resources to that locality. The presence of a second round where the third candidate is not present raises the marginal return to allocating resources to that locality. As a result, candidates A and B solicit more votes from locality J in a two-round election.

For the other localities – where candidate C has promised 0 – the change in allocations is ambiguous. If the increase in the overall budget G is small (large) and/or the increase in allocation to locality J is large (small), then allocations to the other localities may decrease (increase). The magnitude of these changes will depend on the parameters of the model, such as the cost of effort, and the functional form of $u_j(\cdot)$.

Prediction 1.4 states that the overall government budget promised is higher in two-round systems than in single-round systems:

PREDICTION 1.4. *$G^{1R} < G^{2R}$.*

In a single-round system, candidates face lower incentives to invest in *all* localities. The marginal return to investing in any locality is higher in two-round elections. These incentives come from the fact that, in two-round elections, there is a conditionality to winning: To win in the first round, candidates must not only attain the most votes, but must attain a majority of votes, otherwise candidates must compete again in a second round. With higher incentives to invest in all localities, candidates in two-round elections exert more effort to increase the government budget.

The proofs for predictions 1.3 and 1.4 are in Appendix A.1.8.

When comparing inequality in allocations between the two electoral rules, the outcome of interest is how the allocation of resources to locality J – where candidate C has offered the entire government budget – compares to the allocation to the other localities.[14] Prediction 1.5 states that when comparing two-round systems to single-round systems, the level of inequality in government resources promised is ambiguous:

Prediction 1.5. *If* $q_j^{1R} > q_j^{2R}$ *for all* $j \in \{1, \ldots, J-1\}$*, then the level of inequality is lower in two-round systems compared to single-round systems. If* $q_j^{1R} < q_j^{2R}$ *for all* $j \in \{1, \ldots, J-1\}$*, then the change in the level of inequality is ambiguous.*

The increased allocation in two-round elections to locality J is a force to reduce inequality in allocations between J and the other localities. However, if candidates also offer more to the other localities in two-round elections, the change in the level of inequality will depend on whether the increase in resources to the other localities is greater or less than the increase in resources to locality J.

Discussion.– While my model partitions the electorate into geographic localities, voters can be partitioned along other dimensions, such as income, race, or ideology. My model easily translates into these other settings, and the predictions yield the same interpretation. First, two-round systems raise the marginal return to allocating resources to groups of voters

[14] I focus on the difference in allocations between locality J and the other localities. However, depending on the functional form of $u_j(\cdot)$, it is possible that relative allocations between the other $\{1, \ldots, J-1\}$ localities become more or less unequal in two-round elections. If differences in $u_j(\cdot)$ across the localities are sufficiently small, any changes in inequality between the $\{1, \ldots, J-1\}$ localities will be small relative to changes in inequality caused by changes in locality J. In the case that $u_j(\cdot) = u(\cdot)$, then all $\{1, \ldots, J-1\}$ localities receive the same allocation, and the only disparity that matters is between J and the other localities.

that are heavily targeted by other candidates. This results in candidates offering campaign promises that also appeal to these voters. Second, candidates face higher incentives to appeal to all voters, leading them to increase the overall government budget.

My model predicts that two-round elections lead to different outcomes because candidates adapt different campaign strategies by offering policy proposals that appeal to a broader group of voters. However, two-round elections may also lead to different outcomes by (i) causing different types of candidates to enter the race or (ii) causing different types of candidates to win the election. While I do not model candidates' entry decisions, I test for these alternative mechanisms in the empirical analysis. I find evidence suggesting that different outcomes in two-round elections are mostly driven by candidates' strategic responses.

1.3 Institutional context

Municipalities in Brazil are autonomous governmental entities with an elected executive (a mayor, or *prefeito*) and legislative body (a council of legislators, or *camara de vereadores*). Elections for municipal positions are at large and held for all positions simultaneously every four years. Mayors in municipalities with less than 200,000 voters are elected through a single-round system, while in larger municipalities they are elected in a two-round system. In the two-round system in Brazil, if no candidate receives at least 50% of the votes in the first round, then a second round is held 3 weeks later with the top two candidates.[15] Legislators are elected through an open-list proportional system. Voters cast votes for a mayoral candidate, and either for a legislative candidate or a generic vote for the party. Mayors are limited to serving two consecutive terms, while there is no term limit for legislators.

State electoral authorities, the *Tribunais Electorais Regionais*, register citizens and maintain electoral rolls. Several features of Brazilian elections, mandated either in the federal constitution or by law, facilitate voter turnout on election day. First, voter registration is

[15] In my sample, a second round occurs 57% of the time.

compulsory and must be completed at least 151 days prior to the election. Second, voting is compulsory for all literate Brazilian citizens between 18 and 69 years of age.[16] Third, elections are held on the first Sunday in October, a day when few voters are at work.

The timing of the announcement of the electoral rule for mayor has varied. In earlier elections, the electoral rule was announced 3–4 months prior to the election. In more recent elections, the number of registered voters has been regularly published, allowing the electoral rule to be known much earlier.

Brazilian elections are a multi-party system, with over 30 political parties registered in the 2016 municipal elections. Mayoral candidates are associated with a party and often a coalition of parties, which are formed prior to the election.[17] Party and coalition affiliations serve as important linkages to the state and federal levels of government (Brollo and Nannicini, 2012).

Once elected, mayors have a broad mandate to provide public goods, particularly in education, health, and local infrastructure. Municipal revenue is a combination of inter-government transfers and local revenues. The majority of the municipal budget is represented by inter-government transfers (the average within my sample is 65.6% of the budget). The bulk of these transfers come from the state or federal level (on average, 78.0%), and are either constitutional automatic transfers (*Fundo de Participação do Municipios*) or discretionary transfers (*convênios*). Municipalities face considerable flexibility in spending these transfers. Among the automatic transfers, 70% of the funds are unrestricted. While 30% are earmarked, municipalities are only restricted to spending this percentage on health and education.

The allocation of spending on public goods occurs through two main channels. First, the majority of public goods are allocated through the annual budgetary process. Second, mayors and legislators can submit bills requesting specific public works or services. While these actions require joint approval by the mayor and legislature, mayors retain veto power

[16] Voters may provide justification to a local electoral judiciary office. Absent this justification, voters who fail to vote must pay a small fine and those who fail to vote for three consecutive elections are prevented from accessing public services, such as obtaining a passport or government loans.

[17] Seats for the legislative council are allocated based on the number of votes received by candidates or parties in the coalition.

and wield significant influence over the process.

As a result, mayors are important for decisions around both the size of the municipal budget and how municipal funds are allocated. This study focuses on public goods provision in municipal education for several reasons. First, a large fraction of the municipal budget is allocated toward education spending: In 2012, it represented 30.5% of municipal budgets. Second, municipal education is a fairly geographically localized public good. Third, unlike other public goods, for which municipalities share joint responsibility with the state or federal government, the provision of elementary education (*Ensino Fundamental*) is one of the few public goods almost entirely under the jurisdiction of the municipality.

1.4 Empirical strategy

1.4.1 Econometric framework

The threshold rule for mayoral elections provides a natural candidate for a regression discontinuity design (RDD). I exploit the 200,000 registered voter threshold, which determines a municipality's electoral rule, to estimate the impact of the electoral rule.

Denote $Y_i(0)$ and $Y_i(1)$ as the potential outcome in municipality i if assigned a single-round and two-round system, respectively. Assignment of the electoral rule is determined by the running variable, the number of registered voters X_i. The assignment variable D_i takes on the value of $D_i = 0$ if $X_i < 200,000$ and $D_i = 1$ if $X_i \geq 200,000$.

Following Imbens and Lemieux (2008) and Calonico et al. (2014), to estimate the treatment effect at the discontinuity, I use a local linear regression specification:

$$Y_{it} = \beta_1 D_{it} + \beta_2 X_{it} + \beta_3 X_{it} \cdot D_{it} + \gamma_t + \varepsilon_{it} \tag{1.4}$$

where for municipality i in election year t, X_{it} is the running variable, D_{it} is the assignment variable, γ_t is an election-year fixed effect, and Y_{it} is the outcome of interest. Each observation in equation (1.4) represents a municipality and election year, or municipality-year. Equation (1.4) amounts to fitting two linear regressions using municipality-years close to the left and to the right of the threshold. β_1 represents the estimate of the local average treatment effect. Standard errors are clustered at the municipality level.

Because the treatment effect is identified only at the threshold, equation (1.4) is estimated using municipality-years close to the threshold. The main analysis uses a 50,000 registered voter window, but robustness is provided for other bandwidths as well as bandwidths selected using data-driven methods (Calonico et al., 2014 and Imbens and Kalyanaraman, 2012).[18]

1.4.2 Identification

In order for β_1 to represent the causal effect of the electoral rule, the conditional expectation of the potential outcomes must be continuous at the threshold. In the following section, I discuss identification and interpretation of the RDD estimates.

Violations of smoothness.– The first way the smoothness assumption can be violated is if the threshold choice is motivated by political or economic factors. There appears to be little evidence for this. The choice of 200,000 registered voters as the threshold was somewhat arbitrary and mainly reflected practical concerns regarding the cost of holding a second round (Chamon et al., 2018; Fujiwara, 2011). In addition, the 200,000 registered voter threshold was set in the 1988 Federal Constitution and has not changed since then. It is unlikely that politicians chose the threshold anticipating which municipalities would be above or below this threshold in 1996 and later.

A second possibility is if municipalities sort across the threshold. Practically, it is difficult for municipalities to selectively sort, since voter registration is mandatory and handled by state electoral authorities. Visually, there are no discontinuities in the distribution of

[18]Because there is a skewed right tail of municipality sizes due to a few extremely large municipalities such as Rio de Janeiro and São Paulo, the data-driven methods would sometimes select bandwidths larger than the support – ie. larger than 200,000. As a result, the optimal bandwidth is calculated on a subset of elections that lies within the support and is symmetrical around the threshold: 0-400,000 voters.

registered voters at the threshold (Figure 1.2). To test this formally, I estimate the size of the discontinuity in the density of the running variable at the threshold (McCrary, 2008). The size of the discontinuity is both small in magnitude (0.169) and insignificant ($p = 0.389$).[19]

Figure 1.2 Density of elections around the 200,000 registered voter threshold

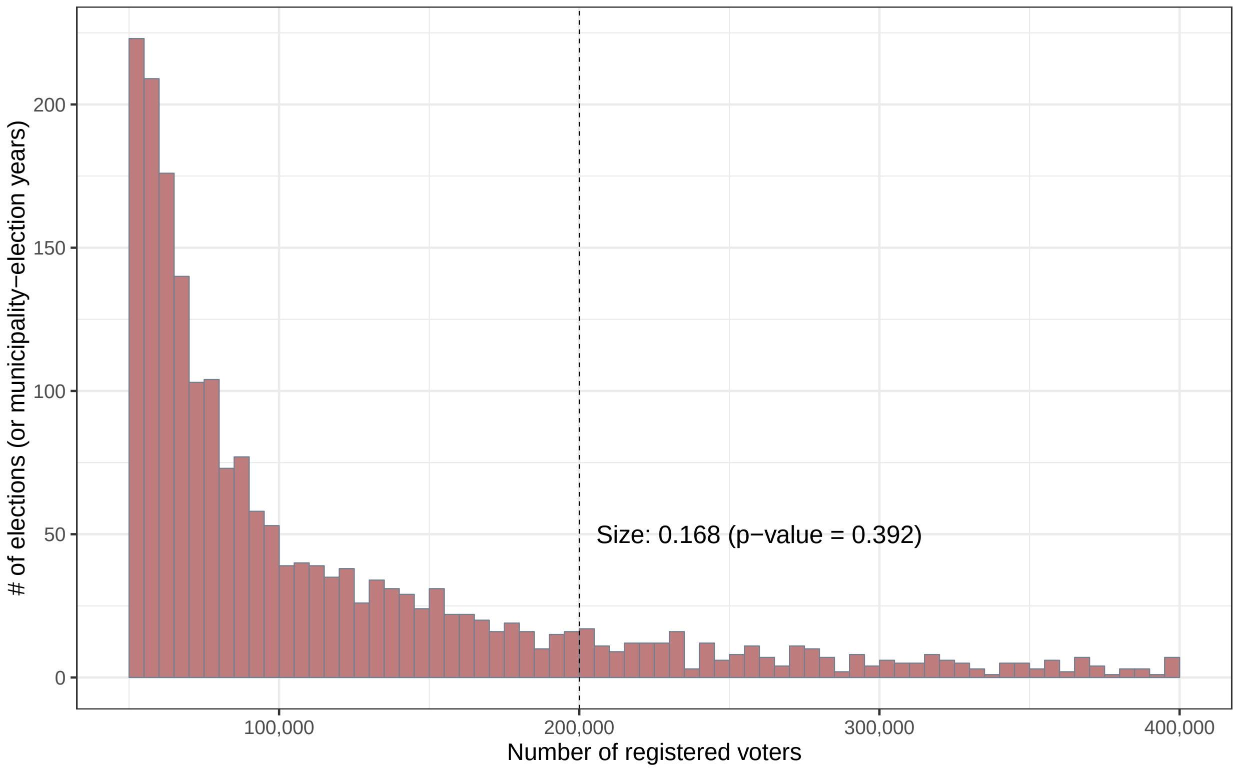

Plot includes only elections with between 50,000 and 400,000 registered voters (6.0% of the universe of elections). Size of the discontinuity in the density of elections uses the McCrary test and is estimated off all elections below the 99.9 percentile in size. An "election" is defined as a municipality-election year. Bin sizes are 10,000 voters.

A third violation is if other policies also change discretely at the threshold. While a number of policies in Brazil are implemented using thresholds, these are based off population counts and not the number of registered voters. Population and the number of registered voters are highly correlated but do not vary one-to-one (Figure A.3). To the best of my knowledge, there are no other policies at the 200,000 registered voter threshold, and most other policy thresholds are far from 200,000 registered voters. Two exceptions are a constitutional

[19]Due to the skewed right tail of municipality sizes, the size of the discontinuity was estimated off a sample of municipality-years excluding those above the 99.9 percentile of registered voters.

amendment in 2000 that places a salary cap for local legislators at 300,000 inhabitants and a constitutional amendment in 2004 that changes the size of the local legislature at 285,714 inhabitants. I address this in two ways. First, I estimate a placebo regression where the electoral rule is assigned at these population thresholds. I show that there are no discontinuities of a similar size in the mayoral electoral outcomes I examine (Tables A.20 and A.21). However, legislator salaries and legislature size can affect economic outcomes, so I do not estimate placebo regressions for my public goods outcomes. Second, I test whether the probability of being above or below these thresholds changes discontinuously at the 200,000 voter threshold. I do not find evidence of a discontinuity (Figure A.4), indicating that any effect of these policies is balanced across the threshold.[20]

The last possibility is that potential confounds change discretely at the threshold. I test this by estimating equation (1.4) on pre-treatment characteristics of municipalities. I discuss this in detail in the following section.

Balance on pre-treatment characteristics.– Since the treatment (here, the two-round election) is determined by the number of registered voters, municipalities can move into the treatment group or be treated multiple times.[21] As a result, I define "pre-treatment" in two ways: 1) prior to the introduction of the electoral system in the 1988 Constitution, and 2) prior to the most recent election in which the municipality was untreated or prior to 1996 if the municipality was never untreated.[22] To measure outcomes for (1), I use the 5% population sample of the 1980 census, the most recent census prior to 1988. To measure outcomes for (2), I use outcomes either from the census prior to the most recent year in a single-round

[20]The regression discontinuity estimate on the probability of being above the 300,000 resident threshold is -0.0350 ($p = 0.735$). The regression discontinuity estimate on the probability of being above the 295,714 resident threshold is 0.0516 ($p = 0.667$).

[21]Municipalities can also move out of treatment. However, while municipalities do experience population decline, none moves below the 200,000 voter threshold.

[22]The earliest electoral data available is 1996, so I cannot observe whether municipalities are treated or untreated prior to 1996. Since there is only one unobserved municipal election after the 1988 Constitution (in 1992) and only 45 municipalities have moved across the 200,000 voter threshold between 1996 and 2016, it is unlikely that many municipalities had their electoral rule change between 1992 and 1996.

system (the 1991, 2000, or 2010 census) or outcomes from the 1991 census, for municipalities that were in a two-round system in 1996.

There is no significant treatment effect on nearly all outcomes measured prior to the 1988 Constitution (Panel A of Table 1.3) or prior to the most recent election in a single-round system (Panel B of Table 1.3). The treatment effect is estimated on economic characteristics (unemployment rate, literacy, and income), segregation levels[23] (along income and demographics), and income inequality. To rule out the concern that there are factors that change discontinuously at the threshold and affect which municipalities move into treatment and the length of treatment, I test whether population growth is discontinuous at the threshold. If municipalities that grow at slower rates are more likely to be under the threshold and remain there longer, it would pose an issue for causal identification. However, I do not find that population growth is discontinuous at the threshold.

One exception is a large and significant effect on population density. However, there are several reasons to believe that this is a false positive. First, dropping one outlier municipality reduces the coefficient by 38.2%, suggesting that a few municipalities are driving this effect (although the estimate is still significant at the 10% level). My main results are robust to dropping this outlier municipality. In addition, there is no visible discontinuity and the effect disappears at larger bandwidths (Figures A.5 and A.6). Second, the estimate ($p = 0.047$) is not significant after Bonferroni adjusting the significance threshold for the number of hypotheses tested. Third, the regression discontinuity coefficients across the ten outcomes (from the most recent single-round election) are not jointly significant ($p = 0.339$).

While it is not clear how an imbalance in population density impacts the economic and political outcomes of interest, this poses an issue if politicians are manipulating the composition and size of the electorate by, say, moving citizens or manipulating municipality borders. This does not appear to be the case for two reasons. First, I find no differences

[23] This is calculated using the entropy index (see Section 1.4.4 for a detailed calculation), which measures how far each census sector is from equal representation of all groups. This was calculated separately for income and demographics. For income, the groups are defined by bins of income relative to the minimum wage. For demographics, the groups are defined by sex, age, and literacy.

Table 1.3 Regression discontinuity estimates on municipality pre-characteristics

Panel A: Characteristics measured prior to the 1988 Constitution					
	% illiterate	% low income	Unempl. rate	Pop. density	
TwoRound	−1.499	−0.751	0.044	-397.678^{**}	
	(1.057)	(1.148)	(0.145)	(184.347)	
Single-round mean	19.029	50.182	1.978	516.067	
Observations	293	293	293	293	
Panel B: Characteristics measured prior to most recent single-round election					
	Muni. area change (%)	Pop. growth (%)	Pop. density	Income seg.	Dem. seg.
TwoRound	−0.601	−0.468	-663.035^{**}	−0.003	0.001
	(2.296)	(1.260)	(332.555)	(0.004)	(0.001)
Single-round mean	-3.003	2.863	1,013.396	0.090	0.028
Observations	231	295	295	231	231
	% illiterate	Income per capita	% low income	Unempl. rate	Gini coefficient
TwoRound	0.014	−18.964	1.211	−0.652	0.004
	(0.590)	(33.328)	(2.290)	(1.114)	(0.009)
Single-round mean	7.409	646.917	36.123	11.671	0.540
Observations	295	295	295	295	295

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Panel A: outcomes from the 1980 census. *Panel B:* outcomes either from the census prior to the most recent election in a single-round system or from the 1991 census. F-stat for all treatment effects in Panel B jointly significant: 1.146 ($p = 0.339$). *Muni. area change* is the percentage change in municipality area from the prior census. *Pop. growth* is the percentage change in population from the prior census. *Pop. density* is population density, per km^2. *Income seg.* and *Dem. seg.* refer to income and demographic segregation, respectively, of census tracts (measured using the entropy index). *Income per capita* is average monthly household income per capita, in *reais*. *% low income* is the fraction of households earning between 0 and 50% of the minimum wage. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Standard errors clustered at the municipality level. *Source:* 1980, 1991, 2000, and 2010 Demographic Census.

in the urbanization rate (a difference of 0.065%, where the single-round mean is 95.191%, $p = 0.943$). Second, this effect is seen in 1980, prior to the introduction of the threshold rule (Panel A of Table 1.3). Third, there is no evidence that changes in municipality area or population growth change discontinuously across the threshold (Panel B of Table 1.3). Nevertheless, I include population density as a control in all specifications:

$$Y_{it} = \beta_1 D_{it} + \beta_2 X_{it} + \beta_3 X_{it} \cdot D_{it} + \beta_4 Z_{it} + \beta_5 Z_{it} \cdot D_{it} + \gamma_t + \varepsilon_{it}$$

where Z_{it} is the municipality's population density in the most recent census prior to the election.

Testing smoothness .– I test whether my outcomes of interest vary smoothly around the discontinuity. To do this, I estimate placebo regressions where the electoral rule is assigned at other registered voter thresholds between 170,000 and 230,000 voters. This tests whether the effects I estimate are concentrated at the actual threshold of 200,000 voters. I show that there is no discontinuity in my outcomes at thresholds where the treatment does not change (Figures A.22, A.23, A.24, A.25, and A.26).

Registered voters as a running variable.– Since treatment (moving into a two-round election) is determined by the number of registered voters, there are interesting implications for the interpretation of the regression discontinuity estimate.

As discussed above, municipalities can move into treatment or be treated multiple times. While this does not invalidate causal identification, it affects whether the treatment effects I estimate are the result of single or multiple treatments.

To investigate this, I test whether the probability that a municipality's *previous* election was a two-round system changes discontinuously at the threshold (Figure A.8). Because the regression discontinuity framework identifies the treatment effect *at* the threshold, treatment effects should be interpreted as the result of the change in this probability *at* the threshold. I find that at the 200,000 registered voter threshold, the probability of having previously been treated is zero and there is no discontinuous jump.[24] As a result, the treatment effects identify the effect of moving into a two-round election for the first time.

Compliance.– While not an issue for causal identification, imperfect compliance with treatment can affect the interpretation of the causal estimates (Angrist et al., 1996). In this context, compliance was perfect (Figure A.7). All municipalities below the threshold or where the top candidate received at least 50% held one round. All municipalities above the threshold and where the top candidate did not receive at least 50% held two rounds.

1.4.3 Data sources

Electoral data.– Data on municipality elections come from Brazil's electoral authority (*Tribunal Superior Eleitoral*, or TSE). These are available for 6 municipal elections between 1996 and 2016. The electoral data provides information on the candidates running, the party and coalition each candidate belongs to, and the number of votes received. In total, the data encompasses 32,767 municipal elections covering 5,568 municipalities.

Electoral results are available for each polling station (*seçao eleitoral*), allowing me to observe at a very fine level the number of votes each candidate receives. I use this to measure the geographic distribution of voters for specific candidates at both an *overall* and *candidate* level (see Section 1.4.4 on the measures used). Baseline results use votes from the first round of elections. Results from the first round are used in order to have similar measures between single- and two-round elections, but results are robust to using votes from the final round (the first round in single-round elections and the second round in two-round elections).

Public goods provision in schools.– To measure public goods provision in municipal public schools, I use data provided in the 1997-2016 School Census (*Censo Escolar*). This census is conducted annually by the research arm of the Ministry of Education, *Instituto Nacional de*

[24] Because no municipalities move under the threshold, all municipalities under the threshold have 0 probability that the previous election was two-round. The regression discontinuity estimate on the probability of being in a two-round system in the previous election is -0.00802 ($p = 0.870$).

Estudos e Pesquisas Educacionais.

I use the Census to observe the level of resources present in schools offering elementary education. I calculate a measure of resources present in each school across two categories: equipment and infrastructure. Equipment includes *movable* elements, such as the number of computers and availability of air conditioning (see Table A.1 for the full list). Infrastructure includes *immovable* elements, such as the number of classrooms, type of sanitation, and availability of a library (see Table A.2 for the full list). I construct indices of these resources, separately for equipment and infrastructure, by taking the principal component of these variables and computing the school's percentile rank within the country for each year.[25]

1.4.4 Measuring concentration of voters

In this section, I use the following notation. In municipality m, there are K_m candidates and I_m polling stations. p_{mk} is the fraction of voters for candidate k in municipality m, p_{imk} is the fraction of voters for candidate k in polling station i, n_{im} is the number of voters in polling station i, and N_m is the number of voters in municipality m.

Overall geographic concentration of voters.– To measure the *overall* geographic concentration of voters, I use three indices that are commonly used in the racial segregation literature to measure multigroup spatial segregation: the coefficient of variation, the fractionalization index, and the entropy index. These indices and their properties are described in more detail in White (1986) and Reardon and O'Sullivan (2004). The indices assume a value of 1 if there is full geographic concentration of voters: Each polling station contains voters for only one candidate. The indices assume a value of 0 if there is full geographic dispersion of voters: Each polling station contains the same composition of voters as the municipality as a whole.

[25] The variables collected in the School Census under each category varied from year to year. This makes it difficult to compare the raw PCA index across years. Calculating a school's percentile rank across all schools for that year allows for valid comparisons across years.

The coefficient of variation, s_m, is defined as:

$$s_m = \frac{1}{K_m - 1} \sum_{k=1}^{K_m} \sum_{i=1}^{I_m} \frac{n_{im}}{N_m} \frac{\left(p_{imk} - p_{mk}\right)^2}{p_{mk}}$$

This index is interpreted as the square deviation of voter composition in polling stations from voter composition in the municipality. Dividing by $K_m - 1$ keeps the index between 0 and 1. When each polling station has the same composition as the municipality ($p_{imk} = p_{mk}$), the index takes a value of 0.

The fractionalization index, f_m, is calculated in two steps. The average fractionalization within each polling station, $\bar{f}_m$, is defined as:

$$\bar{f}_m = \sum_{k=1}^{K_m} \sum_{i=1}^{I_m} \frac{n_{im}}{N_m} p_{imk} \left(1 - p_{imk}\right)$$

The overall fractionalization within the municipality, $\hat{f}_m$, is defined as:

$$\hat{f}_m = \sum_{k=1}^{K_m} p_{mk} \left(1 - p_{mk}\right)$$

The final measure, the fractionalization index, f_m, is defined as:

$$f_m = \frac{\hat{f}_m - \bar{f}_m}{\hat{f}_m}$$

Fractionalization, also known as the interaction index, measures the probability that two members within a population chosen at random are from different groups. There are two ways to interpret f_m. One, f_m is the average concentration within polling stations, normalized by the level in the municipality to keep the index between 0 and 1. Two, f_m is the fraction of concentration in the municipality that is due to differences in voter composition between polling stations. When each polling station has the same concentration as the municipality – or when there are no differences between polling stations and $\hat{f}_m = \bar{f}_m$ – the index takes a value of 0. When each polling station contains only one type of voter – or when there are large differences between polling stations and $\bar{f}_m = 0$ – the index takes a value of 1.

The entropy index, h_m, is also calculated in two steps. The average entropy within each polling station, $\bar{h}_m$, is defined as:

$$\bar{h}_m = -\sum_{k=1}^{K_m}\sum_{i=1}^{I_m}\frac{n_{im}}{N_m}p_{imk}\ln p_{imk}$$

The overall entropy within the municipality, $\hat{h}_m$, is defined as:

$$\hat{h}_m = -\sum_{k=1}^{K_m} p_{mk}\ln p_{mk}$$

The final measure, the entropy index, h_m, is defined as:

$$h_m = \frac{\hat{h}_m - \bar{h}_m}{\hat{h}_m}$$

Entropy measures how far the population is from equal representation of all groups. The interpretation and range of values of the entropy index h_m are the same as that of the fractionalization index f_m.

In my sample, the correlation between these indices is between 0.89 and 0.96.[26] Conceptually, the three indices can be thought of as measures of the average deviation of the composition of polling stations from that of the municipality.

Sensitivity of overall concentration to the number of candidates.– The indices may mechanically change due to the number of candidates. I discuss this briefly below; the majority of this discussion derives from Reardon and O'Sullivan (2004).

For the coefficient of variation, before dividing by $K_m - 1$, the index attains a maximum value of $K_m - 1$ and as a result would depend mechanically on the number of candidates.

[26] The correlation coefficient between the coefficient of variation and fractionalization is 0.956; between the coefficient of variation and entropy is 0.892; and between fractionalization and entropy is 0.929.

Dividing by $K_m - 1$ nets out this mechanical effect.

For the fractionalization index, $\sum_{i=1}^{I_m} p_{imk}(1 - p_{imk})$ ranges from $1/K_m$ to 1. As a result, the polling station average, $\bar{f}_m$, and municipality measure, $\hat{f}_m$, depend mechanically on the number of candidates. Normalizing $\bar{f}_m$ by $\hat{f}_m$ removes part of this mechanical effect.

For the entropy index, $\sum_{i=1}^{I_m} p_{imk} \ln p_{imk}$ ranges from 0 to $\ln K_m$. As a result, the polling station average, $\bar{h}_m$, and municipality measure, $\hat{h}_m$, depend mechanically on the number of candidates. Normalizing $\bar{h}_m$ by $\hat{h}_m$ removes part of this mechanical effect.

While these indices no longer monotonically depend on the number of candidates, these indices may still be affected by the number of candidates. I address this concern with several robustness exercises in Section 1.5.4.

Candidate-level geographic concentration of voters.– To capture a *candidate*-level measure of the spatial distribution of voters, I use the standard deviation in vote shares across polling stations. This measure describes whether a candidate's supporters are spread across many or concentrated within a few polling stations. For candidates who receive votes from many areas in the municipality, we should expect to observe the vote share varying less across polling stations. The standard deviation in a candidate's vote share, σ_{mk}, is defined as:

$$\sigma_{mk} = \left(\frac{1}{I_m - 1} \sum_{i=1}^{I_m} \left(p_{imk} - \frac{1}{I_m} \sum_{i=1}^{I_m} p_{imk} \right)^2 \right)^{1/2}$$

Using only the top two candidates.– In some cases, I calculate the indices using only vote shares from the top two candidates. To do this, I assume that only the top two candidates are in the race and use each candidate's vote share out of the top two. These indices should be interpreted as the geographic distribution of voters for the top two candidates. The advantage of doing this is that it allows me to keep keep the number of candidates fixed and ignore the potential dilution of votes from lower-placed candidates. This is done when comparing to the second round of two-round elections or to show robustness of my results to the number of candidates. Note that, by construction, the standard deviation of votes for the 1st place candidate σ_{m1} is the same as that of the 2nd place candidate σ_{m2}.

1.5 The effect of the two-round system

I present three main results. First, candidates in two-round elections receive broader geographical support. Second, once in office, politicians elected under two-round systems provide more resources to schools and distribute these resources more equitably. These results suggest that politicians in two-round systems are represented by a broader group of voters and that this affects how public goods are provided in the municipality. Third, I find that educational outcomes are improved in two-round municipalities.

1.5.1 The geography of votes

Do candidates in two-round elections secure broader bases of support? I provide evidence that voters are *overall* less geographically concentrated, and that it is the *top two* candidates who receive support from a geographically broader group of voters.

Geographic concentration of voters.– The coefficients on the coefficient of variation, fractionalization index, and entropy index are negative, indicating that voters for specific candidates are overall less geographically concentrated in two-round elections (Figure 1.4 and Panel A of Table 1.5). The composition of voters in polling stations is closer to the composition of voters in the municipality as a whole, indicated by the coefficient of variation, and the composition of voters in polling stations is on average less concentrated, indicated by the fractionalization and entropy indices. Municipalities in two-round systems experience a reduction of 0.0087, 0.0118, and 0.0082 in the coefficient of variation, fractionalization, and entropy of voters, respectively. These coefficients correspond to 45.6%, 43.9%, and 27.4%, respectively, of the average level in single-round municipalities within the bandwidth.

Next, I test whether all or only some candidates obtain support from geographically broader constituencies. I find that voters for the top two candidates are less concentrated in

Figure 1.4 Regression discontinuity plots of overall concentration of voters for specific candidates

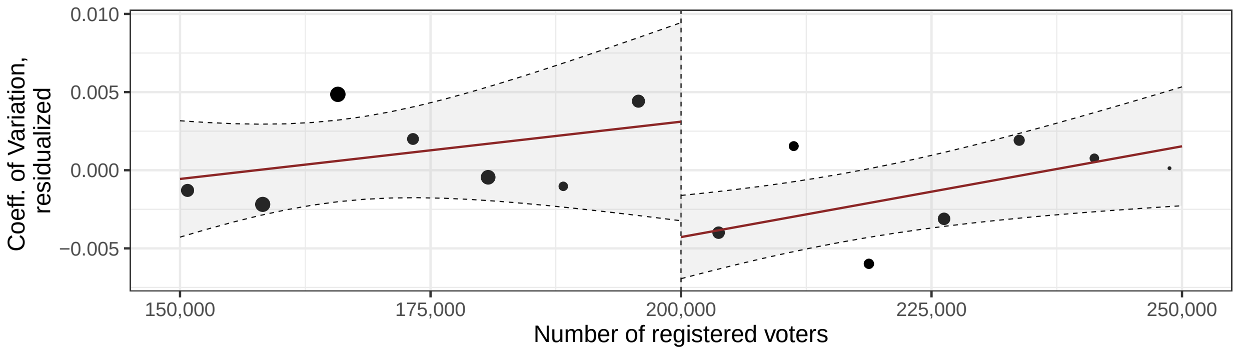

(a) Coefficient of variation

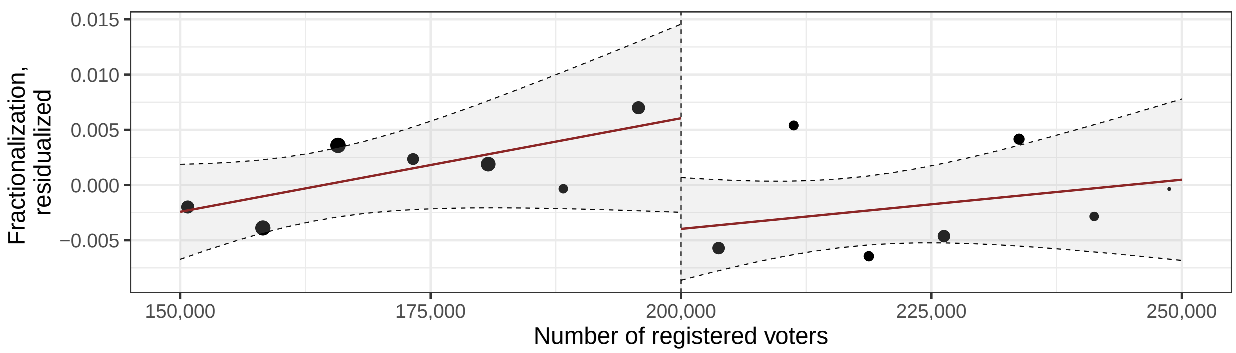

(b) Fractionalization

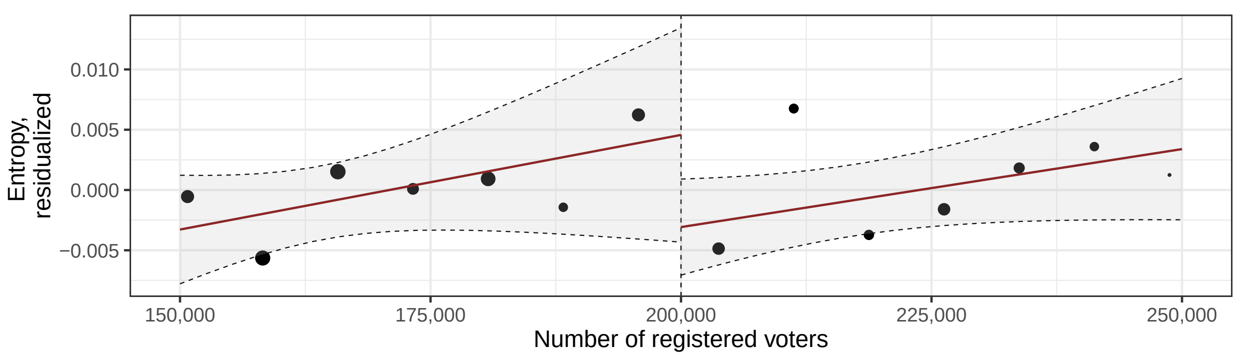

(c) Entropy

Overall concentration of voters for specific candidates, as measured by (a) *Coefficient of variation*, (b) *Fractionalization*, and (c) *Entropy*, using vote counts in polling stations. Vote shares are from the first round. In each panel, each point plots an average value within a 7,500 voter bin. Variables on the vertical axis are residualized by population density and election-year fixed effects. Diameter of the points is proportional to the number of observations. Confidence intervals (dashed lines) represent the 95% confidence intervals of a local linear regression (solid red line) with standard errors clustered at the municipality level.

two-round systems, but that voters for the third and fourth placed candidates in two-round systems are concentrated similarly to voters of the same-placed candidates in single-round

Table 1.5 Regression discontinuity estimates on the geographic concentration of voters

Panel A: Concentration indices of voters for specific candidates				
	Coefficient of variation	Fractionalization	Entropy	
TwoRound	-0.009^{***}	-0.012^{**}	-0.008^{*}	
	(0.003)	(0.005)	(0.005)	
Potential bias	0.0008	–0.0002	–0.0923	
Single-round mean	0.019	0.027	0.030	
Observations	264	264	264	
Panel B: Standard deviation in vote shares for each candidate				
	1st place candidate	2nd place candidate	3rd place candidate	4th place candidate
TwoRound	-0.017^{**}	-0.014^{*}	-0.005	0.004
	(0.007)	(0.008)	(0.007)	(0.004)
Potential bias	–0.0011	–0.0010	–0.0004	–0.0002
Single-round mean	0.080	0.075	0.042	0.023
Observations	264	264	251	216

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Panel A: overall concentration of voters for specific candidates, as measured by coefficient of variation, fractionalization, and entropy of vote counts in polling stations. *Panel B:* candidate-level concentration of voters, measured by standard deviation in a candidate's vote shares (for the 1st-4th place candidate) across polling stations. *Potential bias* is the simulated effect on the outcome from having an additional candidate in every single-round election. Vote shares are from the first round. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

systems (Figure 1.6 and Panel B of Table 1.5). The estimates for the top two candidates are similar in magnitude: The first placed candidate experiences a 0.0167 reduction in variance (20.9% of the single-round mean), and the second placed candidate experiences a 0.0142 reduction (18.9% of the single-round mean). The estimates for the third and fourth placed candidates are close to zero and insignificant.

The effects on overall concentration are significantly stronger when using vote shares from the top two candidates only (Table 1.7).[27] This further indicates that the bases of support of

[27] I compare the coefficients in Table 1.5 with Table 1.7. The p-values for the difference between the estimates are 0.024 (for the coefficient of variation), 0.052 (for fractionalization), and 0.147 (for entropy).

Figure 1.6 Regression discontinuity plots of the candidate-level concentration in voters

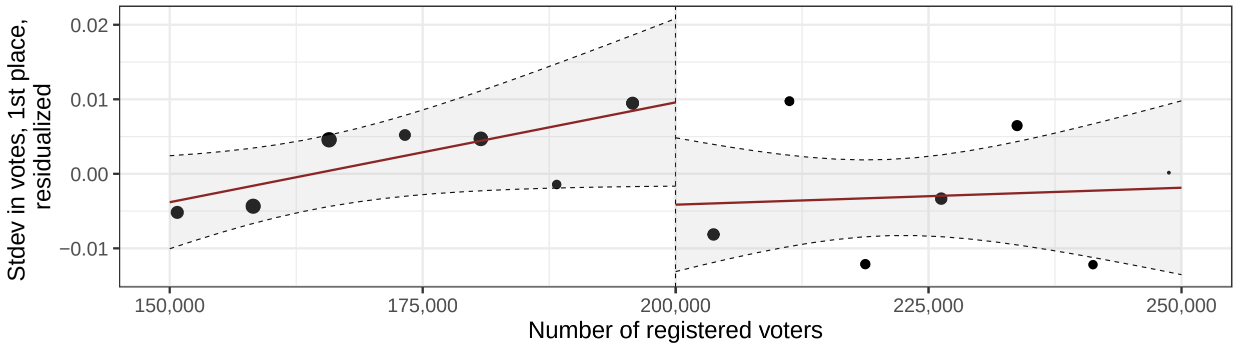

(a) Standard deviation in votes for the 1st place candidate

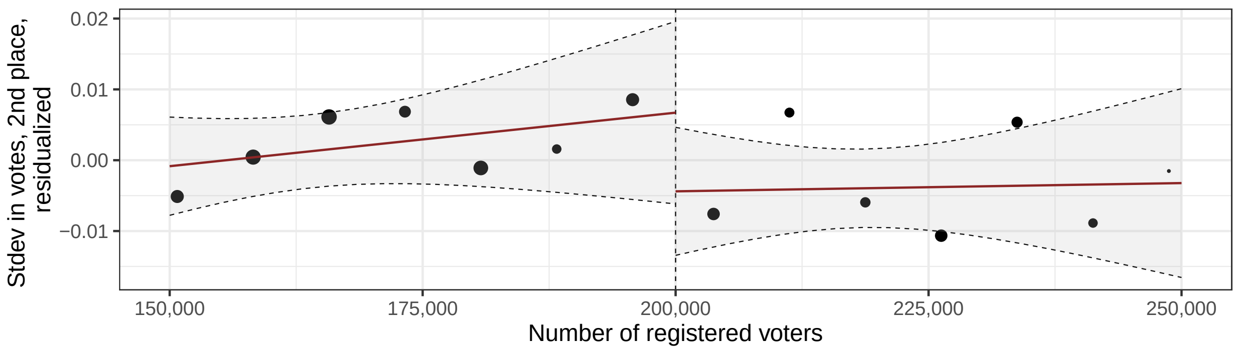

(b) Standard deviation in votes for the 2nd place candidate

Standard deviation in a candidate's vote counts across polling stations, for the (a) 1st place and (b) 2nd place candidate. Vote shares are from the first round. In each panel, each point plots an average value within a 7,500 voter bin. Variables on the vertical axis are residualized by population density and election-year fixed effects. Diameter of the points is proportional to the number of observations. Confidence intervals (dashed lines) represent the 95% confidence intervals of a local linear regression (solid red line) with standard errors clustered at the municipality level.

the top two candidates drive the reduced concentration of voters in two-round elections.

These effects are not limited to vote shares in the first round. Using only vote shares from the top two candidates, I compare concentration in single-round elections with concentration in the second round of two-round elections. Voters in the second round are also less geographically concentrated (Panel A of Table 1.8). Additionally, I compare concentration between the first and second round in two-round municipalities. While these estimates are correlational and not causal, concentration in the second round is lower than in the first round, indicating that candidates consolidate their voter bases between rounds (Panel B of Table 1.8).

Table 1.7 Regression discontinuity estimates on the geographic concentration of voters, using vote shares from the top two candidates only

	Coefficient of variation	Fractionalization	Entropy	Std Dev of 1st place candidate
TwoRound	-0.015^{*}	-0.018^{**}	-0.013^{**}	-0.015
	(0.008)	(0.009)	(0.007)	(0.009)
Single-round mean	0.036	0.038	0.029	0.088
Observations	264	264	264	264

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Coefficient of variation, *Fractionalization* and *Entropy* measure the overall concentration of voters for specific candidates, using vote counts in polling stations. *Standard deviation of 1st place candidate* is the standard deviation in the 1st place candidate's vote counts across polling stations. All outcomes use only vote shares from the top two candidates. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Voter engagement.– These results suggest that two-round elections lead to greater inclusiveness, as voters from more geographical areas are represented. I find inclusiveness along another dimension: voter behavior. Specifically, I find higher rates of voter engagement in two-round elections. While turnout is unaffected (which is expected, as turnout is mandatory in Brazil), the number of blank and invalid ballots is significantly lower in two-round municipalities (Table 1.9). The number of blank and invalid ballots plausibly corresponds to the number of dissatisfied or disinterested voters (Gonzales et al., 2019).[28] The reduction suggests that voters in two-round elections engage in the electoral process at higher rates.

1.5.2 The allocation of municipal resources

I next investigate the impact of the two-round system on public goods provision. If politicians secure broader bases of support in two-round elections, they may also provide

[28]Ballots can be invalid or blank for a number of reasons. For example, municipalities with higher numbers of illiterate voters will have more blank and invalid ballots (Fujiwara, 2015). Since the illiteracy rate is not discontinuous across the threshold and all municipalities used electronic voting by 2000 (which reduced the number of unintentional errors), I interpret the difference in the number of blank and invalid ballots as voter engagement. Gonzales et al. (2019) provide empirical evidence for this interpretation, as they find that forced electoral participation increases the number of blank and invalid ballots cast.

Table 1.8 Estimates on the geographic concentration of voters across different rounds of elections

	Coefficient of variation	Fractionalization	Entropy	Std Dev of 1st place candidate
Panel A: single- versus two-round elections: vote shares from final round (1st round in single-round compared to 2nd round in two-round)				
TwoRound	-0.020^{**}	-0.023^{***}	-0.017^{***}	-0.019^{**}
	(0.008)	(0.008)	(0.006)	(0.009)
Single-round mean	0.036	0.038	0.029	0.088
Observations	216	216	216	216
Panel B: two-round elections: vote shares in first versus second round (1st round in two-round compared to 2nd round in two-round)				
2ndRound	-0.012^{***}	-0.012^{***}	-0.010^{***}	-0.013^{***}
	(0.002)	(0.002)	(0.002)	(0.002)
First round mean	0.049	0.049	0.037	0.103
Observations	432	432	432	432

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Coefficient of variation, *Fractionalization* and *Entropy* measure the overall concentration of voters for specific candidates, using vote counts in polling stations. *Standard deviation of 1st place candidate* is the standard deviation in the 1st place candidate's vote counts across polling stations. All outcomes use only vote shares from the top two candidates. *Panel A* compares the 1st round results (in single-round elections) with the 2nd round results (in two-round elections) and presents the regression discontinuity estimates. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level. *Panel B* compares the 1st round results (in two-round elections) with the 2nd round results (in two-round elections), using the full sample of elections that held two rounds. *Estimation method:* Standard regression with *2ndround* as the regressor and election fixed effects. Standard errors clustered at the municipality level.

public goods differently once in office. I provide evidence that two-round elections impact both the *level* and *distribution* of municipal resources.

Municipal schools.– Municipal schools in two-round municipalities have, on average, more resources than those in single-round municipalities (Figure 1.10 and Columns 1 and 2 in Table 1.12). A school in a two-round municipality is, on average, 0.081 percentiles and 0.057 percentiles higher in the national distribution of equipment resources and infrastructure resources, respectively. The coefficient on infrastructure resources is smaller and less significant; this may be because infrastructure is more difficult to manipulate. Allocating new infrastructure,

Table 1.9 Regression discontinuity estimates on other electoral outcomes

	Turnout	Blank/invalid ballots	# candidates
TwoRound	0.006	-3.821^{**}	1.273^{***}
	(0.008)	(1.670)	(0.339)
Single-round mean	0.843	16.524	4.604
Observations	296	296	296

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Turnout is the fraction of eligible voters who cast a ballot in the election. *Blank/invalid ballots* is the sum of ballots (in thousands) that were either blank or voided, and is also equal to turnout minus valid ballots. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

such as gymnasiums, requires significantly more time and capital than allocating equipment, such as computers. This is reflected when looking at the estimates separately by year in the mayoral term: Infrastructure levels are less responsive to the electoral cycle (Column 1 in Table A.9).

In addition to differences in the overall levels, there is less variance in the resources present in schools in two-round municipalities (Figure 1.11 and Columns 1 and 2 in Table 1.12). The standard deviation in equipment resources is 0.018 percentiles lower in two-round municipalities (15.9% of the single-round mean). Although the estimate on the standard deviation in infrastructure resources is of a similar magnitude (-0.021 percentiles), the difference is not significant.

It is instructive to estimate effects for schools at different parts of the distribution in the municipality. Intuitively, if the variance in resources is lower in two-round municipalities, we should expect to see that schools with the least (most) resources in the municipality have more (less) resources.

I group schools into quartiles, which are defined by first calculating each school's percentile in the municipality distribution prior to the election, and then assigning the school to one of

Figure 1.10 Regression discontinuity plots of the overall level of resources in municipal schools

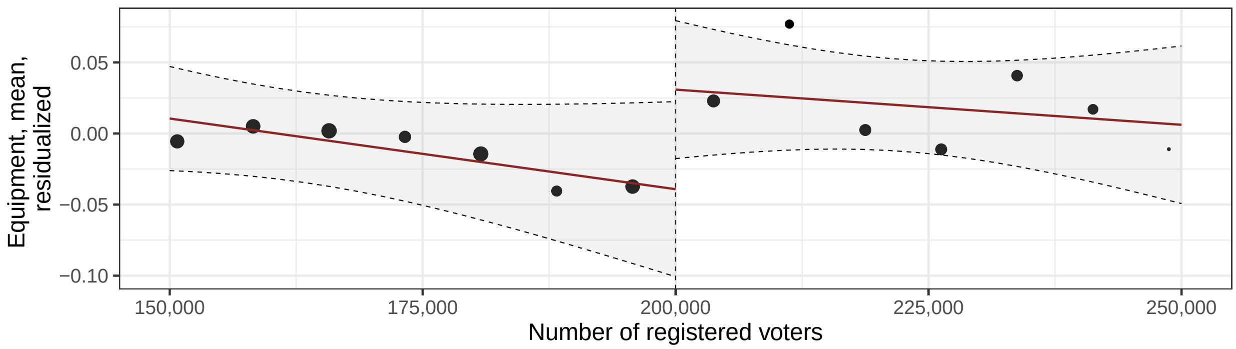

(a) Equipment, mean level of resources

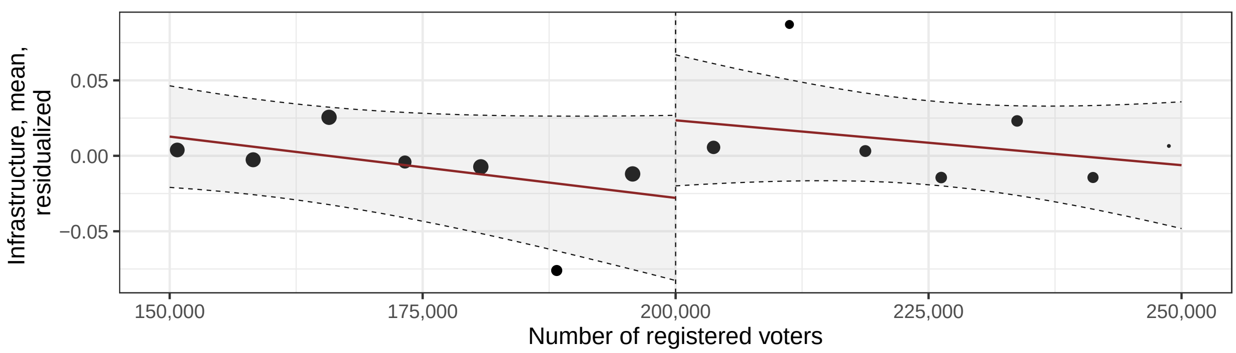

(b) Infrastructure, mean level of resources

Equipment and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. *Mean level of resources* is the mean index level across schools in the municipality for (a) equipment and (b) infrastructure. In each panel, each point plots an average value within a 7,500 voter bin. Variables on the vertical axis are residualized by population density and election-year fixed effects. Diameter of the points is proportional to the number of observations. Confidence intervals (dashed lines) represent the 95% confidence intervals of a local linear regression (solid red line) with standard errors clustered at the municipality level.

four quartiles. The increased level of resources is concentrated in schools located at the lower end of the distribution (Table 1.13).[29] Schools at the bottom 25% of the distribution are 0.082 percentiles higher in equipment resources and 0.116 percentiles higher in infrastructure resources. Schools in the second quartile experience positive, but smaller, gains – 0.066 percentiles in equipment resources and 0.102 percentiles in infrastructure resources. There is no significant difference in resources at the top of the distribution.

[29]The p-values for the difference in the regression discontinuity estimates between schools at the 1st quartile compared to the 4th quartile is 0.057 (for equipment) and 0.038 (for infrastructure).

Figure 1.11 Regression discontinuity plots of the distribution of resources in municipal schools

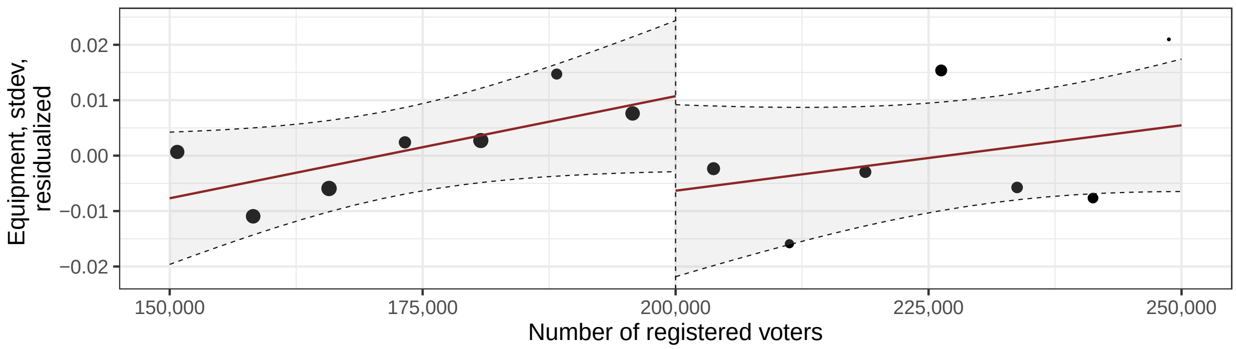

(a) Equipment, standard deviation in resources

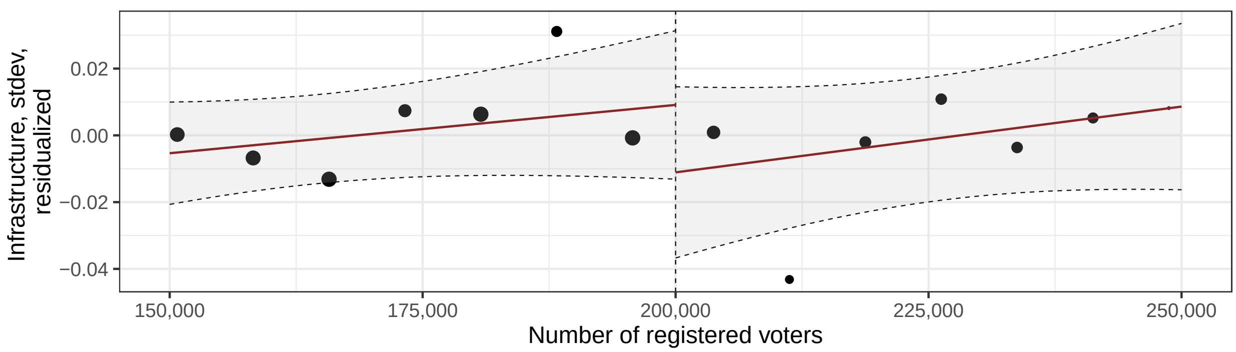

(b) Infrastructure, standard deviation in resources

Equipment and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. *Standard deviation in resources* is the standard deviation in the index across schools in the municipality for (a) equipment and (b) infrastructure. In each panel, each point plots an average value within a 7,500 voter bin. Variables on the vertical axis are residualized by population density and election-year fixed effects. Diameter of the points is proportional to the number of observations. Confidence intervals (dashed lines) represent the 95% confidence intervals of a local linear regression (solid red line) with standard errors clustered at the municipality level.

1.5.3 Downstream outcomes

If mayors provide more public goods and distribute them more equitably, does this translate into downstream economic outcomes? I find, in two-round municipalities, improvements in education outcomes but limited effects on economic outcomes.

Education outcomes.– I measure education outcomes in four ways. Using the School Census, I measure the drop-out, failing, and passing rate in schools across the municipality. Using the

Table 1.12 Regression discontinuity estimates on resources in municipal schools

	Mean level of resources		*Standard deviation in resources*	
	Equipment	Infrastructure	Equipment	Infrastructure
TwoRound	0.081^{**}	0.057^{*}	-0.018^{*}	-0.021
	(0.035)	(0.033)	(0.009)	(0.016)
Single-round mean	0.738	0.731	0.121	0.157
Observations	820	912	820	912

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Equipment and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. The first two columns (*Mean level of resources*) have as the dependent variable the mean index level across schools in the municipality. The last two columns (*Standard deviation in resources*) have as the dependent variable the standard deviation in the index across schools in the municipality. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

2000 and 2010 Demographic Censuses, I measure the literacy rate among cohorts who were of elementary school age during the electoral term. I find that in two-round municipalities, drop-out rates are significantly lower, 1.65 percentage points off a baseline of 3.21 percentage points; literacy rates are also significantly higher among elementary cohorts, 1.20 percentage points off a baseline of 91.45 percentage points (Figure 1.14 and Panel A of Table 1.15). While the drop-out rates are lower, this does not necessarily lead to improved passing rates or worsened failing rates. These results suggest that differential public goods provision in two-round municipalities lead to some improvements in education outcomes.

Economic outcomes.– If two-round municipalities lead to improved education outcomes, this may result in improved economic outcomes. I find some – though limited– improvements on broader economic outcomes (Panel B of Table 1.15). Using the 2000 and 2010 Demographic Census, I measure the fraction of low-income households, income per capita, and the unemployment rate. These outcomes are measured between 2 and 10 years after the election.[30] Using the 1997-2013 NOAA night lights series, I measure the mean night lights

[30]For the 1996 elections, outcomes are observed 4 years later in the 2000 Demographic Census. For the 2000, 2004, and 2008 elections, outcomes are observed 10, 6, and 2 years later in the 2010 Demographic

Table 1.13 Regression discontinuity estimates on resources in municipal schools, for schools at different quartiles in the municipal distribution

	Mean level of resources in schools at different quartiles			
	1st quartile (Bottom 25%)	2nd quartile	3rd quartile	4th quartile (Top 25%)
Panel A: Equipment				
TwoRound	0.082^{**}	0.066^{*}	0.069^{*}	0.038
	(0.035)	(0.037)	(0.042)	(0.029)
Single-round mean	0.652	0.733	0.781	0.856
Observations	700	728	760	748
Panel B: Infrastructure				
TwoRound	0.116^{**}	0.102^{**}	0.056	0.013
	(0.046)	(0.047)	(0.035)	(0.021)
Single-round mean	0.540	0.689	0.814	0.914
Observations	776	764	784	780

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Equipment and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. Dependent variables are the mean index level of equipment (*Panel A*) and infrastructure (*Panel B*) elements, separately by quartiles. Quartiles are defined by the school's percentile in the municipal distribution in the year prior to the election. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

level in the municipality. I find that the fraction of low-income households is significantly lower. This may be a false positive, as the coefficient is large relative to the mean; controlling for the pre-election level reduces the coefficient to -1.477 percentage points ($p = 0.0872$). Coefficients on the other outcomes suggest that two-round municipalities lead to improved economic outcomes (income per capita is higher, unemployment is lower, and night lights is higher), but that this effect is not significant.

While two-round elections may simply not lead to improved economic outcomes, there may be several reasons that this is a false negative. First, the public goods and policies

Census.

Figure 1.14 Regression discontinuity plots of municipal education outcomes

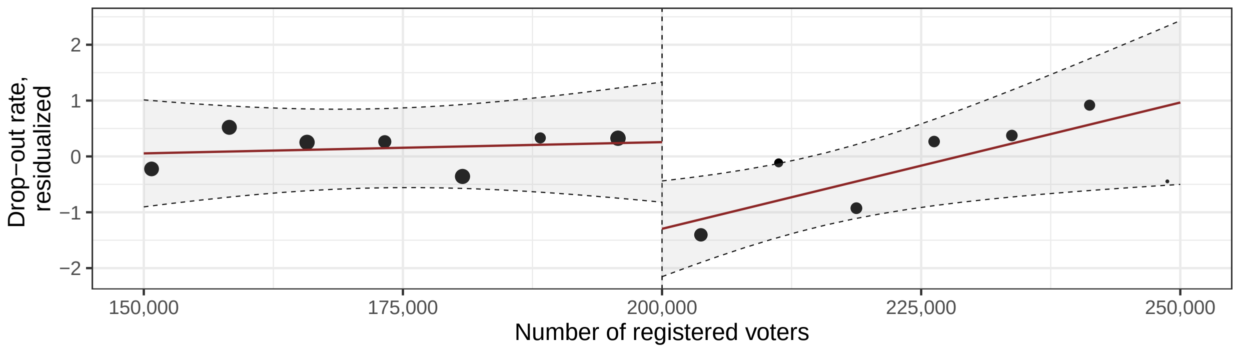

(a) Drop-out rate

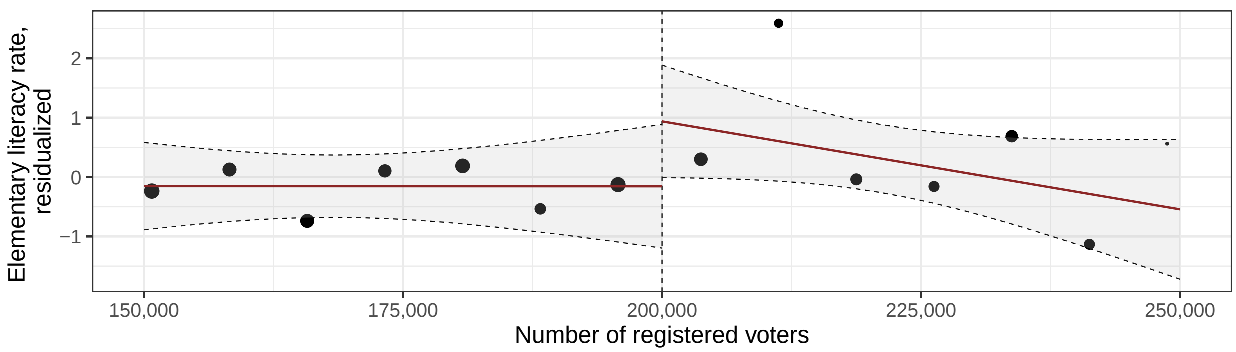

(b) Elementary literacy rate

Drop-out rate is from the School Census. It is the mean rate across schools in the municipality. *Elementary literacy rate* is from the 2000 and 2010 Demographic Census. It is the literacy rate of cohorts who are of elementary school age during the mayoral term. In each panel, each point plots an average value within a 7,500 voter bin. Variables on the vertical axis are residualized by population density and election-year fixed effects. Diameter of the points is proportional to the number of observations. Confidence intervals (dashed lines) represent the 95% confidence intervals of a local linear regression (solid red line) with standard errors clustered at the municipality level.

that Brazilian mayors implement may have limited or no influence on economic outcomes. Outside of health and education, mayors are also responsible for local infrastructure, such as public transportation; urban planning; and public health, such as sanitation. It is possible but unlikely that these policies do not influence outcomes such as income and employment. Second, it could be that there are no improvements in economic outcomes in the short term. Outcomes such as income and night lights may take more than 2 to 10 years to improve. Third, improvements in economic outcomes may not be experienced in aggregate, but only among certain populations. For example, I find that increased school resources are concentrated

Table 1.15 Regression discontinuity estimates on municipal education and economic outcomes

Panel A: Education outcomes				
	Drop-out rate	Failing rate	Passing rate	Elem. literacy rate
TwoRound	-1.649^{**}	-0.747	2.330	1.199^{*}
	(0.667)	(1.115)	(1.459)	(0.710)
Single-round mean	3.211	8.645	88.283	91.445
Observations	909	908	909	177
Panel B: Economic outcomes				
	Low income rate	Income per capita	Unemployment rate	Night lights
TwoRound	-5.186^{*}	64.667	-0.964	2.715
	(3.079)	(61.782)	(0.635)	(3.306)
Single-round mean	27.929	762.417	9.815	22.527
Observations	177	177	177	763

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

<u>*Panel A:*</u> Municipal education outcomes. *Drop-out rate*, *Failing rate*, and *Passing rate* are from the School Census. They are the mean rate across schools in the municipality and should add up to 1 in each school. *Elem. literacy rate* is from the 2000 and 2010 Demographic Census. It is the literacy rate of cohorts who are of elementary school age during the mayoral term. <u>*Panel B:*</u> Municipal economic outcomes. *Low income rate*, *Income per capita*, and *Unemployment rate* are from the 2000 and 2010 Demographic Census. *Low income rate* is the fraction of households earning between 0 and 50% of the minimum wage. *Income per capita* is the average monthly household income per capita, in *reais*. *Night lights* is from the 1997-2013 NOAA night lights series. It is the mean night lights level in the municipality. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

in schools at the bottom of the distribution. This may explain the significant effect on the low-income rate, which reflects improved outcomes for the poorest households, and not on more aggregate economic outcomes.

1.5.4 Robustness of the main results

Bias in measures of concentration.– I address two potential sources of mechanical bias in the concentration indices.

One is the size of the parcels (here, polling stations) used to calculate the indices. The number of voters assigned to each polling station is regulated by the TSE, so in principle all polling stations should be of similar size (on average, there are 272 valid votes at each polling station). Empirically, the difference in the number of valid votes at each polling station is small in magnitude (3.605 votes) and insignificant ($p = 0.764$).

A second concern, although potentially a mechanism, is the increased number of candidates in two-round elections (Table 1.9). While the indices may be affected by the number of candidates, the direction of bias is not monotonic (see Section 1.4.4). Nevertheless, I perform three robustness checks.

First, while the number of candidates is a bad control, as it is an endogenous outcome, controlling for the number of candidates does not affect the qualitative results (Table A.10). Second, I simulate the effect of adding an additional candidate to all single-round elections.[31] For most measures, the estimated bias is small and, for the coefficient of variation, of the wrong sign (see the row "Potential bias" in Table 1.5). However, the bias is substantial for the entropy index. Third, I re-calculate all the outcomes using only the vote shares from the top two candidates, in order to maintain the same number of candidates across single- and two-round elections. Doing so does not substantially change the results (Table 1.7).

Calculating the resource index.– The resource index is constructed by taking the first principal component of each school's resources, then calculating a school's percentile rank in the national distribution. Using a resource index constructed by taking the z-scores of each school's resources, then calculating a school's percentile rank in the national distribution, does not affect the qualitative results (Table A.11).

RDD design.– I investigate the robustness of my results to the regression discontinuity design. First, these results are not driven by the choice of bandwidth, whether fixed or chosen by a data-driven method (Figures A.12 and A.13 for voter concentration outcomes; Figures A.14

[31] I do this by adding in a last-place candidate to each polling station. I assign that candidate the number of votes the average last-place candidate receives (1.5%). I take away a proportionate number of votes from the other candidates, to ensure that the total number of voters remains the same. I then calculate the change between the actual value and the simulated value.

and A.15 for municipal school outcomes; Figure A.16 for education outcomes). The estimates maintain similar magnitudes and mostly retain significance for bandwidths out to 150,000 voters, although the estimates for the standard deviation in school resources decline and are not significant at larger bandwidths. Second, dropping controls from the regression – namely, population density and election-year fixed effects – does not substantially affect the results, although the results are noisier (Tables A.17, A.18, and A.19).

Placebo tests.– As discussed in section 1.4.2, I do not find that policy thresholds at 300,000 inhabitants and at 285,714 inhabitants confound my results. First, in placebo regressions, I do not find similar effects at these thresholds for the mayoral electoral outcomes (Tables A.20 and A.21).[32] Second, I show that any effects of these policies is balanced across both sides of the threshold (Figure A.4).

I also show that there are no discontinuities at placebo thresholds in registered voters (170,000; 180,000; 190,000; 210,000; 220,000; 230,000), indicating that the outcomes are relatively continuous at places where the treatment does not change (Figures A.22 and A.23 for voter concentration outcomes; Figures A.24 and A.25 for municipal school outcomes; Figure A.26 for education outcomes). The treatment effect is isolated to the actual threshold: There are no estimates with the same size and significance as at the actual threshold.

1.6 Discussion

In my model, two-round elections lead to different outcomes because candidates adjust their strategies by offering policies that appeal to a broader group of voters. However, there may be other explanations. Namely, different types of candidates may *enter* two-round

[32] Since the number of inhabitants is not the same as the number of registered voters (nor do they map 1:1), to maintain comparability with the baseline estimates, I use a bandwidth of 125,000 inhabitants. This bandwidth was determined by taking half of the population range of municipalities in my 50,000 voter bandwidth (the smallest municipality is 182,082 inhabitants and the largest 434,474 inhabitants). Since the salary cap was implemented in 2000, I estimate this using elections after 2000. Since the legislature size was implemented in 2004, I estimate this using elections after 2004.

elections, or different types of candidates may *win* two-round elections.[33] In the following section, I explore these explanations and provide suggestive evidence that candidates' strategic behavioral responses explain a larger part of the effect of the two-round system.

1.6.1 Selection in candidates

Candidates in two-round elections may have a broader group of supporters because different types of candidates *enter* electoral races. For example, two-round elections may incentivize candidates who are more competent or more relatable to enter. I do not find that candidates in two-round elections are significantly different from candidates in single-round elections along observable characteristics: There is no significant difference in age, sex, educational attainment, or state of birth (Panel A of Table 1.16). Candidates also do not have different occupational backgrounds (Panel B of Table 1.16).

A second possibility is that different types of candidates *win* in two-round elections. However, I also find no evidence that election winners are observably different along the aforementioned characteristics (Table 1.17).

I do find differences in political affiliation among candidates who enter the elections. There are more candidates from small parties and who previously ran as mayoral candidates in two-round elections (Panel A of Table 1.18), but they are not more likely to win (Panel B of Table 1.18).[34] Incumbent candidates are less likely to win two-round elections, but this result is not robust to other bandwidths. Since smaller parties are more likely to appeal to narrower electorates, it is unlikely that the identity of the candidate's party explains the reduced concentration in vote shares.

Small candidates may enter two-round elections for several reasons. First, they may be

[33] A third possibility is that voters behave differently in two-round elections, either through turnout or strategic voting. As mentioned earlier, in Brazil, turnout is mandatory, and so not a major factor. Regarding strategic voting, Fujiwara (2011) finds that third placed and lower candidates receive higher vote shares in two-round elections and argues that voters behave less strategically. While this paper is not focused on voter behavior, I interpret strategic behavioral responses of candidates as an equilibrium outcome that can arise from the electoral rule directly or indirectly through the electoral rule's impact on voter behavior.

[34] I define a "small party" as any party that is not one of the top 5 parties in Brazil by national membership.

Table 1.16 Regression discontinuity estimates on candidate characteristics

Panel A: Demographic characteristics of candidates

	Age	Female	Univ. degree	Born same state
TwoRound	0.470	−0.065	−0.045	−0.070
	(1.514)	(0.040)	(0.048)	(0.049)
Single-round mean	49.955	0.129	0.796	0.784
Observations	264	263	263	263

Panel B: Previous occupation of candidates

	Public sector	Technical	Business
TwoRound	−0.008	0.014	0.017
	(0.063)	(0.059)	(0.025)
Single-round mean	0.466	0.388	0.047
Observations	263	263	263

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Outcomes are the average characteristics of candidates in elections. Panel A contains demographic characteristics. *Univ. degree* is the fraction of candidates whose highest educational attainment is university or higher. *Born same state* is the fraction of candidates who were born in the same state as the election. Panel B contains the industry of candidates' stated previous occupation. *Public sector* includes occupations such as elected positions, judiciary, and workers in public administration. *Technical* includes occupations such as scientists, technicians, and artists. *Business* includes occupations such as administrative positions, workers in commerce and services, and business owners. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

building support for subsequent elections, which is one reason more candidates with previous campaign experience enter two-round elections. Second, these candidates may seek to gain positions in the elected administration. Third, and most interestingly, these candidates may want to influence the top candidates' platforms. In my model, in the two-round election, the third candidate indirectly influences the top candidates' platforms by leading these candidates to offer policies that appeal to her voters. In reality, the third candidate may influence these policies more directly, either by strategically shaping her own policy or direct bargaining.

Table 1.17 Regression discontinuity estimates on characteristics of the winner

Panel A: Demographic characteristics of winners

	Age	Female	Univ. degree	Born same state
TwoRound	0.160	−0.027	0.023	−0.041
	(2.763)	(0.087)	(0.100)	(0.071)
Single-round mean	51.608	0.112	0.832	0.789
Observations	264	263	263	263

Panel B: Previous occupation of winners

	Public sector	Technical	Business
TwoRound	−0.088	0.069	0.059
	(0.137)	(0.131)	(0.044)
Single-round mean	0.534	0.348	0.043
Observations	263	263	263

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Outcomes are the characteristics of the candidate who won the election. Panel A contains demographic characteristics. *Univ. degree* is an indicator for whether the winner's highest educational attainment is university or higher. *Born same state* is an indicator for whether the winner was born in the same state as the election. Panel B contains the industry of candidates' stated previous occupation. *Public sector* includes occupations such as elected positions, judiciary, and workers in public administration. *Technical* includes occupations such as scientists, technicians, and artists. *Business* includes occupations such as administrative positions, workers in commerce and services, and business owners. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

1.6.2 Strategic behavioral responses by candidates

I find limited evidence that two-round elections cause different types of individuals to enter or win the elections. Instead, I find evidence that is consistent with candidates adjusting their strategies during the election.

Reelection incentives.– I rule out that candidates' different strategic behaviors are caused by changes in reelection incentives. I do not find a significant difference in the treatment effect on municipal school resources between mayors in their first term (eligible for reelection) or their second (Table 1.19).

Table 1.18 Regression discontinuity estimates on political affiliation of candidates

	Previous candidacy	Incumbency	Small party	PT party	Governor's party
Panel A: All candidates					
TwoRound	0.539^{**}	-0.182	0.721^{**}	0.038	-0.084
	(0.259)	(0.155)	(0.334)	(0.100)	(0.136)
Single-round mean	1.652	0.752	2.535	0.636	0.584
Observations	263	263	296	296	263
Panel B: Winner only					
TwoRound	-0.144	-0.201^{*}	-0.008	-0.022	-0.012
	(0.130)	(0.115)	(0.127)	(0.094)	(0.120)
Single-round mean	0.621	0.410	0.369	0.187	0.242
Observations	263	263	296	296	263

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Previous candidacy is whether the candidate ran in a previous mayoral election. *Incumbency* is whether the candidate held the position of mayor in a previous term. *Small party* is any party that is not one of the top 5 parties, by national membership. *PT party* is whether the candidate is from the *Partido dos Trabalhadores*. *Governor's party* is whether the candidate is from the party of the incumbent state governor. Dependent variables are either the number of candidates with that characteristic (*Panel A*) or an indicator for the winner having that characteristic (*Panel B*). *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

However, it may be the case that term limits are not relevant for mayors in my sample. In Brazil, mayors face strong incentives to reward their supporters after the election, even those who are not eligible for reelection. Municipal mayors are viewed as important local operatives who deliver votes for their parties at the state and federal levels, and it is not uncommon for mayors to seek office themselves at the state and federal levels.[35]

This suggests that, in line with the model's predictions, different outcomes in two-round elections come instead from candidates adopting different strategies during the campaign, which I discuss in further detail in the following sections.

[35] In my sample, 46.2% of mayors stand as candidates in a state or federal election after their term, and 83.7% of mayors have stood at this level either before or after their term. In larger cities, the position of the municipal mayor is often viewed as a stepping stone to higher office.

Table 1.19 Regression discontinuity estimates on resources in municipal schools, by possibility of re-election

	Mean level of resources		*Standard deviation in resources*	
	Equipment	Infrastructure	Equipment	Infrastructure
TwoRound	0.113^{**}	0.088^{**}	-0.019	-0.044^{**}
	(0.047)	(0.045)	(0.013)	(0.018)
TwoRound * FirstTerm	-0.034	-0.018	-0.008	0.006
	(0.033)	(0.029)	(0.011)	(0.015)
Single-round mean	0.738	0.731	0.121	0.157
Bandwidth size	50,000	50,000	50,000	50,000
Observations	789	789	789	789

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Heterogeneous treatment effects by re-election incentives. *FirstTerm* is a dummy indicating whether the mayor is a first-term mayor (eligible for reelection). *Equipment* and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. The first two columns (*Mean level of resources*) have as the dependent variable the mean index level across schools in the municipality. The last two columns (*Standard deviation in resources*) have as the dependent variable the standard deviation in the index across schools in the municipality. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Concentration between the first and second round.– In Section 1.5.1, I find that the concentration of voters for the top two candidates decreases between the first and second round of the two-round election. In my model, candidates do not change platforms between rounds. In reality, it is more likely that candidates adjust their strategies but are constrained in the extent to which they can do so. There is empirical evidence that candidates qualifying for the second round rally votes from supporters of the candidates eliminated after the first round (Pons and Tricaud, 2018). The decrease in concentration between rounds suggests that candidates adjust their strategies to rally these voters.

Campaign financing.– For the 2004-2012 elections,[36] I am able to observe one aspect of candidates' strategies: how the campaigns are financed. I find that candidates in two-round

[36] I exclude the 2016 elections from the sample, as a new campaign finance law was passed in 2016 that banned donations from corporations.

elections receive *less* donations, both on average (Panel A of Table 1.20) and between the top two candidates (Panel B of Table 1.20). While the outcomes are noisy, the effects are strongest for donations from corporations: Candidates in two-round elections receive *fewer* donations from corporations.[37]

Table 1.20 Regression discontinuity estimates on campaign donations

	Donation amounts received by candidates		
	Total	From individuals	From corporations
Panel A: Average donations per candidate			
TwoRound	-0.225	-0.491	-0.673^{*}
	(0.286)	(0.310)	(0.404)
Single-round mean	12.844	10.742	11.782
Observations	154	154	154
Panel B: Total donations among top two candidates			
TwoRound	-0.074	-0.614	-1.023^{*}
	(0.335)	(0.492)	(0.541)
Single-round mean	14.053	11.717	12.846
Observations	154	154	154

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Panel A: Outcomes are log average donation levels, in *reais*, received by candidates (total donations in the election divided by the number of candidates). *Panel B:* Outcomes are log total donations, in *reais*, received by the top two candidates. Donors identified as *Individual* and *Corporation* depending on whether the donor provided a CPF (individual identification number) or CNPJ (corporate identification number). *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

These results suggest that candidates in two-round elections run their campaigns differently. For example, candidates may offer policies that appeal to individuals rather than corporations. To the extent that corporations represent a narrower swath of the electorate, the pattern in donations suggests that candidates in two-round elections adopt strategies to appeal more broadly.

[37] Donors are classified as corporations or individuals, depending on whether a CPF (individual identification number) or CNPJ (corporate identification number) was filed for the donation.

1.7 Conclusion

A majority of countries use two-round systems to choose their leaders, and an increasing number of countries are adopting this system over time. This paper studies how the electoral rule leads to the election of politicians who represent broader or narrower groups of voters and distribute state resources differently. To identify the effect of the two-round system, I leverage a unique rule in Brazilian municipal elections: Municipalities above a threshold of registered voters hold two rounds, whereas municipalities below this threshold hold a single round. I find that candidates in two-round municipalities are represented by a geographically broader group of voters. Once in office, mayors elected under two-round systems provide more resources to municipal schools and distribute these resources more evenly across schools. I find evidence that downstream educational outcomes are also improved in two-round municipalities.

I present a model to highlight why two-round elections may lead to these empirical results. In my model, the second round raises the marginal return to allocating resources. This creates incentives to (i) increase the government budget and (ii) appeal to voters that candidates in a single-round system would otherwise ignore. The main intuition is that two-round elections perform two functions. First, they require a candidate to attain at least 50% of the vote in order to win. Second, the second round effectively limits the number of candidates to two.

My model proposes that two-round elections lead to different outcomes because candidates adopt different strategies. I find evidence suggesting that two-round systems cause candidates to adjust their behavior rather than cause different types of candidates to enter the races. First, I find evidence indicating that candidates consolidate their bases of support between rounds. Second, candidates in two-round elections adopt more broadly appealing strategies, resulting in fewer corporate donors for their campaigns.

If two-round systems lead to positive outcomes, why is the two-round system not more widely used? The reality is that there are potential trade-offs. First, it may be costly for voters to vote twice in a short span of time. In Brazil, turnout is *lower* in the second round compared to the first round. Uncovering the reasons for this will help better explain the costs of two-round systems. Second, I find that in two-round systems, individuals at lower parts of the distribution benefit. As a result, there may be opposition to implementing two-round systems by richer households or the elite. Identifying barriers to adopting more inclusive institutions is crucial to understanding the process of political reform. Finally, two-round elections may result in better outcomes only when the electorate is composed of many small groups. Brazil is a multi-party system, and the average single-round election has 4.6 candidates running. Incentivizing candidates to incorporate smaller groups in the coalition may lead to better outcomes. This may not translate to contexts where the electorate is composed of two large groups. Providing more empirical evidence of the causal effect of two-round systems in different contexts would greatly advance our overall understanding of electoral systems.

Chapter 2

The Gendered Impact of Anti-Sweatshop Activism in Indonesia

2.1 Introduction

Is international activism effective at improving the lives of overseas workers? In the 1990s, companies such as Nike, Adidas, and Reebok became targets of anti-sweatshop activism and implemented reforms to improve labor standards in their overseas factories. Due to the gendered workforce of the textile, footwear, and apparel industry, these reforms are particularly important given the evidence that better economic opportunities for women can lead to women's empowerment and better household outcomes (Breierova and Duflo, 2004; Duflo, 2003; Qian, 2008). This paper studies whether international anti-sweatshop activism succeeded in generating sustained improvements in labor standards in the textile, footwear, and apparel industry in Indonesia and whether these reforms improved the well-being of female workers.

To quantify these effects, I use a difference-in-differences strategy, by comparing differences among firms and workers before and after the anti-sweatshop movement in either targeted sectors or locations with targeted firms. I find that the anti-sweatshop movement improved labor standards in the textile, footwear, and apparel (TFA) sector, and that these improvements affected the educational composition of female workers and had modest impacts on downstream marriage and fertility outcomes.

This paper focuses on the anti-sweatshop movement and its impact on the TFA sector in Indonesia. In the 1990s, multinational corporations as well as countries where factories were located became targets of anti-sweatshop activism due to revelations of labor abuses and poor labor standards. Activists pressured companies such as Nike, Adidas, and Reebok to improve labor standards in their overseas factories. As part of their response, companies published codes of conduct where they committed to establishing minimum standards, including improving worker pay, restricting forced overtime, and reducing child employment. In some cases, these labor standards were specifically aimed at improving female welfare, such as prohibiting firing of pregnant female workers.

If labor standards improved in factories in Indonesia, the impact of these improvements can disproportionately affect females, due to the heavily female composition of workers in the TFA sector. The TFA sector employs a large fraction of women in the manufacturing sector and has provided an important source of employment for young women. In particular, to the extent that labor standards specifically aimed at improving female welfare were implemented, these improvements will solely benefit females. This paper empirically tests whether the anti-sweatshop movement improved labor standards in TFA factories in Indonesia, and whether this led to improved outcomes for female TFA workers.

I test these hypotheses by examining both firms and workers in Indonesia between 1980 and 2000. I use a difference-in-differences design with two definitions of treatment. I evaluate impacts on firms and workers before and after the anti-sweatshop movement. In one strategy, I compare outcomes *across* sectors by estimating differences between the TFA sector and other manufacturing sectors. In a second strategy, I compare outcomes *within* the TFA sector and *across* geographic locations by estimating differences between locations where Nike, Adidas, and Reebok factories operated and locations where other TFA factories operated. I obtain four main empirical results.

First, I find higher wage growth in the TFA sector, corroborating the results found by Harrison and Scorse (2010). I measure real base wages paid to production workers in manufacturing firms between 1988-2000. I find that wages in the TFA sector grew 3.6 percent faster relative to the rest of manufacturing and 11.6 percent faster in Nike, Adidas, and Reebok subdistricts. The estimated effects are significantly stronger when restricting to exporting firms. I do not find that these improvements are temporary: these effects remain sustained throughout the 1990s. I show that these results are not driven by differences in pre-trends or region-specific shocks. These results indicate that the anti-sweatshop movement increased wages both at targeted firms and within the broader TFA sector in Indonesia.

Second, labor standards improved for female TFA workers, but only within targeted areas. Specifically, I examine whether the incidence of child employment and forced overtime were reduced, by comparing female TFA workers in 1995-2000 relative to 1980 and 1990. I find that female TFA workers in Nike, Adidas, and Reebok districts were 1.6 percentage points less likely to be underage, a reduction of 23.2% of the mean rate, and worked 1.45 fewer hours in the past week, a reduction of 3.3% of the mean weekly hours. On the other hand, I find the opposite effect within the broader TFA sector: compared to the rest of manufacturing, female TFA workers were more likely to be underage and worked more hours in the past week. While the anti-sweatshop movement succeeded in reducing underage employment and forced overtime in targeted areas, these reforms were not implemented in the broader industry.

Third, the anti-sweatshop movement affected female workers' decisions to enter and exit employment in the TFA sector and in the labor force more broadly. Looking at the broader labor force, between 1992-2000 relative to 1980, female labor force participation increased among women 15-25 years of age in geographic areas more affected by the anti-sweatshop movement. Specifically, female labor force participation was 2.6 percentage points higher in subdistricts with a TFA firm relative to urban areas in other subdistricts and 5.5 percentage points higher in Nike, Adidas, and Reebok subdistricts. Labor force participation among broader age groups was unaffected. Looking at the TFA sector, I find that, for TFA workers compared to the rest of manufacturing in 1998-2000 relative to 1980, educational attainment increased among all workers but decreased among workers younger than 25 years. I do not find effects when examining TFA workers in Nike, Adidas, and Reebok subdistricts. Taken together, these results suggest that the anti-sweatshop movement affected the composition of the broader TFA workforce and, given the effects on younger cohorts, a portion of this effect results from different entry decisions by workers into the sector.

Fourth, I find suggestive evidence of some, but limited, effects on downstream marriage and fertility outcomes. Female TFA workers, compared to female workers in other manufacturing sectors in 1998-2000 relative to 1980, were less likely to be married and less likely to have had children. I find similar, but statistically insignificant effects, when examining TFA workers in Nike, Adidas, and Reebok subdistricts. However, I do not find effects on other female empowerment indicators, including the age at marriage, the rate of contraception use, the age at first child, the number of children born, or children's school attendance. While labor standard reforms due to the anti-sweatshop movement had some impacts on women's marriage and fertility decisions, overall, these reforms led to modest gains in female autonomy and empowerment.

This paper is directly related to a literature on the role of manufacturing in developing countries. Broadly, this literature has examined the development effects of the expansion of manufacturing (Atkin, 2015; Brambilla et al., 2012; Bustos, 2011; Verhoogen, 2008). A related literature has focused on labor standards in manufacturing firms. Generally, the literature has been skeptical regarding the effectiveness of multinationals in enforcing compliance with labor standards in overseas firms (Amengual and Distelhorst, 2019; Besley and Ghatak, 2007; Locke et al., 2009; Locke and Romis, 2010; Locke et al., 2007; Short et al., 2015; Toffel et al., 2015). Even if multinationals are effective, there is a larger question of whether the implementation of these standards benefits firms and workers (Besley and Burgess, 2004; Botero et al., 2004; Dragusanu and Nunn, 2018). Both Boudreau (2020) and Harrison and Scorse (2010) study the implementation of better labor standards in multinational firms and its subsequent effects. Boudreau (2020) shows that private enforcement of labor standards can improve compliance with local labor laws in garment factories in Bangladesh and that enforcement of these standards does not lead to adverse effects on labor productivity, wages, or employment. Harrison and Scorse (2010) also study the anti-sweatshop movement in Indonesia. They find that firms increased wages to workers and that the burden of these wage increases mostly fell on firms, through reduced investment and profits and increased probability of closure, and did not lead to reduced employment. This paper builds on these studies by providing evidence that the anti-sweatshop movement succeeded in both improving wages and working conditions in TFA factories in Indonesia and examining the effect of these reforms on workers.

More broadly, this study contributes to a large literature examining the gendered impact of better economic opportunities for women. These gendered impacts typically take two forms. One, better economic opportunities for women improve gender equality within households due to increased returns to investing in female children (Ashraf et al., 2020; Jensen, 2012; Munshi and Rosenzweig, 2006). Two, better economic opportunities for women lead to changes in bargaining power within the household and women empowerment, and eventually to better household outcomes (Breierova and Duflo, 2004; Duflo, 2003; Qian, 2008). Of interest are Heath and Mobarak (2015), who find that girls remain in school and delay marriage and childbirth when garment jobs arrive to their villages, and Atkin (2009), who shows that women who take export manufacturing jobs report stronger bargaining power in the household and have children who are taller. This paper contributes to this literature by studying how improving the *quality* of employment for women can have direct impacts on female workers, both in the broader labor force as well as those employed in the sector.

The remainder of this paper is organized as follows. Section 2.2 describes the TFA sector in Indonesia and the anti-sweatshop movement. Section 2.3 presents the empirical strategy and data sources used. Section 2.4 examines the impact on wages in firms. Section 2.5 presents the results on female workers. Section 2.6 concludes.

2.2 Context

2.2.1 Female employment in Indonesia

Female labor force participation in Indonesia is considered moderate and was relatively stagnant between 1990 and 2000, despite robust economic growth and large increases in female educational attainment. Between 1990 and 2000, the female labor force participation rate was 50.6% with no sustained growth in the period, while female educational attainment increased from 5.5 years in 1990 to 6.9 years in 2000 (Figure 2.1a). However, these numbers have been argued to mask broader trends. Despite decreasing labor force engagement among women in rural areas, the rise of opportunities in urban wage employment boosted female labor force participation in urban areas (Schaner and Das, 2016). In particular, these jobs pay well: in 2000, women employed in the manufacturing sector reported receiving monthly wages that were 23.0% higher than wages in non-manufacturing sectors in urban areas.

In line with this, TFA emerged as a growing industry in Indonesia during this period. The number of production workers employed in medium- and large-scale TFA firms increased from 646,545 in 1990 to 1,232,670 in 2000.[38] In particular, TFA is a remarkably gendered industry. First, the TFA sector employs a large fraction of the female labor force in manufacturing. While firm-level data on female employment is not available before 1995 for the earlier part of the study period, between 1995-2000, the TFA sector employed 44.9% of female production workers in manufacturing despite comprising 22.0% of firms in manufacturing (Figure 2.2a). Second, within TFA firms, the worker composition is heavily female. Between 1995-2000, the TFA sector employed on average 795,299 female production workers each year, or 66.8% of production workers in the TFA sector (Figure 2.2b). Among exporters, this fraction is 70.8%.

Within the broader labor market, manufacturing provides an important source of female employment. Between 1992-2000, female labor force participation rates were on average

[38] Data on production workers comes from the *Statistik Industri*, which comprises all manufacturing firms with over 20 employees.

Figure 2.1 Trends in female labor force participation

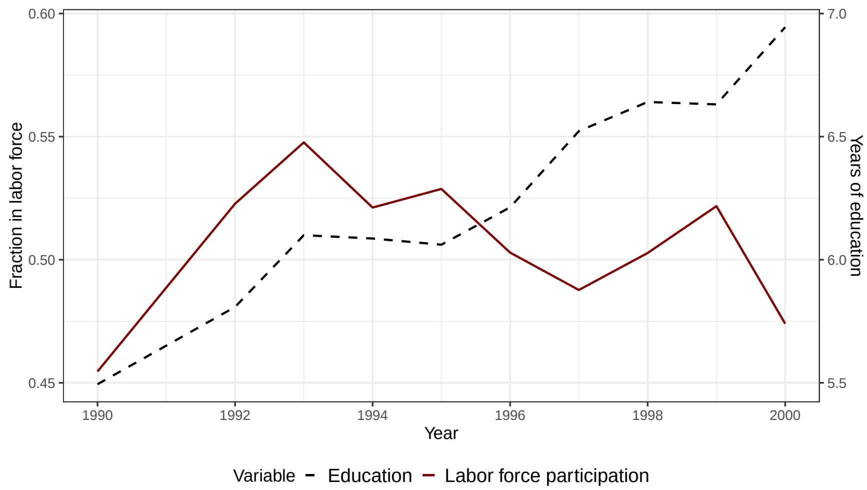

(a) Female labor force participation and education in Indonesia

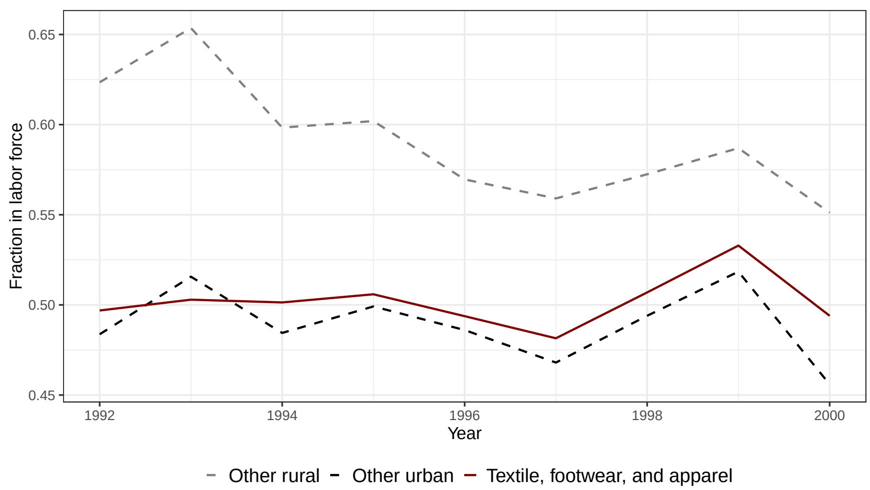

(b) Female labor force participation in subdistricts with textile, footwear, apparel firms relative to other subdistricts

Figures plot the female labor force participation rate in Indonesia. Panel (a) plots the female labor force participation rate (solid red line) and educational attainment (dashed black line). Panel (b) plots the female labor force participation rate in subdistricts with textile, footwear, apparel firms (solid red line) relative to urban areas in other subdistricts (dashed black line) and rural areas in other subdistricts (dashed gray line). *Fraction in labor force* is the number of women whose main activity in the last week was working or looking for work divided by the number of women aged 15-64. *Years of education* is the average years of education among women aged 15-64. Data from the 1990 Population Census and 1992-2000 SUSENAS.

higher in subdistricts where TFA firms operated, at 50.2%, compared to urban areas in other subdistricts, at 48.9% (Figure 2.1b). Within these subdistricts, TFA firms employ a large fraction of the female labor force. While data is only available between 1998-2000,

Figure 2.2 Female employment in textile, footwear, apparel manufacturing

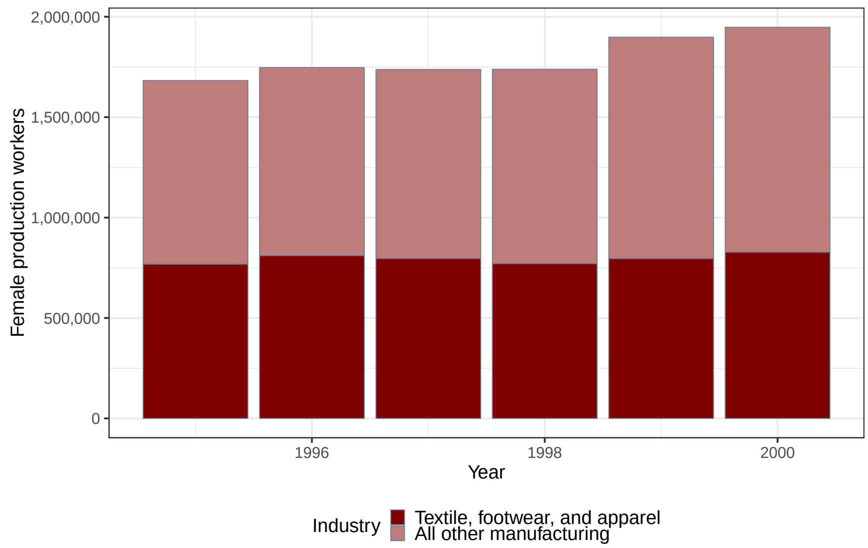

(a) Female production workers in manufacturing

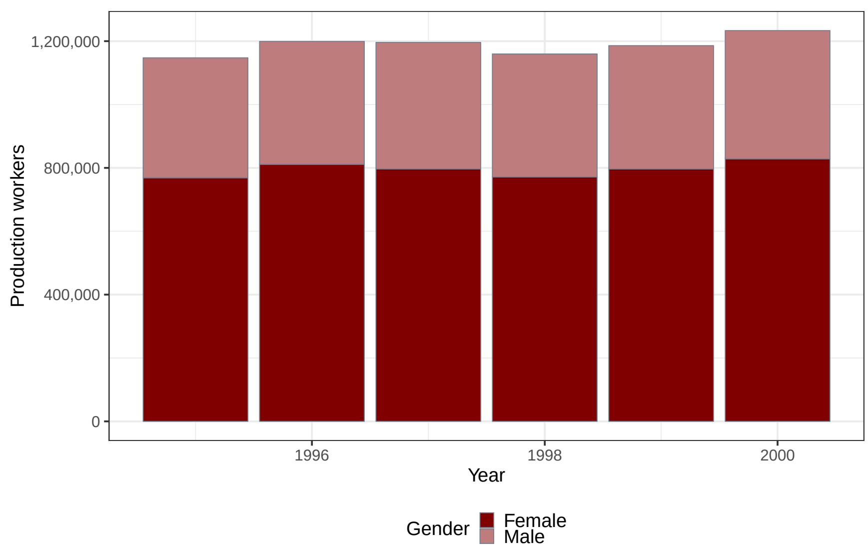

(b) Production workers in textile, footwear, apparel firms

Figures plot the total number of paid production workers employed in manufacturing firms. Panel (a) plots the number of female production workers in textile, footwear, apparel (dark red) vs. all other manufacturing (light red). Panel (b) plots the number of male (light red) vs. female (dark red) production workers in textile, footwear, apparel firms. Data from the 1995-2000 *Statistik Industri*.

in subdistricts where TFA firms operate, 10.4% of the female labor force was employed in the TFA sector and 21.5% was employed in the manufacturing sector. Employment in manufacturing is particularly important for women with a moderate level of education.

Among women with only primary schooling, 15.3% of the female labor force was employed in the TFA sector and 27.7% was employed in the manufacturing sector.

2.2.2 Anti-sweatshop activism

During the 1990s, the TFA sector in Indonesia came under increasingly critical international scrutiny. There had been growing concern through the 1970-1980s among academics and activists regarding labor rights abuses by multinational corporations who had been relocating production to factories in developing countries. International criticism towards Indonesia gained traction in 1989, after several articles about wage protests at Nike factories and a USAID-funded study documenting minimum wage violations in Indonesian factories. Nike in particular became a prominent target after a 1992 article detailed the low wages received by workers.[39] As the anti-sweatshop movement gained steam, governments and international agencies also began exerting pressure, both on companies and countries with lower labor standards. For example, complaints were filed against Indonesia in 1989 under the Generalized Scheme of Preferences to remove Indonesia's tariff privileges with the United States. However, more aggressive action was not carried out until the late 1990s.

The anti-sweatshop movement pushed for the implementation of reforms in multinational corporations in three ways. One was the adoption by corporations of codes of conduct with their suppliers. These codes of conduct were meant to establish the expectation of certain labor standards within factories. Two, the use of external monitors to inspect factories and ensure compliance with the codes of conduct. This was and remains a source of contention, as the independence of external monitors was often questioned. This led to the creation of organizations with purportedly greater independence, such as the Fair Labor Association. Three, allowing workers to organize. In a few cases, this happened directly, where activists directly educated and, in some cases, organized workers in factories. For the most part, this tactic was employed indirectly: activists pushed companies to either sever contracts with

[39] The article "The New Free Trade Heel: Nike's Profits Jump on the Backs of Asian Workers" was published in *Harpers Magazine* in August 1992.

suppliers in countries that did not allow workers to organize or to pressure governments to pass legislation allowing workers to organize.

As the subject of early criticism and pressure, Nike was one of the first movers in response to anti-sweatshop activism. In 1992, Nike formally published a code of conduct to establish minimum labor standards in its factories. In 1994, both Nike and Adidas hired Ernest and Young to conduct factory audits. However, Nike's code of conduct was not fully implemented in its factories until 1995-1996 (Murphy and Mathew, 2001).

The main standards which the anti-sweatshop movement advocated for were centered around wages, work hours, safety standards, and minimum employment ages. Companies were pushed to provide a minimum living wage, often above the legal minimum wage within a country. While not every country did so, some countries, including Indonesia, concurrently increased legal minimum wages due to international pressure.[40] Standards on work hours set out to establish a maximum work week, due to allegations of forced overtime and little overtime pay. Safety standards addressed concerns of indoor air quality, handling of toxic chemicals, and workplace injuries. Finally, reports of underage employment led to demands for minimum employment ages, although companies often set these according to a country's minimum age of employment which could be as young as 14 years in some countries.[41] In its 1992 code of conduct, Nike established a minimum age of 16 years in its apparel factories and of 18 years in its footwear factories. Anti-sweatshop activists also pushed for the implementation of other labor standards, such as prohibiting discrimination and firing of pregnant workers and education programs for workers, although these were adopted by companies to varying degrees.

[40] Minimum wage increases occurred throughout the early 1990s, with the real minimum wage doubling between 1990 and 1995 (Rama, 2001). Indonesia also passed legislation to standardize minimum wages across provinces by establishing a consumption bundle which each province would use to set its minimum wage.

[41] The legal age of employment in Indonesia during the study period was 15 years.

2.3 Empirical strategy

My identification strategy uses a difference-in-differences estimator to quantify the effects of the anti-sweatshop movement. As in Harrison and Scorse (2010), I consider two definitions of treatment: (i) being in the TFA sector, relative to all other manufacturing sectors, and (ii) being in the TFA sector in a location where a Nike, Adidas, or Reebok factory operates, relative to the TFA sector in other locations. Treatment (i) compares across sectors within locations. Treatment (ii) compares within the TFA sector across locations. Locations are defined as either districts or subdistricts, depending on data availability. I compare outcomes before and after 1992, the year Nike published a code of conduct establishing worker wages and labor standards.

I begin with a difference-in-differences regression that compares outcomes before and after 1992 in treated units relative to non-treated units. When estimating the effects on worker wages, the unit is a firm. When estimating the effects on female TFA workers, the unit is an individual. Formally, I estimate the following OLS regression:

$$\log y_{it} = \alpha + \beta_1 \; TREAT_i \cdot POST92_t + \beta_{2d} \; t \cdot \omega_d + \delta_t + \gamma_i + \varepsilon_{it} \tag{2.1}$$

where for a unit i in year t, y_{it} is the outcome variable; $TREAT_i$ is an indicator equal to 1 if a unit is considered treated; $POST92_t$ is an indicator equal to 1 if the year is after 1992; $t \cdot \omega_d$ are district-specific (or province-specific) time trends;[42] δ_t is a time fixed effect; and γ_i is a time-invariant fixed effect. Following Abadie et al. (2017), standard errors are clustered at the level at which treatment is assigned. Sections 2.3.1 and 2.3.2 discuss the fixed effects and standard errors used for each identification strategy in further detail.

The main coefficient of interest is β_1, which indicates the change in outcomes before and after 1992 in treated units relative to non-treated units. β_1 identifies the causal impact of

[42] District-specific time trends are included in analyses at the subdistrict level. Province-specific time trends are included in analyses at the district level.

the adoption of better working standards in TFA factories due to anti-sweatshop activism under the assumption that, absent anti-sweatshop activism, treated units and non-treated units would have had parallel trends.

I also estimate an "event-study" specification where I allow the difference-in-differences estimator to vary by year:

$$\log y_{it} = \alpha + \sum_{t \neq T_0} \beta_{1t}\; TREAT_i \cdot \mathbb{1}[YEAR_t = t] + \beta_{2d}\; t \cdot \omega_d + \delta_t + \gamma_i + \varepsilon_{it} \qquad (2.2)$$

for a reference pre-period T_0, which I define as the first year available prior to 1992. Equation 2.2 allows me both to examine dynamic effects in the post-treatment period and, when multiple pre-treatment periods are available, to test for pre-trends by testing $\beta_{1,t} = 0$ for $t < 1992$.

The inclusion of time-invariant fixed effects allows me to control for unit or location characteristics that are fixed across time. The inclusion of year fixed effects allows me to control for common shocks across all units within each year. I include district-specific (or province-specific) time trends in order to control for long-run trends at the district (or province) level, such as economic growth or demographic changes.

Subdistricts where Nike, Adidas, and Reebok contractors operated are identified by using the factory disclosure lists published by Nike, in 2005, and by Adidas, in 2010.[43] This is an important limitation, as this list reports firms operating after my sample period which may be a biased sample of firms. For example, Nike, Adidas, or Reebok may have canceled contracts with firms unwilling to comply with the new labor standards. As in Harrison and Scorse (2010), I compare my list of firms with those named in newspapers from the time period, and confirm that many of these firms are on my list.

[43] Adidas acquired Reebok in 2005, so the disclosure list in 2010 contains both company's factories.

2.3.1 Effects on wages

To estimate the effect on wages, I define treatment as follows: (i) if a firm is TFA firm, relative to all other manufacturing firms and (ii) if a TFA firm is in a subdistrict (or district) where a Nike, Adidas, or Reebok factory operates, relative to all other TFA firms.

Because the firm-level data is a panel, I estimate a panel regression for treatment (i). In other words, the time-invariant fixed effects are firm fixed effects when estimating equations 2.1 and 2.2. For treatment (ii), the time-invariant fixed effects are subdistrict (or district) fixed effects. Standard errors are clustered at the level at which treatment is assigned: either at the firm level, if treatment (i), or at the subdistrict (or district) level, if treatment (ii).

I measure the real base wage paid to production workers using the 1988-2000 *Statistik Industri*, the annual manufacturing survey conducted by the Indonesian government's statistical agency *Badan Pusat Statistik*. The survey covers all manufacturing firms in Indonesia with over 20 employees. Between 14,664 (in 1988) and 22,997 (in 1996) firms are represented in the survey. The survey is a panel, with each firm having on average 6.2 observations during this period. Real wages are the base wage paid to production workers, deflated by the consumer price index (CPI), where the CPI is equal to 100 in 2010.[44]

Prior to 1992, production worker wages were, on average, 1,008.6 thousand rupiahs (approximately 110 USD) lower in TFA firms compared to the rest of the manufacturing sector (Figure 2.3a). By 1996, the year after Nike fully implemented its new code of conduct, this gap narrowed to 825.6 thousand rupiahs (90 USD). Similarly, within TFA firms prior to 1992, production worker wages were 1042.3 thousand rupiahs (110 USD) lower in subdistricts where Nike, Adidas, and Reebok factories operated (Figure 2.3b). By 1996, firms in these subdistricts were paying production workers *more* – on average, 2642.4 thousand rupiahs (290 USD). Equation 2.1 empirically tests whether these differences narrowed significantly after 1992.

[44]Wages are calculated by dividing the firm's wage/salary expenses for production workers (excluding overtime, gifts, bonuses, etc.) by the total number of paid production workers.

Figure 2.3 Trends in production worker wages

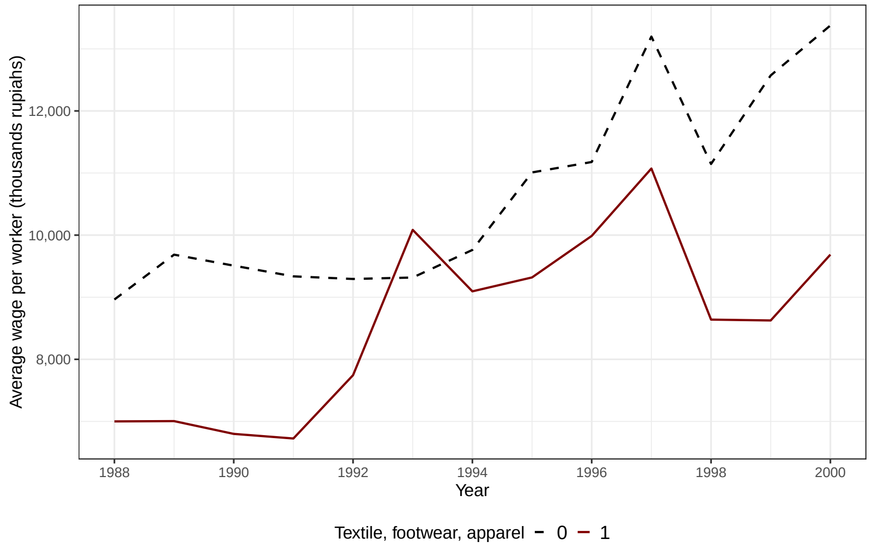

(a) Wages in manufacturing

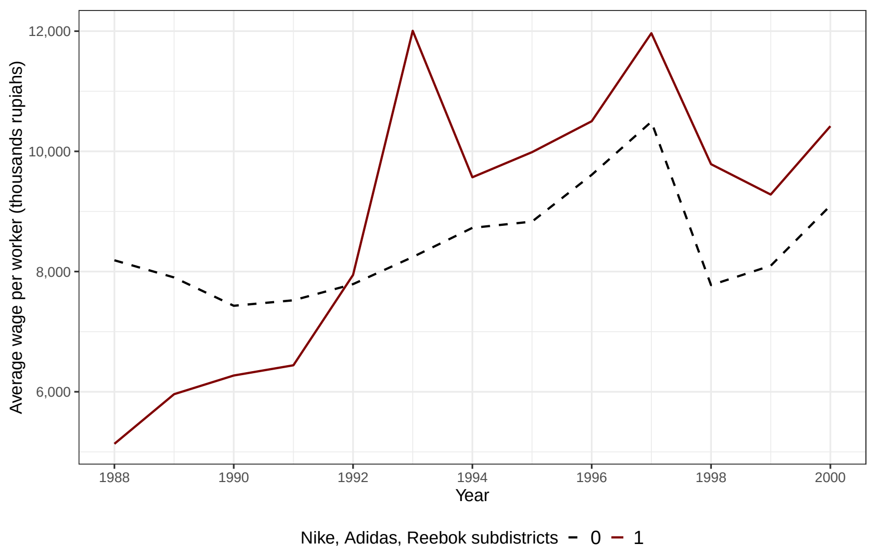

(b) Wages in textile, footwear, apparel firms

Figures plot the real wage paid to production workers. Panel (a) plots the wage in textile, footwear, apparel (dashed line) vs. all other manufacturing (solid line) firms. Panel (b) plots the wage in textile, footwear, apparel firms in Nike, Adidas, Reebok subdistricts (dashed black line) vs. in all other subdistricts (solid red line). A firm is considered in textile, footwear, apparel if the firm reported this as its main product for all years between 1988-1991. Sample comprises firms reporting exporting a positive percentage of output for all years between 1992-1995. Data from the 1988-2000 *Statistik Industri*.

I estimate equations 2.1 and 2.2 separately on the full sample of firms and on a subsample of international firms. To identify international manufacturing firms, I restrict the sample to firms reporting exporting a positive percentage of output for all years between 1992, when Nike established its code of conduct, and 1995, when these reforms were fully implemented.

For treatment (i), I consider a firm to be in textile, footwear, and apparel if this was the main product reported by the firm for all years in the pre-treatment period between 1988-1991. For treatment (ii), because subdistricts are only available in the 1994-1997 *Statistik Industri*, I identify subdistricts where Nike, Adidas, and Reebok factories operated based on the subdistrict a firm reported operating in in 1994. When expanding treatment (ii) to the district level, I identify districts where Nike, Adidas, and Reebok factories operated based on the district a firm reported operating in for that year.

2.3.2 Effects on female textile, footwear, and apparel workers

To estimate effects on female TFA workers, I define treatment as follows: (i) if a woman is employed in the TFA sector, relative to women employed in all other manufacturing sectors within the same subdistrict (or district) and (ii) if a woman employed in the TFA sector resides in a subdistrict (or district) where a Nike, Adidas, or Reebok factory operates, relative to all other women employed in the TFA sector.

For treatment (i), the time-invariant fixed effects are subdistrict-industry (or district-industry) fixed effects when estimating equations 2.1 and 2.2. For treatment (ii), the time-invariant fixed effects are subdistrict (or district) fixed effects. Standard errors are clustered at the level at which treatment is assigned: either at the subdistrict-industry (or district-industry) level, if treatment (i), or at the subdistrict (or district) level, if treatment (ii).

To measure outcomes for female workers, I draw on data from three sources: the 1992-2000 *Survei Sosial Ekonomi Nasional*, also known as SUSENAS; the 1980 and 1990 *Sensus Penduduk*, the decennial population census; and the 1995 *Survei Penduduk Antar Sensus*, the intercensal survey, also known as SUPAS. These are conducted by the Indonesian government's statistical agency *Badan Pusat Statistik*. SUSENAS is a nationally representative household survey comprising around 200,000 households. For the censuses, I use the 5% sample of the 1980 census, the 5% sample of the 1990 census, and the 0.5% sample of the 1995 intercensal survey.

Because each data source contains different variables and different identifiers for geographic location and industry, various subsets of the data are used for each outcome and specification based on the data availability. Table B.1 lists the variables and identifiers available for each data set. The pre-treatment period comprises 1980 for the subdistrict level analysis and comprises 1980 and 1990 for the district level analysis. The availability of subdistrict identifiers is a limitation of the data, as 1980 is the most recent pre-treatment period for the subdistrict level analysis, over 10 years prior to the onset of the anti-sweatshop movement. In addition to this, for analysis at the subdistrict level, I am unable to test for the presence of pre-trends.

I restrict the sample to women i) residing in subdistricts (or districts) where TFA firms operate and ii) employed in the manufacturing sector. Subdistricts (or districts) where TFA firms operate in are identified from the 1994 *Statistik Industri*, which is the earliest survey year with firm subdistrict locations. TFA firms were located across 237 subdistricts in 76 districts, of which 44 subdistricts in 19 districts contained Nike, Adidas, and Reebok factories.

For treatment (i), I consider an individual to be in textile, footwear, and apparel if this was the industry reported as the individual's main job. Necessarily, these individuals are currently working, as an individual's industry is not collected if she is looking for work. For treatment (ii), I identify Nike, Adidas, and Reebok subdistricts (or districts) based on the individual's subdistrict (or district) of residence, ie. where she was surveyed.

When estimating the effects on female labor force participation, I conduct a slightly different analysis. I define treatment as follows: (i) if a woman resides in a subdistrict (or district) where a TFA firm operates, relative to women residing in urban areas in all other subdistricts (or districts) and (ii) if a woman resides in a subdistrict (or district) where a Nike, Adidas, or Reebok factory operates, relative to women residing in subdistricts (or districts) where a TFA firm operates.

2.3.3 Identification concerns

Each identification strategy identifies a specific treatment effect and has its advantages and disadvantages. Treatment (i) identifies the effect of being in the TFA sector relative to other manufacturing sectors. By comparing firms or workers in the same subdistrict (or district), treatment (i) keeps constant confounding factors that are location-specific, such as access to markets or migration decisions. However, selection into industries (for both firms and workers) is non-random and comparing firms and workers in different industries may not yield the average treatment effect if the sectoral bias leads to differential trends between sectors. Treatment (ii) identifies the effect of being in a location where a Nike, Adidas, or Reebok factory operates relative to other locations. By comparing firms or workers in the TFA sector, treatment (ii) eliminates the sectoral bias. However, treatment (ii) has two disadvantages. First, firms and workers may choose to locate in one subdistrict (or district) over another. If this locational bias leads to differential trends between locations, comparing firms and workers in subdistricts (or districts) where Nike, Adidas, or Reebok operate will not identify the average treatment effect. Second, if markets are sufficiently integrated across locations, the effects of the anti-sweatshop movement may spill over into other firms in the TFA sector, which will bias any estimated effects towards zero.

One concern is that despite including time fixed effects, unit or location fixed effects, and district-specific (or province-specific) time trends, omitted variable bias remains. Specifically, there may be time-varying characteristics that are correlated with treatment and that lead to differential trends between treated and non-treated units. I address this by estimating an alternative specification where I replace the time fixed effects with province-by-year (or island-by-year) fixed effects.[45] Doing so allows me to control for common regional shocks, such as weather shocks.

[45] Province-by-year fixed effects are included in analyses at the subdistrict level. Province-by-island fixed effects are included in analyses at the district level.

2.4 The impact of anti-sweatshop activism on firms

Did the anti-sweatshop movement succeed in increasing wages in firms manufacturing textile, footwear, and apparel? In this section, I present evidence that anti-sweatshop activism increased wages paid to production workers in the textile, footwear, and apparel sector, and significantly so in firms located in Nike, Adidas, and Reebok subdistricts.

2.4.1 Effects on wages

I find that wages grew significantly more between 1992-2000 relative to 1988-1991 in the TFA sector relative to the rest of manufacturing (Table 2.4). Wages in the TFA sector grew 3.6 percent faster relative to the rest of manufacturing (Column 1). This effect is particularly pronounced among exporters (Column 2) a.nd firms located in subdistricts where Nike, Adidas, and Reebok factories operate (Columns 3 and 4). Among exporters, wages in the TFA sector grew 13.8 percent faster and, in Nike, Adidas, and Reebok subdistricts, 16.8 percent faster. Since anti-sweatshop activism was aimed primarily at firms exporting products to international markets, it is not surprising that the impacts on wage growth are concentrated among exporting firms.

Wage increases in the TFA sector were small to zero in 1992 but began taking effect in 1993 and remained sustained throughout the period until 2000 (Figure 2.5). Among exporters, the treatment effects are lower after 1998, suggesting that while wage increases were sustained in the TFA sector, the growth became more modest in later years. It is also possible that more modest growth after 1998 is due to the 1997 Asian financial crisis.

I next turn to estimating wage growth among TFA firms, but focus on comparing TFA firms in areas more heavily targeted by anti-sweatshop activism with those in other areas. I find that wage growth was significantly higher in locations with Nike, Adidas, and Reebok factories (Table 2.6). Wages in TFA firms grew 11.6 percent faster in Nike, Adidas, and Reebok subdistricts compared to other subdistricts (Column 1). Among exporters, the

Table 2.4 Estimates on production worker wages: all manufacturing

	Treatment variable: Firms producing textile, footwear, apparel			
Sample	Full sample	Exporters	Nike, Adidas, Reebok subdistricts	Exporters in Nike, Adidas, Reebok subdistricts
	(1)	(2)	(3)	(4)
TREAT × POST92	0.036***	0.138***	0.083***	0.168**
	(0.012)	(0.052)	(0.029)	(0.077)
Firms	40,103	2,677	3,029	486
Observations	169,575	13,537	23,259	3,065
Dep. var. mean	8.627	8.923	8.995	9.079

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variable is the log real base wage paid to production workers. Treatment is firms producing textile, footwear, apparel for all years between 1988-1991, relative to the rest of manufacturing. *Full sample* is the full sample of firms. *Exporters* is the sample of firms reporting exporting a positive percentage of output for all years between 1992-1995. *Nike, Adidas, Reebok subdistricts* is the sample of firms located in subdistricts where Nike, Adidas, Reebok factories operate. *F-stat* for coefficient on TREAT × POST92 in (1) = (2) is 9.51 ($p = 0.002$), in (1) = (3) is 10.90 ($p = 0.001$), in (2) = (4) is 7.32 ($p = 0.007$), and in (3) = (4) is 7.34 ($p = 0.007$). *Estimation method:* Regression includes year fixed effects, firm fixed effects, and district-specific time trends. Standard errors clustered at the firm level.

treatment effect is significantly larger: wage growth was 21.1 percent higher in Nike, Adidas and Reebok subdistricts (Column 2). This mirrors the result in the previous analysis, in that treatment effects are stronger among firms exporting products to international markets. I find similar effects when expanding the geographical area to the district level, by comparing TFA firms in districts where Nike, Adidas, and Reebok operate with TFA firms in other districts. Wages grew 5.6 percent and 7.2 percent more in Nike, Adidas, and Reebok districts among all firms and among exporting firms, respectively, although the treatment effects are not statistically significant at the 10% level among exporters (Columns 3 and 4).

To test the validity of the difference-in-differences estimates, I test for differences in pre-trends by estimating the difference-in-differences estimate separately by year among the sample of exporting firms. I do not find evidence of statistically significant differences in pre-trends prior to 1992 when the treatment is defined as the TFA sector (Figure 2.5) or being

Figure 2.5 Estimates by year of production worker wages: all manufacturing

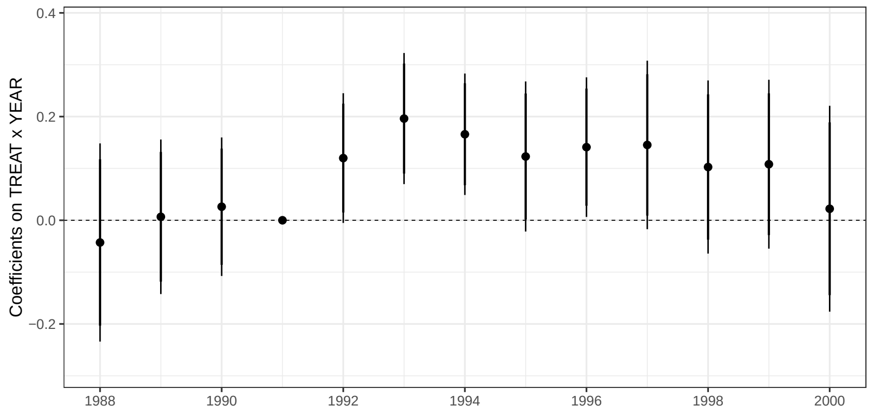

(a) Exporters

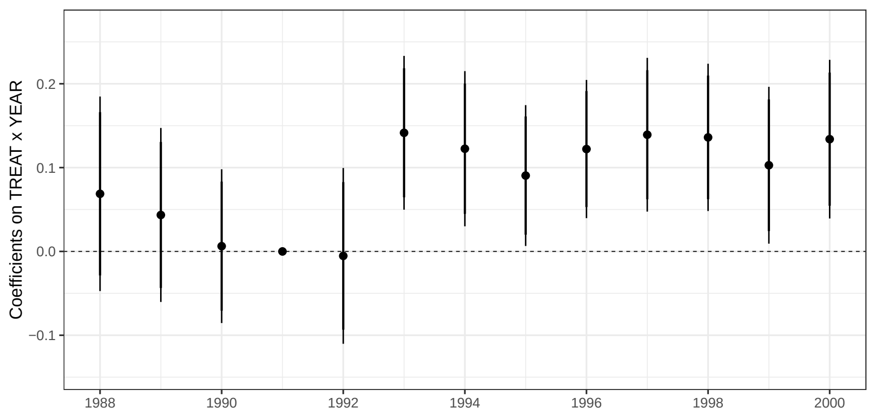

(b) Firms in Nike, Adidas, Reebok subdistricts

Dependent variable is the log real base wage paid to production workers. Treatment is firms producing textile, footwear, apparel for all years between 1988-1991, relative to the rest of manufacturing. The thicker vertical lines represent the 90% confidence interval and the thinner vertical lines represent the 95% confidence interval. *Exporters* is the sample of firms reporting exporting a positive percentage of output for all years between 1992-1995. *Nike, Adidas, Reebok subdistricts* is the sample of firms located in subdistricts where Nike, Adidas, Reebok factories operate. *Estimation method:* Regression includes year fixed effects, firm fixed effects, and district-specific time trends. Standard errors clustered at the firm level.

located in a Nike, Adidas, and Reebok subdistrict (Figure 2.7). Further, the estimated effects are of similar magnitude and significance after including province-by-year (or island-by-year) fixed effects to account for regional shocks (Tables B.2 and B.3).

Table 2.6 Estimates on production worker wages: textile, footwear, apparel

	Treatment variable: Firms in Nike, Adidas, Reebok subdistricts		*Treatment variable:* Firms in Nike, Adidas, Reebok districts	
Sample	All TFA (1)	TFA exporters (2)	All TFA (3)	TFA exporters (4)
TREAT × POST92	0.116***	0.211**		
	(0.032)	(0.103)		
TREAT × POST92			0.056*	0.072
			(0.029)	(0.081)
Firms	4,966	345	4,966	345
Observations	32,555	2,802	37,898	3,092
Dep. var. mean	8.598	8.824	8.598	8.824

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variable is the log real base wage paid to production workers. For Columns 1-2, treatment is textile, footwear, apparel firms operating in Nike, Adidas, Reebok subdistricts, relative to textile, footwear, apparel firms in other subdistricts. For Columns 3-4, treatment is textile, footwear, apparel firms operating in Nike, Adidas, Reebok districts, relative to textile, footwear, apparel firms in other districts. *All TFA* is the full sample of firms producing textile, footwear, apparel for all years between 1988-1991. *TFA exporters* is the sample of textile, footwear, apparel firms reporting exporting a positive percentage of output for all years between 1992-1995. *F-stat* for coefficient on TREAT × POST92 in (1) = (2) is 8.71 ($p = 0.003$) and in (3) = (4) is 1.94 ($p = 0.165$). *Estimation method:* Regression includes year fixed effects, subdistrict fixed effects (in Columns 1-2), district fixed effects (in Columns 3-4), district-specific time trends (in Columns 1-2), and province-specific time trends (in Columns 3-4). Standard errors clustered at the subdistrict level (Columns 1-2) or district level (Columns 3-4).

2.5 The impact of anti-sweatshop activism on female workers

In this section, I ask three questions. One, did the anti-sweatshop movement improve the working conditions of female workers in textile, footwear, and apparel? In Section 2.5.1, I find that the movement succeeded in achieving its goals of reducing underage employment and forced overtime, but only in areas targeted by the activism and not in the textile, footwear, and apparel industry more broadly. Two, did the anti-sweatshop movement affect the composition of workers in the sector? Given that wages and labor standards improved, different types of workers may choose to enter and exit employment in textile, footwear, and apparel and the labor force more broadly. In Section 2.5.2, I find that female labor force

Figure 2.7 Estimates by year of production worker wages: textile, footwear, apparel

(a) All TFA

(b) TFA exporters

Dependent variable is the log real base wage paid to production workers. Treatment is textile, footwear, apparel firms operating in Nike, Adidas, Reebok subdistricts, relative to textile, footwear, apparel firms in other subdistricts. The thicker vertical lines represent the 90% confidence interval and the thinner vertical lines represent the 95% confidence interval. *All TFA* is the full sample of firms producing textile, footwear, apparel for all years between 1988-1991. *TFA exporters* is the sample of textile, footwear, apparel firms reporting exporting a positive percentage of output for all years between 1992-1995. *Estimation method:* Regression includes year fixed effects, subdistrict fixed effects, and district-specific time trends. Standard errors clustered at the subdistrict level.

participation increased in locations affected by the anti-sweatshop activism and that the educational composition of workers in the broader textile, footwear, and apparel sector was affected. Specifically, educational attainment increased among all workers but decreased among younger cohorts of workers. Three, did the anti-sweatshop movement affect more downstream outcomes on female autonomy and empowerment? Improved wages and labor standards can lead to both higher household income and empowerment of women within the household. In particular, specific labor standards that were advocated, such as maternity policies, may directly affect women's marriage and fertility decisions. In Section 2.5.3, I find some, but limited, effects on women's marriage decisions and fertility outcomes. Specifically, I find that fewer female textile, footwear, and apparel workers are married and have children.

2.5.1 Labor standards

Section 2.4.1 demonstrated the anti-sweatshop movement succeeded in increasing wages paid to production workers in TFA firms. In this section, I directly examine female TFA workers and ask: did the anti-sweatshop movement achieve its other goals around labor standards? Outside of increasing wages, the anti-sweatshop movement called for implementing minimum employment ages and reducing forced overtime. I find that while the broader TFA sector did not achieve these reforms, there is suggestive evidence that these reforms were achieved in areas targeted by the activism.

When comparing women employed in the TFA sector relative to the rest of manufacturing, I find that women employed in the TFA sector were *more* likely to be underage and worked *more* hours in the past week (Columns 1 and 2 of Table 2.8) – effects that go in the opposite direction of the anti-sweatshop movement's goals. These effects are much weaker and statistically insignificant when comparing within districts (Column 2), but the district-level analysis is unable to control for time-invariant characteristics at a fine geographic level. As a result, I rely on the subdistrict-level analysis: compared to women employed in other manufacturing sectors, women employed in the TFA sector were 2.3 percentage points more likely to be underage and worked 2.94 more hours in the past week in 1998-2000 compared to 1980 (Column 1).

These results suggest that some of the labor standards advocated for by anti-sweatshop activists were not implemented in the broader TFA sector and likely worsened. I next examine labor standards in areas targeted by the activism by comparing female TFA workers

Table 2.8 Estimates on labor standards among female textile, footwear, apparel workers

	Treatment variable: Employed in textile, footwear, apparel		*Treatment variable:* Location with a Nike, Adidas, Reebok factory	
Geographic level of analysis	Subdistrict (1)	District (2)	Subdistrict (3)	District (4)
Panel A *Dependent variable: Younger than 16 years of age*				
TREAT × POST92	0.023**	−0.005		
	(0.011)	(0.004)		
TREAT × POST92			−0.009	−0.016**
			(0.026)	(0.008)
Observations	25,689	169,129	11,564	67,397
Dep. var. mean	0.075	0.073	0.059	0.069
Panel B *Dependent variable: Hours worked in past week*				
TREAT × POST92	2.939***	0.313		
	(0.846)	(0.511)		
TREAT × POST92			−0.280	−1.449*
			(1.578)	(0.755)
Observations	25,113	165,019	11,293	65,768
Dep. var. mean	42.927	41.352	43.995	43.538
Pre-period	1980	1980, 1990	1980	1980, 1990
Post-period	1998-2000	1995, 1998-2000	1998-2000	1995, 1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variables are an indicator equal to 1 if the individual is younger than 16 years of age (Panel A) and number of hours worked in the past week (Panel B). For Columns 1-2, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Columns 3-4, treatment is a female textile, footwear, apparel worker residing in a subdistrict or district where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in subdistricts or districts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes year fixed effects, subdistrict-industry fixed effects (Column 1), district-industry fixed effects (Column 2), subdistrict fixed effects (Column 3), district fixed effects (Column 4), district-specific time trends (Columns 1 and 3), and province-specific time trends (Columns 2 and 4). Standard errors clustered at the subdistrict-industry (Column 1), district-industry (Column 2), subdistrict (Column 3), or district (Column 4) level.

in locations with a Nike, Adidas, or Reebok factory to female TFA workers in other locations. I find that labor standards improved in districts with Nike, Adidas, and Reebok factories (Columns 3 and 4 in Table 2.8). In 1995 and 1998-2000 compared to 1980 and 1990, female TFA workers were 1.6 percentage points *less* likely to be underage and worked 1.45 *fewer* hours in the past week (Column 3). These are sizeable improvements: Nike, Adidas, and Reebok districts experienced reductions in the fraction of underage workers by 23.2% of the mean rate and in the number of hours worked by 3.3% of the mean weekly hours.

However, I caveat that the latter results are suggestive for two reasons. One, while the estimates are of the same sign, the treatment effects are weaker and statistically insignificant when comparing subdistricts (Column 3) rather than districts (Column 4). The subdistrict results may be underpowered or the district results may be confounded by differential trends among districts. Second, I find evidence of significant differences in pre-trends in the district-level analysis (Figure 2.9). While the pre-trends are of the wrong sign, caution is needed in interpreting the district-level results. Reassuringly, the results are robust to replacing the year fixed effects with province-by-year or island-by-year fixed effects, which indicate that the results are not driven by region-specific shocks (Table B.4).

2.5.2 Composition effects on workers

Increased wages and better labor standards in the TFA sector may affect workers' decisions to enter and exit employment in the TFA sector and the labor force more broadly. While I cannot observe the entry and exit of workers into the TFA sector, I examine whether there are differences in labor force participation and the composition of workers in the TFA sector. I first examine effects on female labor force participation in geographic areas more affected by the anti-sweatshop movement. I next examine whether there are changes in the educational composition of TFA workers.

Female labor force participation.– I find positive effects on female labor force participation rates among women 15-25 years of age (Table 2.10). Between 1992-2000 compared to 1980, labor force participation among women 15-25 years of age was 2.6 percentage points higher in subdistricts with a TFA firm relative to urban areas in other subdistricts (Column 2) and 5.5 percentage points higher in subdistricts with a Nike, Adidas, or Reebok factory relative to subdistricts with a TFA firm (Column 4). I do not find differences in female labor force

Figure 2.9 Estimates by year on labor standards in locations with a Nike, Adidas, Reebok factory

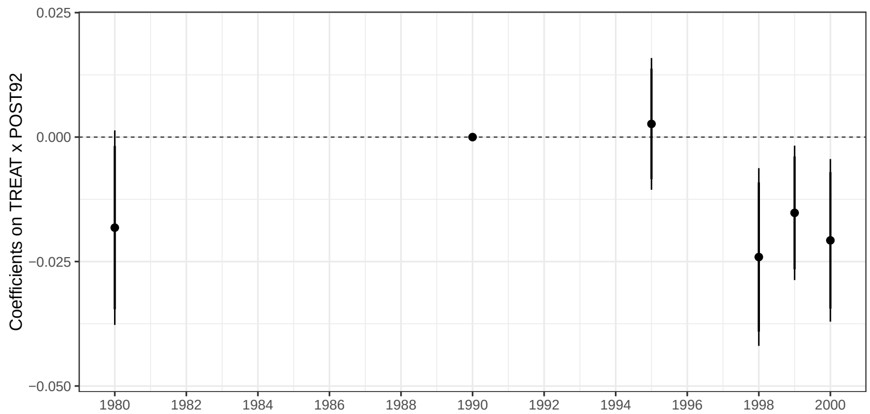

(a) Dependent variable: Younger than 16 years of age

(b) Dependent variable: Hours worked in past week

Dependent variable in Panel (a) is an indicator equal to 1 if the individual is younger than 16 years of age. Dependent variable in Panel (b) is the number of hours worked in the past week. Treatment is a female textile, footwear, apparel worker residing in a district where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in districts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes year fixed effects, district fixed effects, and province-specific time trends. Standard errors clustered at the district level.

participation rates among broader age groups: the effects are much smaller in magnitude and statistically insignificant (Columns 1 and 3).

However, estimates from comparing subdistricts are not consistent with estimates from comparing districts (Table B.5). Districts with TFA firms experienced *lower* rates of female labor force participation in 1992-2000 compared to 1980-1990 (Columns 1-2). Examining the

Table 2.10 Estimates on female labor force participation

	Treatment variable: Subdistrict with a textile, footwear, apparel firm		*Treatment variable:* Subdistrict with a Nike, Adidas, Reebok factory	
Age group	15-64 (1)	15-25 (2)	15-64 (3)	15-25 (4)
TREAT × POST92	0.005 (0.008)	0.026*** (0.009)		
TREAT × POST92			0.026 (0.022)	0.055** (0.023)
Observations	2,084,724	804,999	460,141	188,530
Dep. var. mean	0.402	0.339	0.373	0.352
Pre-period	1980	1980	1980	1980
Post-period	1992-2000	1992-2000	1992-2000	1992-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variable is an indicator equal to 1 if the individual is in the labor force. For Columns 1-2, treatment is residing in a subdistrict where a textile, footwear, apparel firm operates, relative to residing in urban areas in other subdistricts. For Columns 3-4, treatment is residing in a subdistrict where a Nike, Adidas, Reebok factory operates, relative to residing in subdistricts where a textile, footwear, apparel firm operates. An individual is considered in the labor force if the main activity in the last week was working or looking for work. *Estimation method:* Regression includes year fixed effects, subdistrict fixed effects, and district-specific time trends. Standard errors clustered at the subdistrict level.

estimates by year suggests there are statistically significant differences in pre-trends at the district level (Figures B.6). Because the specification in Table 2.10 controls for district-specific time trends, this increases confidence in the estimates from comparing subdistricts. However, because the only pre-period available is 1980, I cannot test for pre-trends at the subdistrict level. Nevertheless, it is unlikely that the results at the subdistrict level are confounded by broader differential trends in female labor force participation, as I do not find effects among women 15-64 years and only find effects among women 15-25 years of age, who are the predominant age group among females employed in TFA.[46]

[46]This is true for manufacturing as well. 46.4% of women employed in manufacturing and 56.5% of women employed in TFA were younger than 25 years of age. Data from the 1990 Demographic Census.

Educational composition of workers.– Next, I examine the educational composition of TFA workers. I measure the highest educational level attained and examine effects separately for all female workers and younger female workers. Any impacts observed on the educational composition of all female workers will reflect both entry and exit decisions of workers. Analyzing younger cohorts of workers allows me to better isolate the impact of the anti-sweatshop movement on entry decisions.

First, I find effects on the educational composition of workers in the broader TFA sector but no effects in areas targeted by the anti-sweatshop movement (Columns 1 and 3 of Table 2.11). In 1998-2000 compared to 1980, females employed in the TFA sector were 5.5 percentage points more likely to have completed secondary schooling (Column 1 of Panel C). This is a large effect, comprising an increase of 42.0% on the mean rate of secondary schooling. While not statistically significant, the fraction of female workers with only primary schooling decreased by 3.6 percentage points (Column 1 of Panel B), suggesting that the composition of TFA workers shifted from lower to higher levels of educational attainment. I do not find statistically significant differences in the educational composition of TFA workers in Nike, Adidas, and Reebok subdistricts compared to TFA workers in other subdistricts (Column 3). However, as the magnitudes of these estimates are not close to zero, I cannot rule out that the educational composition of TFA workers in Nike, Adidas, and Reebok subdistricts did not shift towards more highly educated workers.

Second, I find significantly different effects among younger cohorts of female workers in the broader TFA sector (Column 2 of Table 2.11). While educational attainment of workers overall increased, educational attainment of younger workers decreased: in 1998-2000 compared to 1980, women younger than 25 years in the TFA sector were 11.5 percentage points *more* likely to have no schooling (Column 2 of Panel A) and 13.3 percentage points *less* likely to have completed primary schooling (Column 2 of Panel B), compared to women younger than 25 years in other manufacturing sectors.

I find qualitatively similar results in the district-level analysis (Table B.7). In particular,

Table 2.11 Estimates on educational composition of female textile, footwear, apparel workers

	Treatment variable: Employed in textile, footwear, apparel		*Treatment variable:* Subdistrict with a Nike, Adidas, Reebok factory	
Age group	All (1)	$\leq$ 25 years (2)	All (3)	$\leq$ 25 years (4)
Panel A *Dependent variable: No schooling*				
TREAT × POST92	0.016	0.115***		
	(0.023)	(0.025)		
TREAT × POST92			−0.022	−0.026
			(0.100)	(0.098)
Observations	25,690	15,052	11,565	6,949
Dep. var. mean	0.460	0.410	0.390	0.337
Panel B *Dependent variable: Primary schooling*				
TREAT × POST92	−0.036	−0.133***		
	(0.029)	(0.038)		
TREAT × POST92			−0.009	0.008
			(0.126)	(0.111)
Observations	25,690	15,052	11,565	6,949
Dep. var. mean	0.391	0.449	0.475	0.535
Panel C *Dependent variable: Secondary schooling*				
TREAT × POST92	0.055***	0.036		
	(0.017)	(0.027)		
TREAT × POST92			0.042	0.019
			(0.040)	(0.039)
Observations	25,690	15,052	11,565	6,949
Dep. var. mean	0.131	0.132	0.125	0.123
Pre-period	1980	1980	1980	1980
Post-period	1998-2000	1998-2000	1998-2000	1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variables are indicators for the highest level of schooling attained: none (Panel A), primary (Panel B), and secondary (Panel C). For Columns 1-2, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Columns 3-4, treatment is a female textile, footwear, apparel worker residing in a subdistrict where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in subdistricts where a textile, footwear, apparel firm operates. *F-stat* for coefficient on TREAT × POST92 in (1) = (2) is 8.71 ($p = 0.003$) for no schooling (Panel A), 7.29 ($p = 0.007$) for primary (Panel B), and 4.63 ($p = 0.032$) for secondary (Panel C). *Estimation method:* Regression includes year fixed effects, subdistrict-industry fixed effects (Columns 1-2), subdistrict fixed effects (Columns 3-4), and district-specific time trends. Standard errors clustered at the subdistrict-industry (Columns 1-2) or subdistrict (Columns 3-4) level.

the shift of workers to higher levels of schooling in the broader TFA sector is stronger: in 1995 and 1998-2000 compared to 1980 and 1990, female TFA workers were *more* likely to have secondary schooling and *less* likely to have primary schooling (Column 1 of Panels B and C). The subdistrict-level results are robust to replacing year fixed effects with province-by-year fixed effects (Table B.8).

These results suggest that higher wages in the TFA sector affected the composition of workers, but that better labor standards did not do so. Wages increased both in the broader TFA sector and in areas targeted by anti-sweatshop activism (Section 2.4.1), but labor standards did not improve in the broader TFA sector and only improved in areas targeted by anti-sweatshop activism (Section 2.5.1). I find differences in the educational composition of workers in the broader TFA sector but not in areas targeted by anti-sweatshop activism, suggesting that better labor standards did not lead to composition differences. Further, the results suggest that these reforms affected entry decisions by women into the sector. Labor force participation increased and educational attainment decreased among younger women in subdistricts with TFA firms. Taken together, these results imply that wage increases have stronger impacts on the composition of workers, and that a portion of these composition effects results from different entry decisions of workers into the sector.

2.5.3 Downstream outcomes for female workers

Before discussing the results, it should be noted that due to the nature of the data and outcomes measured, all results on downstream outcomes for female workers are suggestive. This is for three reasons. First, the data is a repeated cross-section and not a panel. As a result, I only observe outcomes for an individual at a single point in time. Second, it is plausible that any treatment effect on the outcomes measured will not be observed at the current point in time but instead in several years. This could be because outcomes evolve slowly (for example, marriage decisions are unlikely to happen immediately) or because outcomes happen at certain points in the life cycle (for example, the total number of children born is not known for certainty until women are past childbearing age). This is particularly true for outcomes related to marriage decisions and fertility. While this issue will not confound the results, it makes it more difficult to detect a treatment effect. Third, I only observe an individual's current industry of employment. If the anti-sweatshop movement affected workers' decisions to exit the TFA sector, either directly through changes within the industry (because of better wages and labor standards) or indirectly through women's marriage and fertility decisions (for example, if women drop out of the labor force after having a child), then I only measure outcomes for women who have remained in the TFA sector. This issue can confound the results as I am only ablve to observe outcomes for women who decide to remain in the sector. These issues would be mitigated if I observed outcomes in a panel data setting or if I observed women's employment histories. However, due to data limitations, I am unable to to measure outcomes several years later or account for movements into and out of the TFA sector.

Despite these caveats, estimating effects on these outcomes can still be informative, particularly for outcomes that may be affected in the short-term and for women who did not exit the sector. I examine effects on women's marriage and fertility outcomes, as these may be related to the higher household income and potential empowerment of women within the household as a result of improved wages and labor standards. These outcomes are also directly related to specific labor standards targeting maternity policies, such as forbidding firing of pregnant workers and providing childcare.

Marriage outcomes.– To begin the analysis on downstream outcomes, I examine women's marriage decisions and behavior in marriage and I find some, but limited, effects (Table 2.12): the anti-sweatshop movement reduced the marriage rate among female TFA workers (Columns 1-2), but did not affect the age at which they married (Columns 3-4) nor the rates of contraception use (Columns 5-6). Specifically, in 1998-2000 compared to 1980, the fraction of female workers ever married was 4.9 percentage points lower among TFA workers compared to other manufacturing workers. While not statistically significant, the effect in

Nike, Adidas, and Reebok subdistricts is comparable in magnitude: 4.2 percentage points lower. Despite marriage rates being lower among female workers, I do not find effects on the age at first marriage. The magnitude of the effects are small (0.053 years younger, when comparing TFA to all manufacturing, and 0.124 years older, when comparing Nike, Adidas, and Reebok subdistricts to other subdistricts). It is unlikely that this is due to marriage not being observed yet for many women, as the average age at marriage in the sample is around 18 years and 77.4% of women in TFA are over 18 years old. However, if women exit the TFA sector or the labor force when married, this could bias the treatment effect towards zero.

Because data on contraception use is only collected for women who are married, I observe one aspect of women's empowerment in marriage: contraception use. I do not find evidence that the anti-sweatshop movement had an impact on the rate at which women ever use contraception – the magnitudes are close to zero and statistically insignificant.

Treatment effects in the district-level analysis present contradictory results for marriage rates and age at first marriage (Table B.9). There is a strong positive effect on the marriage rate in Nike, Adidas, and Reebok districts (Column 2) and a strong positive effect on the age at first marriage among TFA workers (Column 3). However, there is evidence of statistically significant pre-trends (Figure B.10) which suggests that the district-level effects may be confounded by differential trends at the district level. Reassuringly, the subdistrict-level analysis is robust to replacing the year fixed effects with province-by-year fixed effects, which indicates that the results are not confounded by region-specific shocks (Table B.11).

Fertility outcomes.– I next test whether the anti-sweatshop movement had an effect on women's fertility outcomes, as fertility may be directly impacted by better labor standards or indirectly related to marriage decisions, and children's welfare, which may result from women's increased wages and improved working conditions. These results should be taken as suggestive, as the issues highlighted at the beginning of this section are particularly salient for these outcomes. With this caveat in mind, I find some effects on fertility but no effects on children's outcomes (Table 2.13).

Table 2.12 Estimates on downstream marriage outcomes for female textile, footwear, apparel workers

	Dependent variable:					
	Ever married		Age at marriage		Contraception use	
	(1)	(2)	(3)	(4)	(5)	(6)
<u>*Panel A*</u> *Treatment variable: Employed in textile, footwear, apparel*						
TREAT × POST92	-0.049^{**}		-0.053		-0.001	
	(0.024)		(0.230)		(0.031)	
<u>*Panel B*</u> *Treatment variable: Subdistrict with a Nike, Adidas, Reebok factory*						
TREAT × POST92		-0.042		0.124		0.004
		(0.048)		(0.480)		(0.056)
Observations	25,690	11,565	13,782	6,002	10,464	4,610
Dep. var. mean	0.548	0.529	18.430	18.548	0.426	0.441
Pre-period	1980	1980	1980	1980	1980	1980
Post-period	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Ever married is an indicator equal to 1 if the marital status is married, widowed, or divorced. *Age at marriage* is the individual's age at her first marriage. *Contraception use* is an indicator equal to 1 if the individual ever used a contraceptive method. For Panel A, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Panel B, treatment is a female textile, footwear, apparel worker residing in a subdistrict where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in subdistricts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes year fixed effects, subdistrict-industry fixed effects (Panel A), subdistrict fixed effects (Panel B), and district-specific time trends. Standard errors clustered at the subdistrict-industry (Panel A) or subdistrict (Panel B) level.

Table 2.13 Estimates on downstream fertility outcomes for female textile, footwear, apparel workers

	Dependent variable:							
	Ever had children		Age at first child		Number children born		Fraction children in school	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>*Panel A*</u> *Treatment variable: Employed in textile, footwear, apparel*								
TREAT × POST92	-0.058^{**}		-0.614^{*}		-0.162^{*}		-0.038	
	(0.024)		(0.338)		(0.095)		(0.034)	
<u>Panel B</u> *Treatment variable: Subdistrict with a Nike, Adidas, Reebok factory*								
TREAT × POST92		-0.054		0.211		0.011		-0.029
		(0.054)		(0.755)		(0.215)		(0.076)
Observations	25,690	11,565	9,869	4,280	25,690	11,565	6,292	2,743
Dep. var. mean	0.442	0.422	22.280	22.215	1.466	1.327	0.584	0.593
Pre-period	1980	1980	1980	1980	1980	1980	1980	1980
Post-period	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Ever had children is an indicator equal to 1 if the individual ever had a child. *Age at first child* is the individual's age at which she had her first child. *Number children born* is the number of children ever born. *Fraction children in school* is the fraction of the individual's children between 5-18 years who are currently attending school. For Panel A, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Panel B, treatment is a female textile, footwear, apparel worker residing in a subdistrict where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in subdistricts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes year fixed effects, subdistrict-industry fixed effects (Panel A), subdistrict fixed effects (Panel B), and district-specific time trends. Standard errors clustered at the subdistrict-industry (Panel A) or subdistrict (Panel B) level.

Fewer women report ever having children both in the broader TFA sector (Column 1) and in Nike, Adidas, and Reebok subdistricts (Column 2). Specifically, between 1998-2000 compared to 1980, the fraction of women who ever had children is 5.8 percentage points lower in the TFA sector compared to the rest of manufacturing and 5.4 percentage points lower in Nike, Adidas, and Reebok subdistricts compared to other subdistricts, although the latter effect is not statistically significant. Correspondingly, women have 0.162 fewer children in the TFA sector compared to the rest of manufacturing (Column 5). On the other hand, women in the broader TFA sector are *younger* when having their first child (Column 3). However, this effect is likely spurious, as there is no comparable effect among Nike, Adidas, and Reebok subdistricts (Column 4) nor in the district-level analysis (Table B.12).

While I do not find evidence of female empowerment along contraception use (Section 2.5.3), I test whether the anti-sweatshop movement affected empowerment along other dimensions as increased wages may lead to improved outcomes for children in the household. However, I do not find statistically significant effects on school attendance rates among the children of female TFA workers (Columns 7 and 8). In particular, the direction of the effects suggests that outcomes *worsened* for children of female TFA workers: a lower fraction of women's children between 5-18 years are in school. Because the estimates are not close to zero, I cannot rule out a negative effect on children's outcomes. While this result is suggestive, it is puzzling and warrants further investigation.

The district-level analysis presents contradictory results for the fraction of women ever having children and the number of children born (Table B.12). There is a strong positive effect on the fraction of women ever having children (Column 2) and a strong positive effect on the number of children born (Column 3) in Nike, Adidas, and Reebok districts. However, there is evidence of statistically significant pre-trends (Figure B.13) which suggests that the effects estimated in the district-level analysis may be confounded by differential trends at the district level. Reassuringly, the subdistrict-level analysis is robust to replacing the year fixed effects with province-by-year fixed effects, which indicates that the results are not confounded by region-specific shocks (Table B.14).

2.6 Conclusion

As more firms relocate factories to developing countries to take advantage of lower production costs, this brings a central question regarding the labor standards and welfare of workers in these factories. This is particularly important given that both public (governments in developing country) and private (multinational firms and consumers) actors may be limited in their ability to enforce labor standards. This paper studies the anti-sweatshop movement in the 1990s which targeted the labor practices of companies such as Nike, Adidas, and Reebok. I quantify the effect of the anti-sweatshop movement on the textile, footwear, and apparel industry in Indonesia using a difference-in-differences strategy, by comparing differences among firms and workers before and after the anti-sweatshop movement in either targeted sectors or in locations with targeted firms. I test whether the anti-sweatshop movement improved labor standards in TFA factories in Indonesia and whether this led to improved outcomes for female workers.

I find that labor standards improved in factories in Indonesia due to the anti-sweatshop movement along multiple dimensions. First, wage growth was higher in affected sectors and locations. Second, the fraction of underage workers and the number of hours worked were reduced in affected locations. These findings suggest that international activism and reputational concerns can induce multinational firms to effectively enforce better labor standards in overseas factories.

The implementation of better labor standards had impacts on female workers in the industry. First, the anti-sweatshop movement affected entry and exit decisions in the TFA sector and in the labor force more broadly. Female labor force participation increased among younger women in geographic areas more affected by the movement. I also find that educational attainment increased among female TFA workers, but that it decreased among younger cohorts of TFA workers. I provide suggestive evidence that entry movements were influenced more by higher wages and less by better working conditions. Second, I find some, but limited, effects on women's marriage and fertility decisions.

While the anti-sweatshop movement was effective in improving labor standards, these improvements resulted in modest gains in female autonomy and empowerment. This leads to a broader question of whether improving labor standards can lead to better economic development and gender equality. One possibility is that longer term studies are needed to fully rule out that the anti-sweatshop movement did not lead to better outcomes for female workers. A second possibility is that the standards that were the focus of the movement – notably, wages, work hours, underage employment, and workplace safety – are not effective in improving female outcomes. Labor standards more directly targeted to improving women's outcomes, such as maternity leave, childcare, and education programs, may be more effective. Future studies disentangling the impacts of different labor standards can help shed light on which policies are most effective at improving female welfare.

Chapter 3

Uncovering Household Self-Targeting with Machine Learning

3.1 Introduction

Accurately targeting households for social welfare programs is a constant challenge for governments. In developing countries, where the poor are often concentrated in the informal sector, this challenge takes on increased significance as governments must deal with issues surrounding costly income verification and agency problems among those who determine eligibility (Niehaus and Sukhtankar, 2013; Olken and Pande, 2012; Reinikka and Svensson, 2004). This paper studies both of these issues in the context of a proxy-means testing system in Colombia, SISBÉN. In particular, this paper documents how households may self-target by examining household behavior before and after the public release of the mechanism determining beneficiary status.

To quantify agency issues, I examine patterns in household eligibility status and find that households manipulated their eligibility status at higher rates after the public release of the targeting mechanism. To quantify targeting accuracy, I train a machine learning algorithm on an external data set to predict household poverty and compare its performance with that of the government's poverty prediction. I find that the government predicts household poverty poorly and predicts particularly poorly for households far below the poverty line. Taken together, these results suggest that households face high incentives to manipulate their eligibility and I find that, after the public release of the targeting mechanism, households self-targeted by manipulating their eligibility.

Between 1990-2000, government social spending in Colombia increased from 6.8% to

11.7% of GDP, reaching a peak of 14.5% of GDP in 1996 and far outpacing the growth in the rest of Latin America (Figure 3.1 and Huber and Stephens, 2012). As more countries turn to social programs to decrease poverty, designing targeting systems to allocate social spending in contexts where income is not readily observable and the government has limited capacity to monitor distribution is increasingly important. Countries grappling with these issues often turn to proxy-means testing systems (PMT) to identify beneficiaries. In a PMT, household information, such as asset ownership and demographics, is collected and used to predict a household's poverty status and determine eligibility. However, these systems may be subject to manipulation, as households may engage in lying or fraud to become eligible, with or without the cooperation of government officials. This manipulation can increase or decrease targeting accuracy, depending on whether households who qualify through manipulation are more or less likely to be eligible.

While many PMTs use sophisticated econometric techniques, there is far from consensus on their prediction quality (Coady et al., 2004b; Kidd and Wylde, 2011). PMTs rely on techniques that (i) are un-regularized, which contributes to over-fitting of the data, and (ii) do not identify interactions in the data, which leads to model mis-specification. These issues can result in poor out-of-sample predictions and have motivated growing interest in using machine learning methods to predict household poverty in order to reduce errors in targeting (Blumenstock, 2016; Mullainathan and Spiess, 2017).[47]

In 1994, the Colombian government launched SISBÉN I, a national system to identify beneficiaries for social programs administered by the government. SISBÉN is a PMT that provides targeting information through an index called the poverty score. In 1997, the Colombian government disseminated the algorithm for the poverty score to local officials. Because the benefits of being classified as eligible in SISBÉN are high, households faced high incentives to manipulate their way into eligibility. This paper (i) studies the implementation

[47] While outside the scope of this paper, there is also an issue surrounding manipulation and difficulty of verifying household survey answers. This has motivated work on using information collected from external sources, such as satellite images or cell phone data, to predict poverty and income (Blumenstock et al., 2015; Glaeser et al., 2018).

Figure 3.1 Government social spending in Latin America, 1990-2000

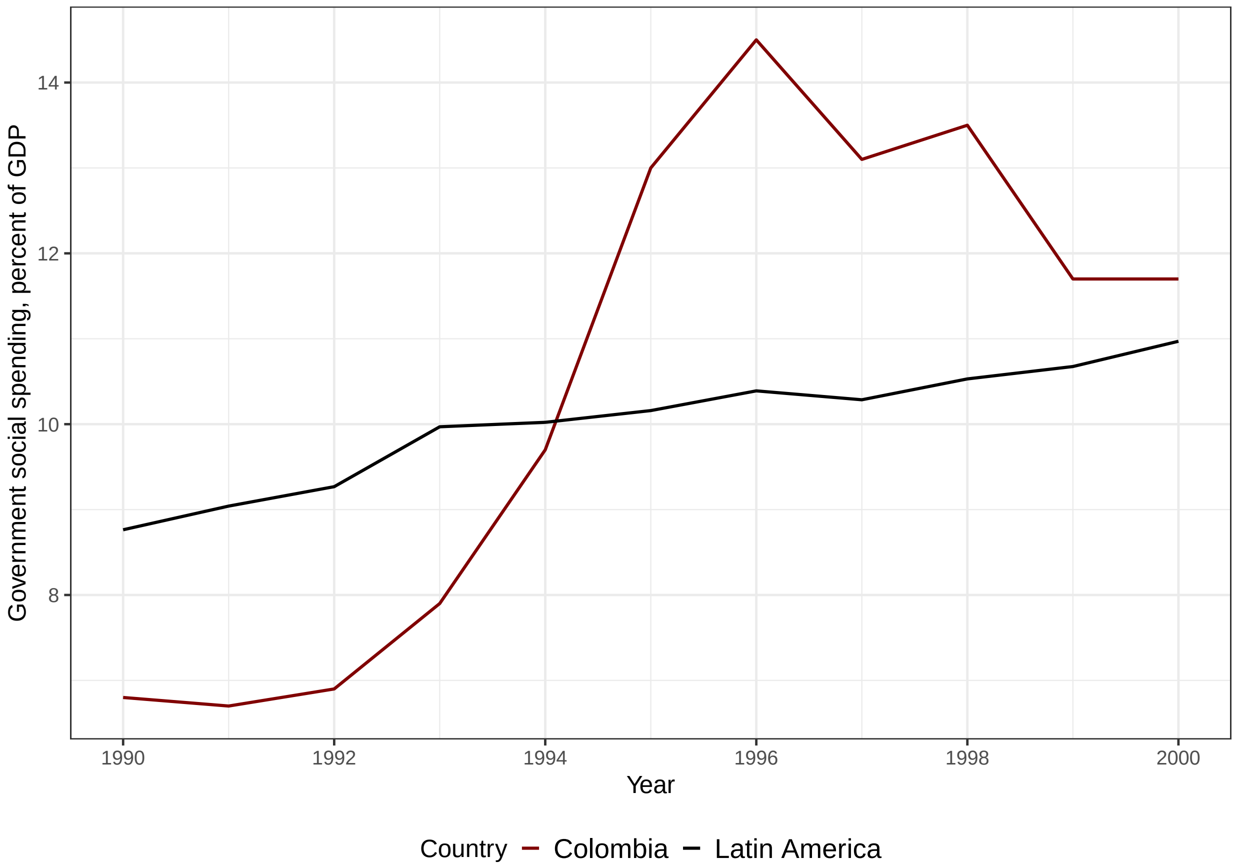

Figure plots government social spending as a percent of GDP. Social spending defined as government expenditure on education, health, social security, and welfare. Series for *Latin America* includes 23 countries in the Latin America and Caribbean region. Data from Huber and Stephens (2012).

of SISBÉN and the manipulation behavior of households before and after the algorithm for the poverty score was made public and (ii) evaluates the prediction quality of SISBÉN by applying machine learning methods to predict household poverty. I use the machine learning predictions to examine whether households self-target by manipulating their poverty scores.

I first provide evidence of manipulation of poverty scores in SISBÉN. Manipulation occurred both before and after the release of the poverty score algorithm, but the release of the algorithm affected the type of manipulation households engaged in. First, I find a large discontinuity in the density of poverty scores right at the eligibility threshold. This discontinuity only appears after the release of the poverty score algorithm in 1997, corroborating the results found by Camacho and Conover (2011). The magnitude of the discontinuity is large: a 2.08 percentage point difference in the density of households at the threshold. Second, I document that one method which households use to manipulate their eligibility involves directly altering their poverty score. Nearly 300,000 households (5.7% of households in SISBÉN) engage in this type of manipulation. The pattern of score changes indicates that these discrepancies resulted from non-random manipulation of poverty scores. However, I find that the fraction of households altering their poverty scores is higher *before* the release of the score algorithm, suggesting that households switched from altering their poverty scores to other methods of manipulation once the algorithm was released.

Second, the government poverty score predicts household poverty poorly compared to machine learning predictions. I use external household surveys, the Quality of Life surveys, to train machine learning algorithms to predict the probability of a household being below the poverty line. Prediction quality, as measured by log-loss, is significantly worse for the poverty score. When examining prediction performance at different poverty levels, the poverty score performs poorly relative to the machine learning predictions and performs particularly poorly for the households it is meant to target: the gap in performance between the poverty score and machine learning predictions is highest for households far below the poverty line. Further, this leads to high levels of targeting errors: on average, the poverty score only considers households at least 100,000 pesos (58.1% of the minimum wage in 1997) below the poverty line as eligible for subsidies.

I examine the outputs of the machine learning algorithms to better understand why the poverty score predicts poorly. First, the poverty score places weight on housing characteristics, such as housing material and access to utilities, despite there being little variation in these variables between households. Second, relative to the poverty score, the machine learning algorithms place much higher weight on characteristics of the household head and educational attainment. The poverty score also uses several variables that are non-linear transformations of the raw variables, but the machine learning algorithms did not select these transformed variables as predictive. Third, the machine learning predictions placed high weight on algorithms that are local predictors, indicating that interactions in the data are important for predicting household poverty. Taken together, this suggests that the frequently cited criticisms of PMTs – lack of regularization, model selection, and inability to detect interactions in the data – contribute to the poor performance of the poverty score.

Finally, I provide suggestive evidence that households self-target by manipulating their poverty scores, but only when they have access to information on the targeting rule. I quantify self-targeting by measuring the predicted poverty levels of households in SISBÉN and test whether households who alter their poverty scores are more likely to be eligible. Prior to the release of the poverty score algorithm, households with altered poverty scores are $5.9-7.4$ percentage points *less* likely to be predicted as poor than all other households, suggesting *negative* self-targeting. After the release of the poverty score algorithm, households with altered poverty scores are $0.7-3.6$ percentage points *more* likely to be predicted as poor, suggesting *positive* self-targeting. I do not find evidence that this is explained by a higher rate of survey answer manipulation among households who alter their poverty scores.

This paper is directly related to a literature documenting manipulation in government social programs. The literature has typically focused on corruption by politicians or bureaucrats who misallocate or redirect resources (Niehaus and Sukhtankar, 2013; Olken, 2006; Reinikka and Svensson, 2004) and less on household manipulation. Camacho and Conover (2011) study manipulation by households and do so in the context of SISBÉN as well. While Camacho and Conover (2011) document household manipulation, they focus on the political determinants of household manipulation. This paper contributes to this literature by both documenting household manipulation and evaluating whether this manipulation improved targeting accuracy.

More broadly, this paper connects to a large literature debating the design of government targeting systems (Alatas et al., 2012; Coady et al., 2004a,b; Klasen and Lange, 2016). While far less widely used than PMTs, the idea of self-targeting mechanisms has received discussion (Alatas et al., 2016; Besley and Coate, 1992; Heckman and Smith, 2004; Kleven and Kopczuk, 2011). Self-targeting typically takes the form of imposing requirements for program eligibility that are less costly for poor households, such as manual labor requirements (as in NREGA in India) or administrative complexity. This paper's contribution is to evaluate the targeting accuracy of a PMT relative to machine learning predictions, and to examine whether households can self-target through illicit methods.

This paper is organized as follows. Section 3.2 describes the SISBÉN program in Colombia. Section 3.3 presents the empirical strategy to quantify manipulation and the procedure to train the machine learning predictions. Section 3.4 presents the results on household manipulation of poverty scores. Section 3.5 examines the prediction quality of the machine learning predictions. Section 3.6 discusses the implications for targeting accuracy. Section 3.7 concludes.

3.2 Context

SISBÉN (*El Sistema de Selección de Beneficiarios para Programas Sociales*) is the national system of identification of social program beneficiaries in Colombia. This paper studies the first iteration of the program, SISBÉN I, which was implemented between 1994 and 2003.[48]

SISBÉN is a registry of over 25 million individuals and is used by a variety of government social programs to determine eligibility. The national subsidized health system is the largest social program relying on SISBÉN, but it is also used by programs administering benefits such as scholarships, housing subsidies, and conditional cash transfers. SISBÉN is a PMT providing targeting information through an index called the poverty score that summarizes a household's level of economic well-being. A household's poverty score is calculated through an algorithm that weights different household socioeconomic variables collected through a household survey. Officials used an initial survey of 25,000 households to both (i) select the subset of variables to include in the algorithm and (ii) determine the weights for each variable,

[48] As of writing, the program is in its fourth iteration, SISBÉN IV.

through principal component analysis. Weights were developed separately for urban and rural households. The result is a poverty score that ranges from 0 (most poor) to 100 (least poor).[49] Poverty scores are then used by social programs to identify beneficiaries, with most social programs employing a threshold of 47 for eligibility.[50] Thresholds were determined using probit estimates to maximize the likelihood of correctly identifying households below the poverty line.

While SISBÉN is a national program, SISBÉN was initiated in the context of a decentralization effort that began in 1991 and was designed to be primarily used and implemented by local governments. The survey instruments (*Ficha de Caracterización Socioeconómica*) and computer program to calculate the poverty score are developed by the National Planning Department, *Departamento Nacional de Planeación*, in the central government. Municipalities are then responsible for the logistical administration of the surveys – which includes hiring enumerators and determining the geographic areas targeted and the timing of surveys – and maintenance of the database. Generally, surveys are geographically targeted to poorer areas using the neighborhood strata system, *Estratificación Socioeconómica*. The neighborhood strata are local assessments of the exterior characteristics of neighborhoods and dwellings, ranging from 1 (most poor) to 6 (least poor). Households in targeted neighborhoods, typically those in strata three and below, are surveyed based on a census-style approach, where all households in a targeted neighborhood are interviewed. This represents the majority of households in the database. Households in non-targeted neighborhoods can request an interview.

I study the public release of the poverty score algorithm in late 1997. The algorithm was released to municipal administrators through an instructional presentation and pamphlet. Because the algorithm is fairly complex, it is unlikely that households became aware of the full calculation procedure, although it is possible that some households did. The release

[49] For a detailed description of the development of the poverty score algorithm, see Castañeda (2005) and Vélez et al. (1998).

[50] This is the threshold for urban households. The most common threshold is 30 for rural households.

of the algorithm appears to have primarily resulted in making households aware that the purpose of the surveys was to assign a household poverty score and that a score below a certain threshold determined eligibility for social programs.

3.3 Empirical strategy

3.3.1 Government poverty score

A household's poverty score is calculated from four subcomponents: housing, utilities, education, and demographics. The housing and utilities subcomponents are calculated from household-level variables, such as building characteristics and access to sanitation. The education and demographics subcomponents are calculated from individual-level variables, such as household members' educational attainment or working status. Table C.1 provides a summary of the variables used in each subcomponent. The SISBÉN survey contains 62 questions, although I am only able to access 26 variables. 19 of these variables are used to calculate the poverty score.

While some variables are used as-is to calculate the poverty score (for example, roofing material is directly used to calculate the score), other variables are first transformed (for example, individuals' educational attainment is used to first calculate the average years of education of household members older than 12 years, which then enters the score calculation). As a result, it is unlikely that households could predict their poverty score without some knowledge of the algorithm.

I use the administrative SISBÉN survey provided by Camacho and Conover (2011). The data comprises the sample of urban households living in neighborhood strata level three and below, where the majority of surveys were targeted, who were surveyed between 1994 and 2003.[51] This sample comprises 18,176,019 individuals and 5,245,018 households across 785

[51]The data includes households surveyed between 1993 and 2006. 1993 is excluded as it was a pilot year. 2004-2006 are excluded as the Colombian government had begun phasing out SISBÉN I for SISBÉN II.

municipalities.

Score reconstruction.– To identify households who alter their poverty scores, I reconstruct a household's poverty score. I do this in two ways. One, using the poverty score variables, I re-calculate a household's poverty score using the government's algorithm. Two, I re-calculate a household's poverty score using a household's subcomponent scores. In the cases that a household's poverty score has been altered, this allows me to pinpoint whether altering the score occurred when calculating the subcomponents or when adding the subcomponents.

Estimating sorting.– To formally estimate the size of the discontinuity in the density of poverty scores around the eligibility threshold, I use a McCrary test (McCrary, 2008). The McCrary test is typically used in regression discontinuity designs to test for sorting of the running variable by estimating the size of a discontinuity in the density of the running variable around the threshold. However, an issue arises when the running variable is discrete, as in this case with the poverty scores in SISBÉN.[52] Because there are many observations that are exactly at the threshold, these tests may be biased. To alleviate this bias, I follow Eggers et al. (2015) and set the threshold to 47.5 (where the actual threshold is 47, as all families with a score of 47 and below are considered eligible) and check that the bin size is not exactly a multiple of 0.5.

3.3.2 Using machine learning to predict poverty

The poverty score is the government's prediction of the probability that a household's consumption level lies below the poverty line.[53] Because SISBÉN does not contain measures

[52]The calculated poverty score is not discrete, as it relies on the addition of subcomponents with decimal values. However, the assigned poverty score is discrete as it is the floor of the calculated poverty score.

[53]Actually, it is the probability that a household's consumption level is *above* the poverty line, since the score ranges from 0, most poor, to 100, least poor. To compare the poverty score to the machine learning predictions, which predict the probability that a household's consumption level is *below* the poverty line, I transform the government poverty score, $POVSCORE_g$, as follows:

$$POVSCORE_a = (1 - POVSCORE_g/100)$$

of household consumption levels, I use the 1993, 1997, and 2003 Quality of Life (*Encuesta Nacional de Calidad de Vida*) surveys, henceforth QOL, to train machine learning algorithms to predict household poverty levels in order to i) evaluate the performance of the government's poverty score and ii) obtain poverty predictions for households in the SISBÉN survey. The QOL surveys are nationally representative surveys carried out by Colombia's national statistical agency, *Departamento Administrativo Nacional de Estadística.*

Constructing the training data.– In order to train machine learning algorithms to predict poverty, I construct a training data set as follows. Since I plan to apply the machine learning predictions, which are trained on the QOL data, to households in SISBÉN, I identify variables common between the SISBÉN survey and QOL surveys. Because the SISBÉN sample is restricted to households in urban areas, I restrict the QOL sample to households in urban areas. While the SISBÉN sample is restricted to households living in neighborhood strata level three and below, I do not restrict the QOL sample to these strata. I follow the urban-rural restriction as it is plausible that the true prediction function is different for households living in urban areas versus rural areas (for example, access to sanitation will predict poverty differently for households in urban versus rural areas). On the other hand, it is less clear whether the true prediction function is different for households living in different neighborhood strata in urban areas. These sample restrictions result in a final sample of 44,977 households in the QOL surveys.[54]

I next split the QOL data into a training and a holdout set. Choosing the size of the holdout set entails a trade-off. A larger holdout set increases the precision of calculating the loss, or the prediction quality. On the other hand, a larger holdout set decreases the size of the training sample, potentially reducing prediction quality. Because I use 10-fold cross-validation in the training procedure, I choose the size of the holdout sample to be equal to the size of a cross-validation sample. I randomly assign households to the training

[54]The distribution across the QOL survey years is: 21,133 households from the 1993 survey, 5,367 households in the 1997 survey, and 18,477 households in the 2003 survey.

set and the holdout set, where the randomization is stratified by year and neighborhood strata. Randomization is stratified according to how households are sampled in SISBÉN. The resulting split is 40,938 households in the training set and 4,039 households in the holdout set. To evaluate the randomization, I test that households in the training and holdout set are similar along various household characteristics (Table C.2). There are small differences in some household characteristics – namely, marital status of the household head, refrigerator ownership, and floor material of the house – but the differences across all characteristics is not jointly significant, indicating that households in the training and holdout sets are mostly similar.

Table 3.2 compares household characteristics between the SISBÉN sample and the QOL training set. In general, households in SISBÉN appear poorer than households in the QOL training set, both in the full training set and among households living in neighborhood strata level three and below. For example, the household head has on average 5.26 years of schooling in SISBÉN but has 7.88 years of schooling among households strata three and below in the QOL training set. Similarly, 32.8% of households in SISBÉN own a refrigerator compared to 52.3% of households in strata three and below in the QOL training set. In particular, a higher fraction of households are considered poor in SISBÉN compared to in the QOL training set.[55] These differences may reflect the non-random targeting of SISBÉN surveys, as local mayors targeted interviews to areas with high poverty, or a higher rate of household manipulation of survey answers in SISBÉN. Manipulation of survey answers is discussed in more detail in Section 3.3.2.

Because the QOL surveys are at the individual level and the poverty measures (ie. whether a household is below the poverty line) are at the household level, I collapse the QOL surveys down to the household level. Doing so generates new variables, such as the number of children 0-4 years, the number of employed household members, or the educational attainment of the highest earner in the household. All categorical variables are transformed into indicator

[55]This is defined as having a poverty score below 47 in SISBÉN or having a household consumption level below the national poverty line in QOL.

Table 3.2 Household characteristics in SISBÉN and Quality of Life training set

	SISBÉN (1)	Full training set (2)	Strata 3 and lower (3)
Number of households	5,245,018	40,938	34,222
Below the poverty line	0.607	0.561	0.632
Characteristics of household head			
Female	0.360	0.339	0.332
Age	39.947	42.391	41.684
Years of schooling	5.261	8.582	7.883
Married or in union	0.537	0.623	0.629
Characteristics of household			
Household size	3.422	4.053	4.123
Number of children under 4 years	0.438	0.391	0.421
Total rooms	1.903	3.587	3.392
Owns refrigerator	0.328	0.513	0.523
Wall material: mud, adobe, brick	0.910	0.976	0.972
Floor material: cement, tile, rug, marble	0.840	0.928	0.922
Water connected	0.903	0.970	0.967
Toilet connected	0.772	0.921	0.909
Electric lighting	0.980	0.996	0.996

Summary characteristics of households in SISBÉN (Column 1) and in the Quality of Life surveys (Columns 2-3). For households in the Quality of Life surveys, sample is either the full training set (Column 2) or households in the training set in neighborhood strata level 3 or below (Column 3). *Below poverty line* is defined as having a poverty score below 47 for households in SISBÉN and is defined as having a consumption level below the national poverty line for households in the Quality of Life surveys. *Water connected* is an indicator equal to 1 if a household's main water source is piped water. *Toilet connected* is an indicator equal to 1 if a household's toilet is connected to the sewage system. *Electric lighting* is an indicator equal to 1 if a household's main source of lighting is electricity.

variables. After dropping features with zero variance and highly correlated predictors, the resulting data set contains 154 features.

Training procedure.– I tune a random forest, boosted tree, and elastic-net regression, then estimate an ensemble of the predictions from the three algorithms to generate the machine learning poverty predictions. Because the poverty score is the probability that a household is below the poverty line, I use a log-loss function to tune the machine learning algorithm parameters, which is fairly standard when predicting probabilities. I use 10-fold cross-validation on the training set to tune the random forest, boosted tree, and elastic-net regression, which is also fairly standard in the literature. Cross-validation folds are stratified on survey year and neighborhood strata. Below, I discuss the tuning procedure for the random forest, boosted tree, elastic-net regression, and ensemble in more detail.

In general, I chose standard tuning parameters for the random forest, boosted tree, and elastic-net regression. Values for the tuning parameters across all three algorithms were chosen to minimize the cross-validated loss. The tuning parameters for the random forest were the number of trees and the number of variables to randomly sample at each split in the tree. These tuning parameters aim to reduce the correlation between trees and reduce variance in the predictions. The tuning parameters for the boosted tree were the number of trees or boosting iterations, the shrinkage parameter, and the depth of each tree. I did not tune the minimum number of observations in each leaf, but imposed a minimum of ten observations. These tuning parameters aim to reduce overfitting and bias. The tuning parameters for the elastic-net regression were alpha and lambda. For alpha, I allowed the algorithm to choose between two values: 0 (a ridge regression) and 1 (a lasso regression). For lambda, I chose the value within 1 standard deviation of the value that minimized the cross-validated loss.

After tuning the random forest, boosted tree, and elastic-net regression, I generate the machine learning predictions using an ensemble of the three algorithms. An ensemble allows me to combine predictions from a local predictor (random forest and boosted tree) and a global predictor (elastic-net regression).

Manipulation of survey answers in SISBÉN.– Due to the high-stakes nature of the SISBÉN survey, it is likely that households attempted to manipulate their survey answers to influence their poverty scores. On the other hand, this manipulation is unlikely to be present in the QOL surveys, as these are statistical surveys and households do not receive any material benefit based on how the survey is answered. This may result in the true prediction functions for poverty being different across the two surveys. This poses a problem, as a prediction function trained on the QOL surveys may not generate accurate predictions when applied to SISBÉN. Unfortunately, SISBÉN does not contain household consumption measures so I cannot use SISBÉN to train a prediction algorithm. While these issues can be alleviated if SISBÉN could be linked to other surveys with less biased measures of poverty or with access to the full SISBÉN survey (which contained survey answers not used in the calculation of the poverty score), this is a limitation of the current study.

To address this issue, I train a second prediction algorithm using a limited subset of the covariates in SISBÉN. The limited covariate list excludes variables that are difficult to verify or can be easily manipulated, such as small household assets and employment variables. The list of variables included in the full covariate list and in the limited covariate list is in Table C.1. The training set for the limited covariate list contains 102 features.

Evaluating performance.– To quantify the prediction quality of the machine learning predictions, I estimate the mean loss in the holdout set and compute conventional standard errors. Let $\hat{l}_i$ denote the loss from holdout observation i. Then the standard error is given by

$$\left(n^{-1} Var\ \hat{l}_i\right)^{1/2}.$$

Since each observation in the holdout set is randomly sampled according to the same data generating process as the full QOL sample, this will yield unbiased measures of the prediction loss. To compare the prediction loss across different predictions, I use standard errors that take into account the correlation in loss between predictions. Let $\hat{l}_{i,a}$ denote the loss from holdout observation i when using the prediction function a and $\hat{l}_{i,b}$ denote the loss from holdout observation i when using the prediction function b. Then the sample variance of the difference in losses is given by

$$\left(n^{-1}\left[Var\ \hat{l}_{i,a} + Var\ \hat{l}_{i,b} - 2Cov(\hat{l}_{i,a}, \hat{l}_{i,b})\right]\right)^{1/2}.$$

3.4 Evidence of manipulation

In this section, I provide evidence that households manipulated their poverty scores in order to gain eligibility to social programs. One, I show that a significant discontinuity in the density of scores appears at the eligibility threshold and that no discontinuity appears at other placebo thresholds. Two, I show that this discontinuity appears after 1997, when the poverty score formula was released to local officials.

Visually, there is a significant discontinuity in poverty scores assigned to households at the eligibility threshold (left panel of Figure 3.3). To test this formally, I estimate the size of the discontinuity in the density of households using a McCrary test and find a significant and large discontinuity at the eligibility threshold (Panel C of Table 3.4). The estimated size of the discontinuity is -0.446 log points, which roughly corresponds to a 1.42 percentage point (about 74,304 households) drop in the density of households at the threshold. In other words, households are manipulating their scores in order to fall below the threshold and gain eligibility to social programs.

Examining the discontinuity estimates by year indicates that this discontinuity does not appear before 1998. For households surveyed in 1994-1997, there is a *positive* discontinuity, or a higher mass of households right above the threshold (Panel A of Table 3.4). On the other hand, for households surveyed in 1998-2003, the discontinuity is *negative* and becomes increasingly negative in later years, indicating that there is a higher mass of households right below the threshold (Panel B of Table 3.4). There is a meaningful difference in the magnitude of the estimates: the discontinuity is much larger after 1997 compared to before 1997. The estimated discontinuity at the threshold is 0.124 log points (approximately 0.43 percentage points) for households surveyed between 1994-1997 and is -0.678 log points (approximately -2.08 percentage points) for households surveyed between 1998-2003. The timeline of the discontinuity estimates – switching from positive to negative between 1997 and 1998, and becoming increasingly negative through 2003 – coincides with the timing of the release of the

Figure 3.3 Density of assigned poverty scores

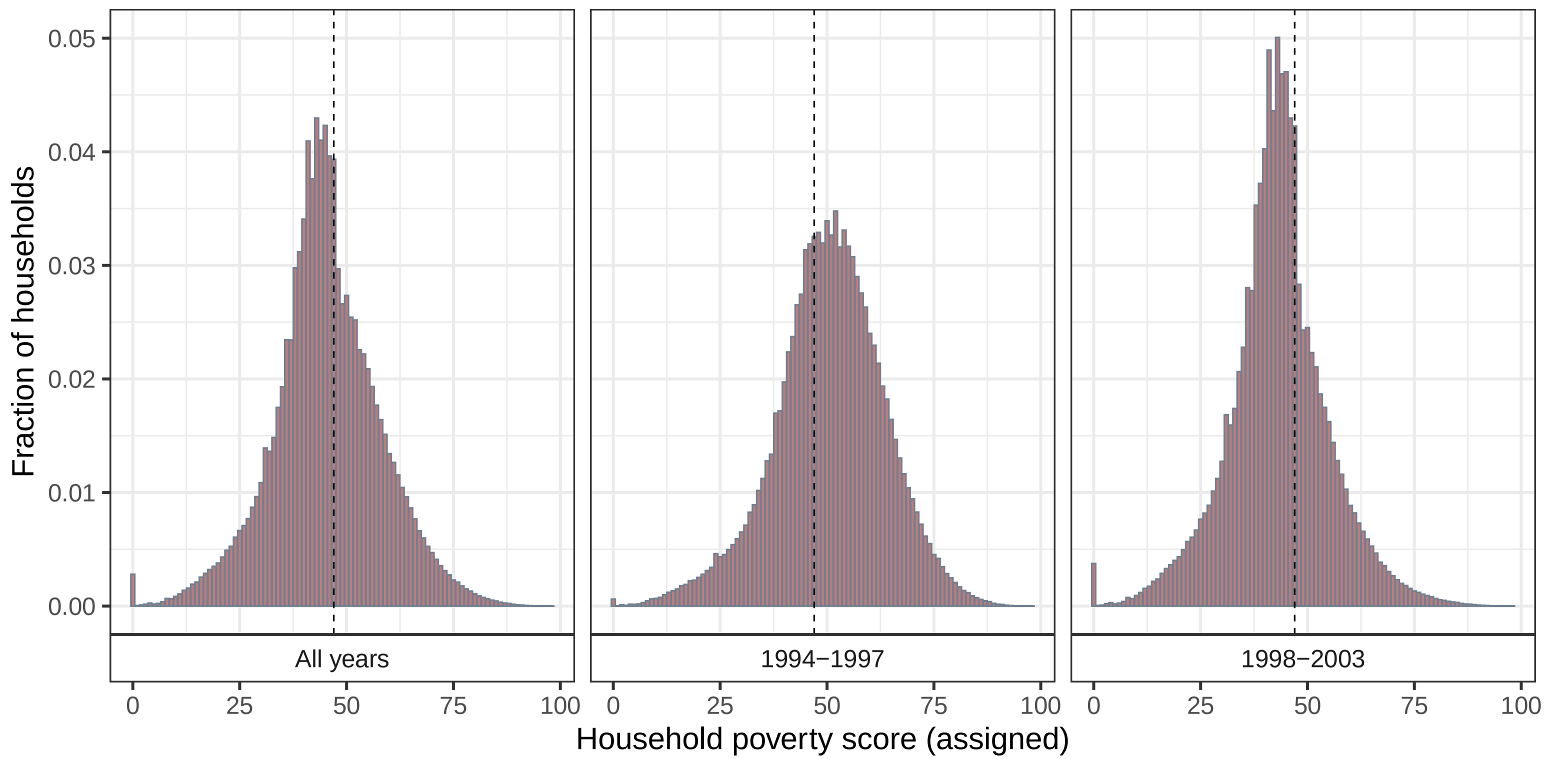

Figures plot the density of poverty scores assigned to households in SISBÉN. Distribution plotted separately for households surveyed across all years (left panel), 1994-1997 (middle panel), and 1998-2003 (right panel). The poverty score algorithm was released to local officials in late 1997. Households with a poverty score below 47 (vertical dotted line) are eligible for a variety of social programs. Data from SISBÉN survey.

Table 3.4 Estimated size of discontinuity in density of assigned poverty scores at eligibility threshold

Years	Estimate	Standard error	p-value	Bandwidth
Panel A: Households surveyed in 1994-1997				
1994	0.066^{***}	[0.008]	< 0.001	27.060
1995	0.159^{***}	[0.004]	< 0.001	23.724
1996	0.053^{***}	[0.006]	< 0.001	26.095
1997	0.123^{***}	[0.005]	< 0.001	24.713
1994-1997	0.124^{***}	[0.003]	< 0.001	24.910
Panel B: Households surveyed in 1998-2003				
1998	-0.299^{***}	[0.005]	< 0.001	29.403
1999	-0.512^{***}	[0.004]	< 0.001	24.479
2000	-0.533^{***}	[0.003]	< 0.001	29.533
2001	-0.740^{***}	[0.004]	< 0.001	24.269
2002	-0.944^{***}	[0.005]	< 0.001	19.401
2003	-0.914^{***}	[0.009]	< 0.001	18.780
1998-2003	-0.678^{***}	[0.002]	< 0.001	24.793
Panel C: Households across all years				
All	-0.446^{***}	[0.001]	< 0.001	29.374

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Estimates of the size of the discontinuity in the density of assigned poverty scores at the eligibility threshold, using a McCrary test. The discontinuity estimate is the estimated log difference in heights of the density function. The poverty score algorithm was released to local officials in late 1997. Households with a poverty score below 47 are eligible for a variety of social programs. Data from SISBÉN survey.

poverty score algorithm.

To evaluate the robustness of the 1998-2003 discontinuity estimate, I estimate the size of the discontinuity at other thresholds around the actual threshold. I do not find discontinuity estimates of comparable magnitude at other thresholds, and most estimates are close to or above zero (Figure C.3).

3.4.1 Score alteration

I document that one method which households use to manipulate their eligibility involves directly altering their poverty score. I reconstruct a household's poverty score and find that for 297,978 households (5.7% of households), the assigned poverty score does not match the reconstructed score, either in calculating the subcomponents or in adding the subcomponents. I provide three pieces of evidence that these discrepancies likely resulted from non-random manipulation and not from random calculation errors.

First, nearly all discrepancies (99.4%) results in a lower assigned poverty score to households (Figure 3.5). If these discrepancies are due to random errors, we would expect to see score changes more evenly distributed around zero. As nearly all poverty scores are altered in the direction of eligibility, the discrepancies are unlikely to be due to random calculation errors. While the difference is small, as the fraction of lowered poverty scores is so high, score alterations are significantly more likely to be negative after the release of the poverty score algorithm (0.67 percentage points, $p < 0.001$).

Two, the distribution of reconstructed poverty scores among households with altered scores, compared to the distribution among households with unaltered scores, is shifted towards higher values and is concentrated above the eligibility threshold (Figure 3.6). Both the median and distribution of reconstructed poverty scores among households with altered poverty scores are significantly different from those among households with unaltered poverty scores ($p < 0.001$ for the difference in medians and $p < 0.001$ for the difference in distributions).[56] The median reconstructed poverty score is 53 among households with an altered score while it is 45 among households with an unaltered score. If score alteration was random, it is more likely that the distribution of reconstructed scores would be similar among households with altered and unaltered poverty scores. However, households with altered scores have reconstructed scores distributed towards higher values that are above the eligibility threshold. Given this pattern, it is more likely that score alteration was in response to expectations that, absent alteration, the household would be above the eligibility threshold.

Three, households which became eligible with an altered score (in other words, a reconstructed poverty score above the threshold and an assigned poverty score below the threshold)

[56]The p-values are computed using a permutation test, where the test statistic is the difference in medians (for testing difference in medians) or the Kolmogorov-Smirnov statistic (for testing difference in distributions).

Figure 3.5 Density of changes to assigned poverty scores

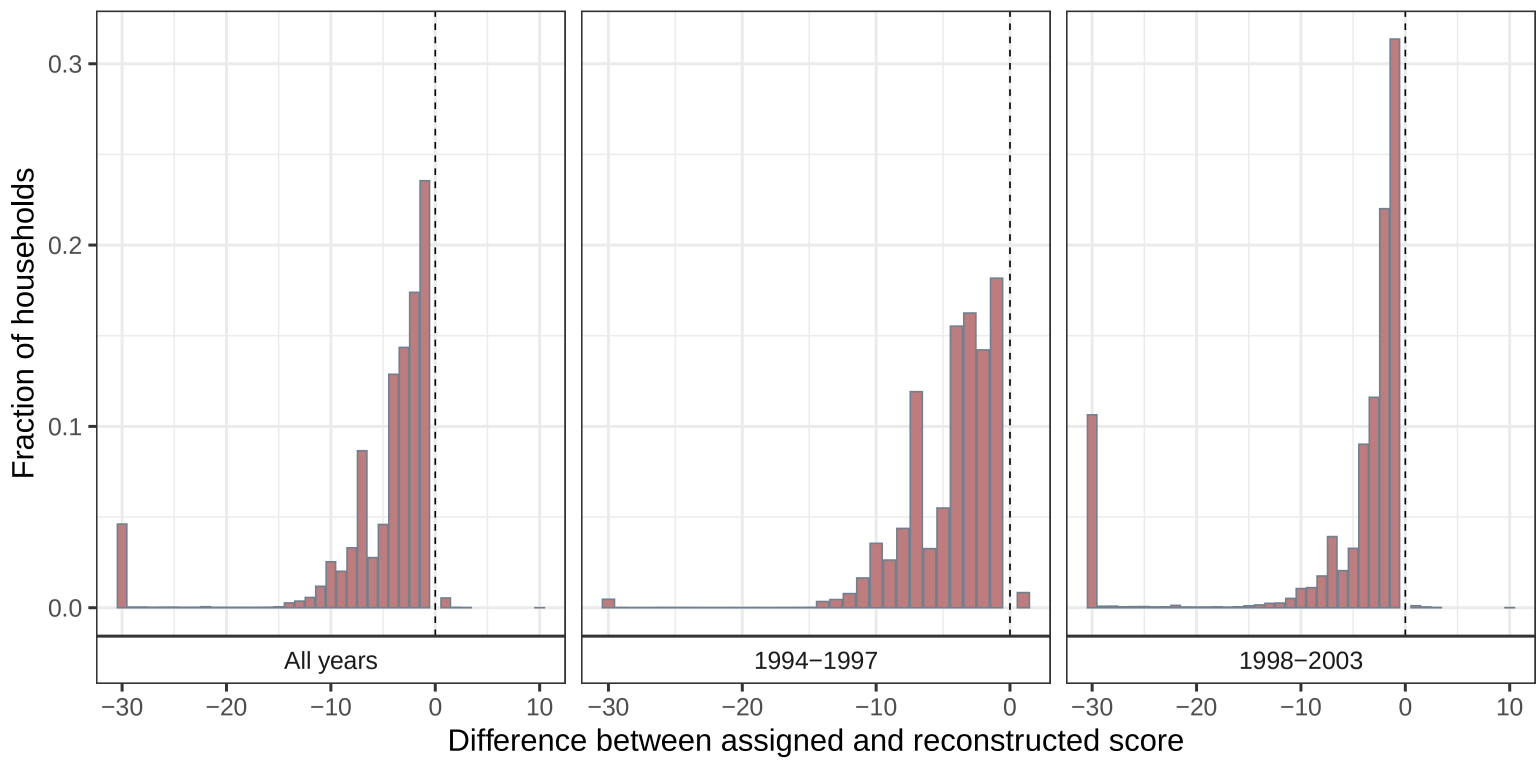

Figures plot the density of score changes among households with an altered poverty score. The score change is defined as the difference between a household's assigned and reconstructed poverty score, where the poverty score is reconstructed from the household's survey answers or the sum of the poverty score subcomponents. Distribution plotted separately for households surveyed across all years (left panel), 1994-1997 (middle panel), and 1998-2003 (right panel). The poverty score algorithm was released to local officials in late 1997. Score changes below 0 (vertical dotted line) indicate that households were assigned a lower score. For ease of viewing, a score change of -30 includes all score changes below -30. Data from SISBÉN survey.

Figure 3.6 Density of reconstructed poverty scores, among households with altered poverty scores

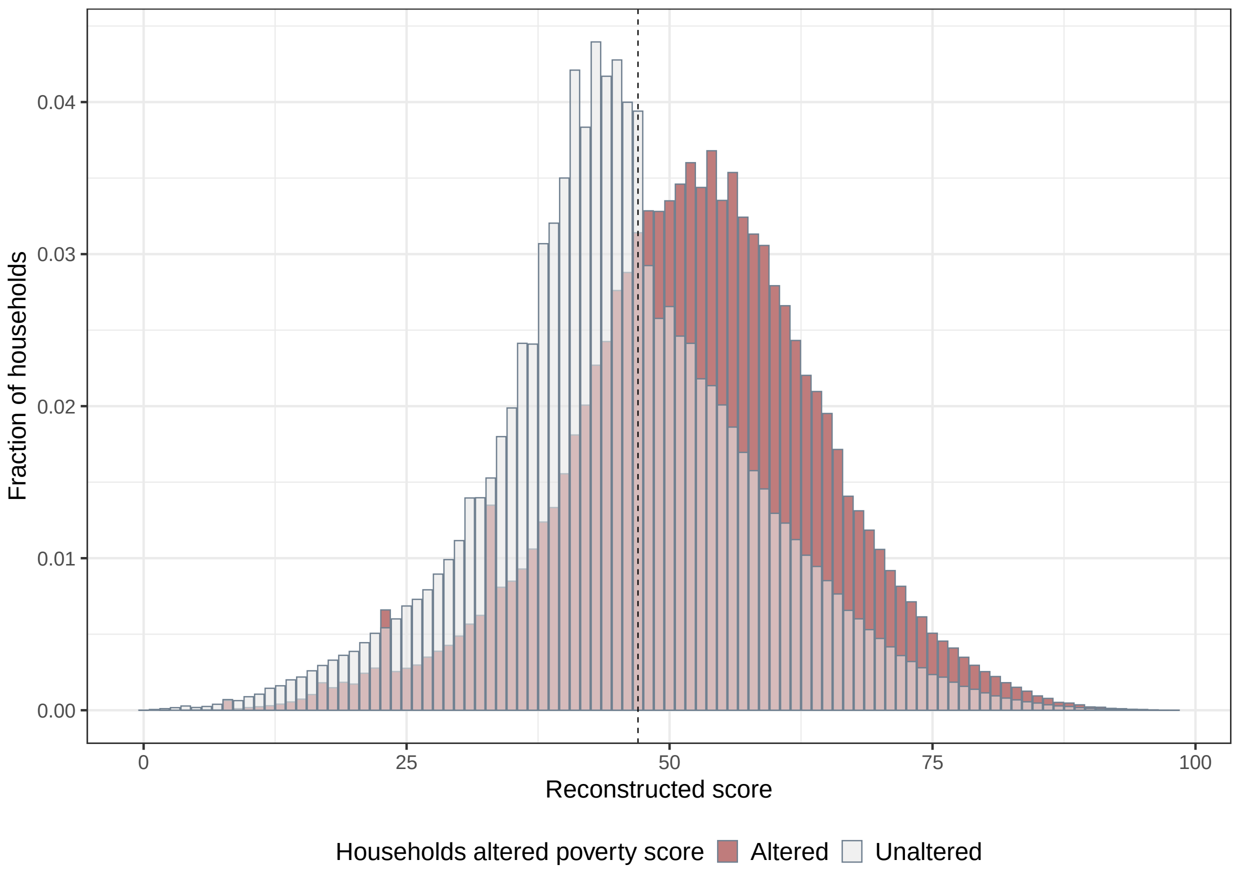

Figures plot the density of reconstructed poverty scores among households with an altered score (red bars) and an unaltered score (gray bars). The poverty score is reconstructed from the household's survey answers or the sum of the poverty score subcomponents. The poverty score algorithm was released to local officials in late 1997. Households with a poverty score below 47 (vertical dotted line) are eligible for a variety of social programs. Data from SISBÉN survey.

have reconstructed poverty scores largely concentrated *just above* the eligibility threshold (Figure 3.7). 17.3% of households have altered poverty scores that allowed them to fall below the eligibility threshold. I check whether the concentration of reconstructed poverty scores just above the threshold is a general feature and do not find that this is the case. Neither households with altered scores but whose score change did not change their eligibility status (Panel (a)) nor households with unaltered scores (Panel (b)) have reconstructed poverty scores concentrated just above the threshold. While it is unlikely that households knew whether they were exactly above or below the threshold, the concentration of reconstructed scores just above the threshold suggests that households were aware that they were close to the threshold. Given that these households have altered scores which allowed them to become eligible, this further suggests that manipulation was non-random.

Examining the pattern of score alterations, nearly all score changes (99.0%) occurred when calculating the subcomponents and not in adding the subcomponents. Interestingly, a higher fraction of score alterations occurred *prior* to the release of the poverty score algorithm. In 1994-1997, 176,524 households (11.1% of surveyed households) had altered poverty scores while in 1998-2003, 121,454 households (3.3% of surveyed households) had altered poverty scores. This suggests that, with the release of the poverty score algorithm, households shifted the manipulation of their poverty scores to less overt methods. This is discussed in more detail in the next section.

3.4.2 Other manipulation methods

While a large number of households have directly altered scores, it is unlikely that this was the only method households used to manipulate their scores. In fact, a significant discontinuity at the eligibility threshold remains even among the reconstructed scores (Figure C.4). While the estimated size of the discontinuity is smaller than that in the assigned scores, it is still sizable (Table C.5). This suggests that score alteration is not the main method which households used to manipulate their scores in order to gain eligibility.

More likely, households manipulated their scores earlier in the process by manipulating their survey answers. One possibility is that households lie to the enumerator. For example, the survey respondent may report that a household member had not worked the previous week when she in fact had. This is a common issue in PMTs, which is why PMTs often rely on collecting verifiable observable characteristics rather than relying on self-reporting. Nevertheless, households may still manipulate their survey answers in other ways. A second possibility is that households hide or otherwise manipulate observable characteristics. For example, households may hide assets such as a TV. Third, households may directly coordinate

Figure 3.7 Density of reconstructed poverty scores, among households who became eligible with an altered assigned poverty score

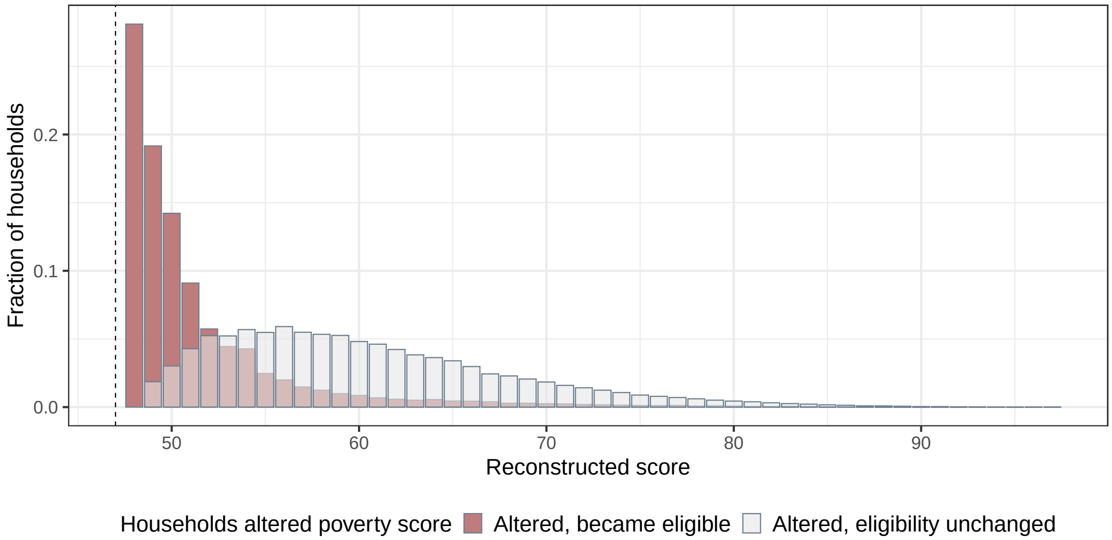

(a) Compared to households whose assigned poverty scores did not change eligibility status

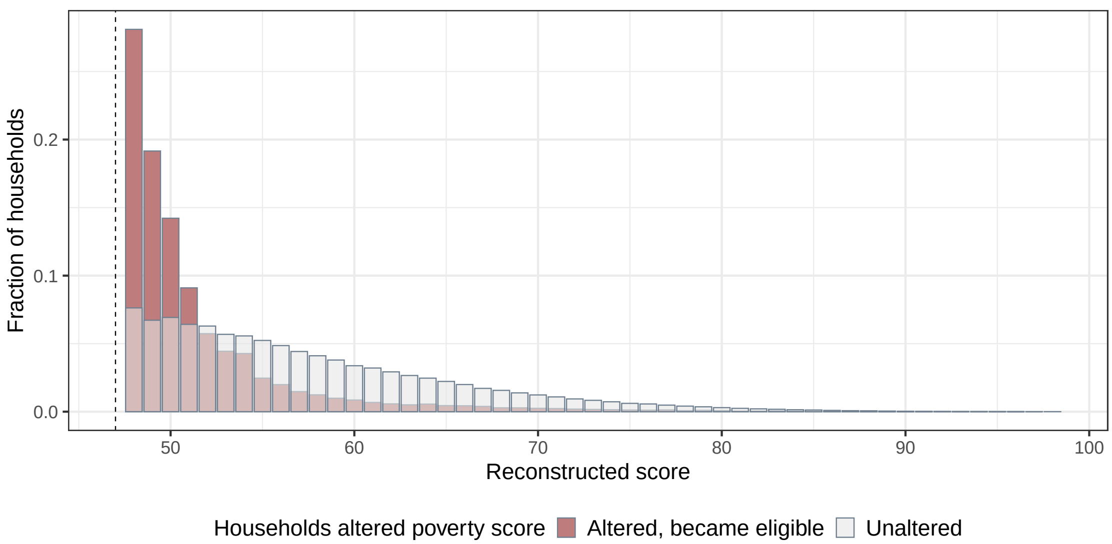

(b) Compared to households with unaltered poverty scores

Figures plot the density of reconstructed poverty scores among households with an altered poverty score that changed their eligibility status (red bars in Panels (a) and (b)), compared to households with an altered poverty score that did not change their eligibility status (gray bars in Panel (a)) and to households with an unaltered poverty score (gray bars in Panel (b)). The poverty score is reconstructed from the household's survey answers or the sum of the poverty score subcomponents. The poverty score algorithm was released to local officials in late 1997. Figures only plot households with reconstructed scores above 47, the eligibility threshold (vertical dotted line). Data from SISBÉN survey.

with enumerators or other administrative officials to alter their survey answers. For example, households may bribe the enumerators or may have been coached by enumerators on the correct answers. In this paper, I focus on score alteration as the main manipulation method. A more detailed study of different manipulation methods can be found in Camacho and Conover (2011).

3.5 Predicting poverty

In this section, I discuss the prediction performance of the machine learning predictions. In Section 3.5.1, I evaluate the performance of the machine learning predictions and the government's poverty score in predicting household poverty in the Quality of Life surveys. The poverty score performs worse than the machine learning predictions. In particular, the gap in performance is highest for the *poorest* households, indicating that the poverty score predicts poorly for the households which it is meant to accurately target. In Section 3.5.2, I discuss the outputs from tuning the machine learning algorithms and the covariates selected as most predictive. Compared to the poverty score, the machine learning algorithms (i) place greater weight on educational attainment and lower weight on household characteristics and (ii) rely more on non-linear and local predictors. This suggests that likely contributors to the poor performance of the poverty score are model selection, unregularized variable selection, and the limited ability to identify interactions in the data.

3.5.1 Prediction performance

To evaluate and compare the prediction quality of the machine learning predictions and the government poverty score, I first examine the prediction loss in the QOL surveys (Table 3.8). As expected, the in-sample loss is smaller than the out-of-sample loss for the machine learning predictions (Column 1 and 2 in Panel A), reflecting the tendency to over-fit in-sample. The machine learning predictions are more accurate than the poverty score, both when using predictions trained on the full covariate list and on the limited covariate list: mean log-loss is 0.368 when using the full covariate list, 0.384 when using the limited covariate list, and 0.639 when using the poverty score (Panel A). Differences in prediction quality between the machine learning predictions and the poverty score are significant at the 1% level (Panel B). Interestingly, while the poverty predictions using the limited covariate list perform worse than when using the full covariate list (a difference of 0.016 log points), the difference in prediction quality is small in comparison to the difference in prediction quality with the poverty score (0.271 log points, for the full covariate list, and 0.255 log points, for the limited covariate list). This suggests that the worse performance of the poverty score arises more from functional form assumptions rather than the limited covariates the government can observe.

Table 3.8 Prediction loss of machine learning predictions and government poverty score

	Training set, in-sample loss	Holdout set, out-of-sample loss
Panel A: Log-loss of prediction		
(1) Full covariate list	0.294	0.368
	(0.002)	(0.008)
(2) Limited covariate list	0.319	0.384
	(0.002)	(0.008)
(3) Poverty score	0.649	0.639
	(0.002)	(0.005)
Panel B: Difference in log-loss between prediction algorithms		
(1) - (2)		-0.016^{***}
		(0.003)
(1) - (3)		-0.271^{***}
		(0.009)
(2) - (3)		-0.255^{***}
		(0.009)

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Panel A: Prediction loss when using the machine learning predictions trained on the full covariate list (row 1), the machine learning predictions trained on the limited covariate list (row 2), and the government poverty score (row 3). Prediction loss calculated both in-sample on the training set (Column 1) and out-of-sample on the holdout set (Column 2). *Panel B:* Difference in out-of-sample prediction loss between prediction algorithms. Prediction loss is the sample mean of the log-loss. Standard errors below the prediction loss. Data from Quality of Life surveys.

Figure 3.9 Poverty predictions from the machine learning predictions and government poverty score, by actual household consumption

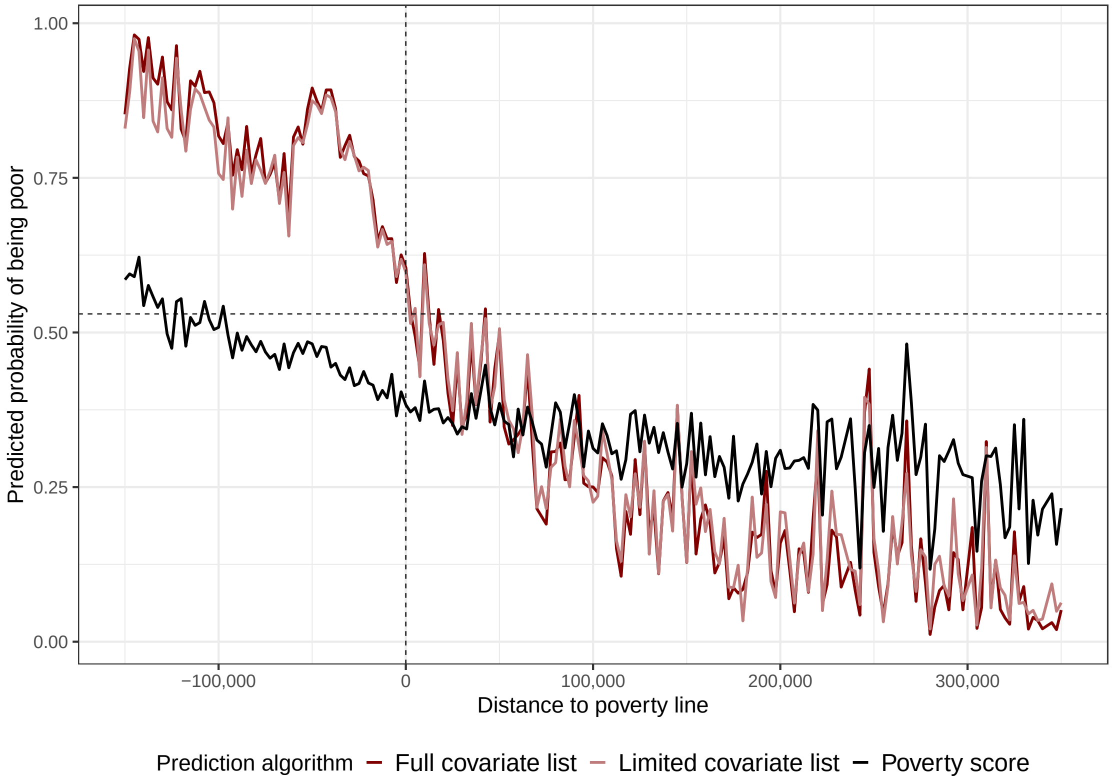

Figure plots out-of-sample poverty predictions (vertical axis) from the machine learning predictions using the full covariate list (dark red), the machine learning predictions using the limited covariate list (light red), and the government poverty score (black) against households' actual consumption levels (horizontal axis). Actual household consumption is defined as the distance to the poverty line for that year. Observations binned into 2,500 peso bins. For ease of viewing, values below $-150,000$ pesos from the poverty line and above 350,000 pesos from the poverty line are combined. Households to the right (left) of the dashed vertical line are above (below) the poverty line. Households above (below) the dashed horizontal line would (not) be considered eligible for social programs. Data from Quality of Life surveys.

Because log-loss is difficult to interpret, using households in the QOL holdout set, Figure 3.9 plots the predicted probability that a household is under the poverty line, using the machine learning predictions and the poverty score, against households' actual distance to the poverty line. A prediction function that performs well should have a predicted probability close to 1 for households far below the poverty line, as this indicates that the function classifies these households as having a high probability of being poor. Households far above the poverty line should have a predicted probability close to 0, as these households have a low probability of being poor. Households right at the poverty line should have a predicted probability close to 0.5, as these households have equal probability of being poor and not poor.

The machine learning predictions appear to perform well. For households below the poverty line (to the left of the vertical dashed line), the predicted probability increases with distance from the poverty line and is close to 1 for households far below the poverty line, indicating that the algorithm classifies households farther from the poverty line as more likely to be poor. For households above the poverty line (to the right of the vertical dashed line), the predicted probability decreases with distance from the poverty line towards 0. For households right at the poverty line, the predicted probability is close to 0.5. In addition, machine learning predictions trained on the full covariate list and on the limited covariate list track each other closely, reflecting the similarity in their loss measures in Table 3.8.

The poverty score, on the other hand, does not perform well. On the one hand, the poverty score is at least monotonic and has a negative slope: the poverty score increases with distance from the threshold, for households below the poverty line, and decreases with distance from the threshold, for households above the poverty line. However, even for the poorest households, the predicted probability is close to 0.5 and for the richest households, the predicted probability is close to 0.25. On average, only households over 100,000 pesos (58.1% of the 1997 minimum wage) below the poverty line would be classified as eligible in SISBÉN (above the horizontal dotted line). Further, the gap between the machine learning predictions and the poverty score is largest for the *poorest* households and much smaller for richer households. While not performing poorly for rich households, the poverty score performs poorly for the subset of households which the PMT is meant to accurately target.

Applying the machine learning predictions to households in SISBÉN, the mean predicted household poverty is 0.81 and 0.73 for the machine learning predictions trained on the full covariate list and the limited covariate list, respectively. The higher predicted poverty when using the full covariate list suggests that there may be some manipulation of the survey answers and that there may be an upward bias in the predictions trained on the full covariate list.

Table 3.10 Variables selected in the machine learning algorithm

Variable	Poverty score	Machine learning algorithm[†] Random forest	Boosted tree	Elastic-net
Year	–			L
Neighborhood strata	–	F,L	F,L	L
Household variables				
Material of housing exterior	X	–	–	–
Material of floor	X	–	–	–
Material of roof	X	–	–	–
Main source of lighting in home	–	–	–	–
Sanitation system	X	–	–	–
Main water source	X	–	–	–
Trash disposal system	X	–	–	–
Household owns: refrigerator	X			
Household owns: television	X			
Household owns: fan	X			
Household owns: blender	X			
Household owns: washing machine	X			
Household size	–	F,L	F,L	F,L
Total rooms in house	–			
Total bedrooms in house	–			
Average rooms per person	X	–	–	–
Household head / highest wage earner variables				
Age	–			L
Sex	–			
Marital status	–		F,L	F,L
Highest educational level	X	F,L	F,L	F,L
Years of schooling	–			
Main economic activity	X	F	L	L

[†] F = selected in full covariate list, L = selected in limited covariate list

Continued on next page

Table 3.10 *(continued)*

Variable	Poverty score	Machine learning algorithm† Random forest	Boosted tree	Elastic-net
Income from main job	–			F
Number of employees at main job	X		F	
Social security coverage	X	F	F	F
Individual variables (collapsed to household level)				
Relationship to head of household	–			
Average age	–			F
Age: 0-4 years	–			L
Age: 0-6 years	–		L	
Age: 5-18 years	–		F,L	
Age: 5+ years	–			
Age: 12+ years	–			
Age: 18-60 years	–			
Age: 60+ years	–			
Fraction of members 0-6 years	X			
Sex	–			
Marital status	–		F,L	
School enrollment status	–			L
School enrollment status (5-18 years)	–			
Highest educational level	–			F,L
Highest educational level (5-18 years)	–		F,L	F,L
Highest educational level (18-60 years)	–		F	L
Avg years of education	X			L
Max years of education	–			L
Median years of education	–			L
Variance in years of education	–	F,L	L	F
Avg years of education (5-18 years)	–	F,L	L	F

† F = selected in full covariate list, L = selected in limited covariate list

Continued on next page

Table 3.10 *(continued)*

Variable	Poverty score	Machine learning algorithm† Random forest	Boosted tree	Elastic-net
Avg years of education (18-60 years)	–	F,L	F	F
Main economic activity	–			F
Main economic activity (18-60 years)	–		F	F,L
Fraction of members employed	X	F		
Avg income from main job	–			F
Avg income from main job (18-60 years)	–	F		F
Avg income from main job, relative to min. wage	X			
Total income from main job	–			
Total income from main job (18-60 years)	–			
Number of employees at main job	–		F	
Number of employees at main job (18-60 years)	–			
Social security coverage	–			
Social security coverage (18-60 years)	–			F

† F = selected in full covariate list, L = selected in limited covariate list

Table indicates the 20 variables selected as most important for prediction by the machine learning algorithms. Variable importance defined using the GCV criterion for the random forest and boosted tree and using the coefficient magnitude for the elastic-net regression.

3.5.2 Predictors of poverty

Prior to tuning the machine learning algorithms, in line with standard practice, I removed variables with little variation and highly correlated predictors. A large number of the variables removed involved housing characteristics (housing material and access to utilities), which encompass two of the subcomponents of the poverty score. This suggests that variable selection and over-fitting due to unregularized procedures likely contribute to the poor performance of the poverty score, as the poverty score weighted these variables fairly heavily despite there being little variation in the data.

Table 3.10 presents the variables selected as most predictive in the random forest, boosted tree, and elastic-net models. In general, variables related to educational attainment and characteristics of the household head are very predictive across the three algorithms. This is partially in line with the poverty score, which has one subcomponent for characteristics of the household head. However, the machine learning algorithms place much more weight than the poverty score on the educational attainment of household members. In addition, several variables used in the poverty score are not selected by any of the machine learning algorithms: average rooms per person, fraction of household members between 0-6 years, and income per capita relative to the minimum wage. Interestingly, these variables involve several non-linear transformations of the raw variables (for example, dividing the number of rooms by the number of household members or dividing by the minimum wage for that year). This suggests that incorrect functional form is another potential contributor to the poor performance of the poverty score.

Turning to the ensemble weights, both when using the full covariate list and the limited covariate list, the ensemble places large weight on predictions from the boosted tree, moderate weight on predictions from the random forest, and the least weight on predictions from the elastic-net regression, which was the lasso regression (Table 3.11). In fact, for algorithms trained on the limited covariate list, the ensemble places close to zero weight on the elastic-net regression. This suggests that local predictors, which allow for detection of much finer interactions in the data, can greatly improve the performance of poverty predictions.

3.6 Implications for targeting

If the government poverty score predicts poverty poorly, do households self-target by manipulating their poverty scores? I examine whether households who have their scores altered in order to fall under the eligibility threshold have different machine learning predictions. On average, households who have their scores altered have similar machine learning predictions

Table 3.11 Prediction weights on machine learning algorithms

	Full covariate list	Limited covariate list
Random forest	0.173***	0.156***
	(0.019)	(0.019)
Boosted tree	0.740***	0.845***
	(0.024)	(0.025)
Elastic-net	0.094***	0.004
	(0.021)	(0.022)

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Weights on machine learning algorithms using algorithms trained on the full covariate list and on the limited covariate list. Weights obtained from a regression of an indicator for poverty on predictions from the tuned random forest, boosted tree, and elastic-net regression in the training set. Data from Quality of Life surveys.

as households whose eligibility status did not change (left panel of Figures 3.12, using the full covariate list, and 3.13, using the limited covariate list). The exception is households with reconstructed poverty scores right above the eligibility threshold; for these households, households with altered scores are predicted as *less* poor than all other households.

However, a different pattern emerges when separately examining households surveyed before and after the release of the poverty score algorithm. For households surveyed between 1994-1997, households with altered scores are *less* poor than all other households (middle panel of Figures 3.12 and 3.13). On average, households with an altered score are predicted as 7.4 and 5.9 percentage points less likely to be poor, when using predictions trained on the full covariate list and the limited covariate list respectively (Table 3.14). For households surveyed between 1998-2003, the opposite is true: households with altered scores are *more* poor than all other households (right panel of Figures 3.12 and 3.13). On average, households with an altered score are predicted as 0.7 and 3.6 percentage points more likely to be poor, when using predictions trained on the full covariate list and the limited covariate list respectively (Table 3.14). These differences are significant at the 1% level.

Observing the opposite correlation between score alteration and household poverty among households surveyed before and after 1997 suggests that different types of households altered

Figure 3.12 Machine learning predictions using the full covariate list of households with altered poverty scores

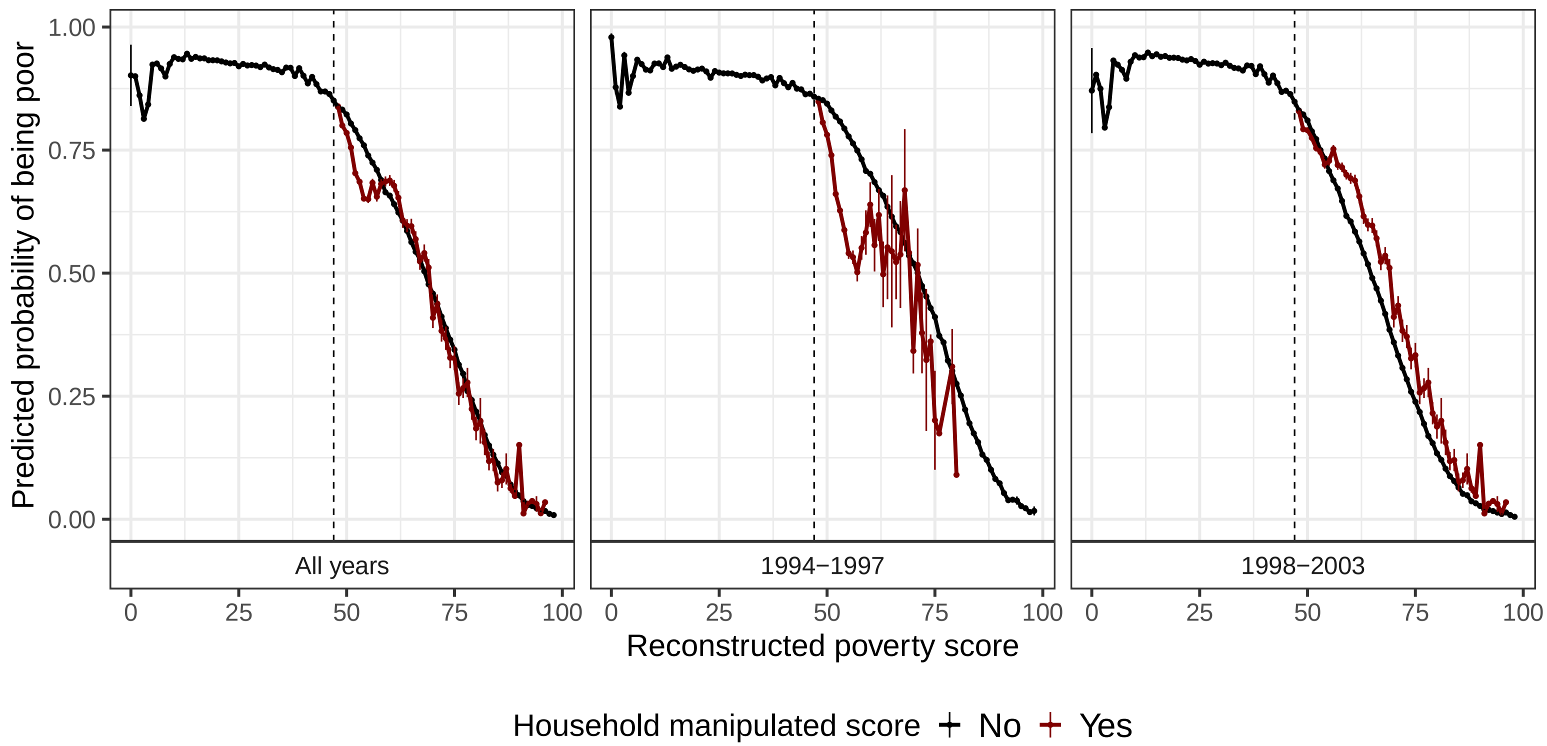

Figures plot the average predicted poverty by reconstructed poverty score of households with an altered poverty score that changed their eligibility status (red line) and all other households (black line). Plotted separately for households surveyed across all years (left panel), 1994-1997 (middle panel), and 1998-2003 (right panel). Vertical axis is the household machine learning prediction trained on the full covariate list. Horizontal axis is the reconstructed poverty score, where the poverty score is reconstructed from the household's survey answers or the sum of the poverty score subcomponents. The vertical lines represent standard error bars around the estimate. The poverty score algorithm was released to local officials in late 1997. Households with a poverty score below 47 (vertical dotted line) are eligible for a variety of social programs. Data from SISBÉN survey.

Figure 3.13 Machine learning predictions using the limited covariate list of households with altered poverty scores

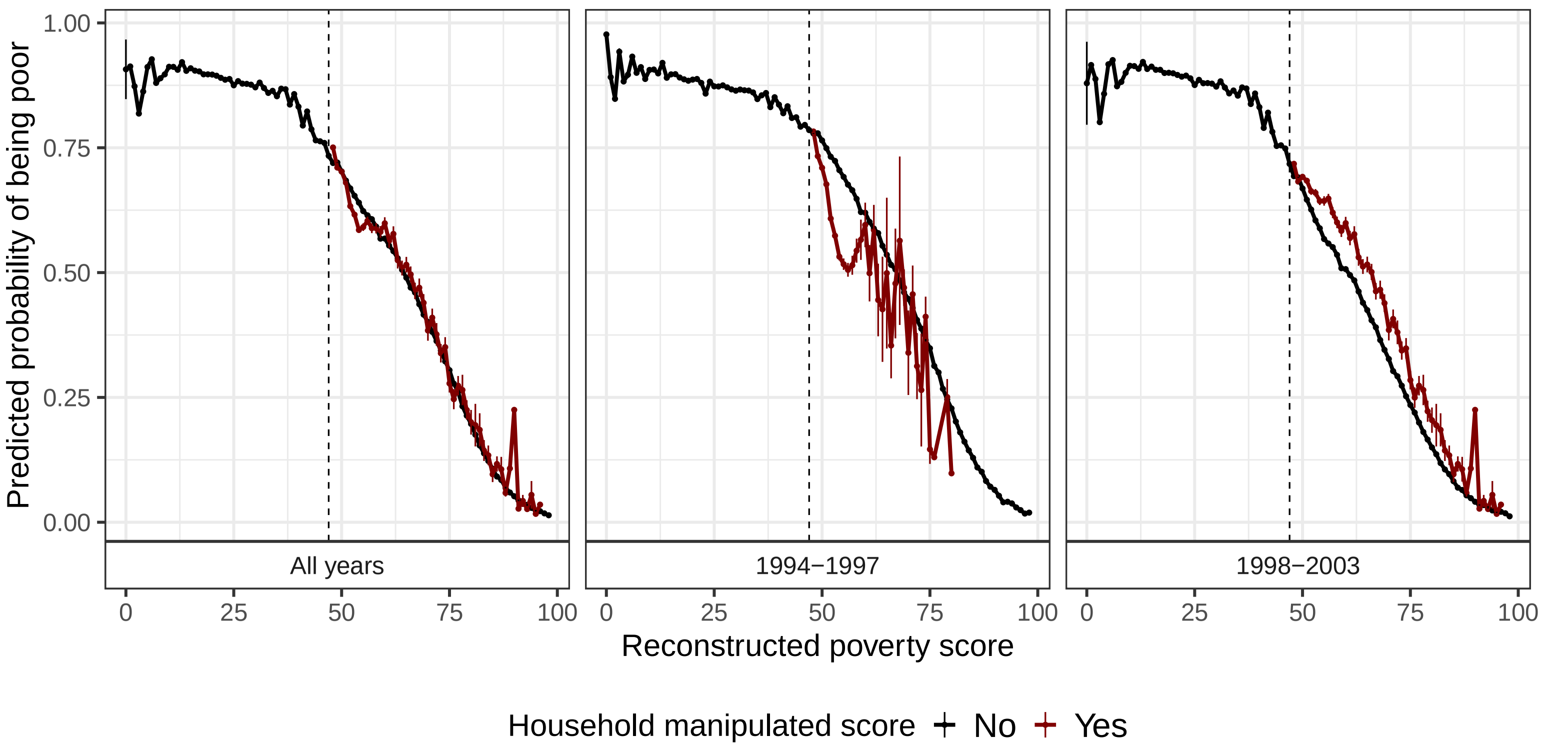

Figures plot the average predicted poverty by reconstructed poverty score of households with an altered poverty score that changed their eligibility status (red line) and all other households (black line). Plotted separately for households surveyed across all years (left panel), 1994-1997 (middle panel), and 1998-2003 (right panel). Vertical axis is the household machine learning prediction trained on the limited covariate list. Horizontal axis is the reconstructed poverty score, where the poverty score is reconstructed from the household's survey answers or the sum of the poverty score subcomponents. The vertical lines represent standard error bars around the estimate. The poverty score algorithm was released to local officials in late 1997. Households with a poverty score below 47 (vertical dotted line) are eligible for a variety of social programs. Data from SISBÉN survey.

Table 3.14 Machine learning predictions of households with altered poverty scores

	Machine learning predictions using	
	Full covariate list	Limited covariate list
	(1)	(2)
Altered score	-0.074^{***}	-0.059^{***}
	(0.002)	(0.002)
Altered score × Post-1998	0.081^{***}	0.095^{***}
	(0.002)	(0.002)
Altered score + Altered score × Post-1998	0.007^{***}	0.036^{***}
	(0.000)	(0.000)
N	2,098,162	2,098,162

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variable is household machine learning predictions using the full covariate list (Column 1) and the limited covariate list (Column 2). *Altered score* is an indicator equal to 1 if the household's reconstructed poverty score does not match the assigned poverty score and the altered score changes their eligibility status. *Post-1998* is an indicator equal to 1 if the household was surveyed between 1998-2003 and 0 if the household was surveyed between 1994-1997. The poverty score is reconstructed from the household's survey answers or the sum of the poverty score subcomponents. The poverty score algorithm was released to local officials in late 1997. Households with a poverty score below 47 are eligible for a variety of social programs. *Estimation method*: Regression includes fixed effects for the value of the reconstructed poverty score. Data from SISBÉN.

their scores. In other words, when the poverty score algorithm was private, households *negatively* self-targeted: richer households altered their score. When the poverty score algorithm became public, households *positively* self-targeted: poorer households altered their score. One explanation is that obtaining information on the targeting mechanism was costly between 1994-1997, as it was private during this period, and that the cost of obtaining information decreased once the algorithm was made public. Between 1994-1997, only households who were able to obtain this information altered their poverty scores, resulting in the negative correlation between score alteration and poverty.

An alternative explanation is that households surveyed between 1998-2003 manipulated their survey answers at higher rates and that manipulating survey answers is positively correlated with altering the poverty score. This will lead to a positive correlation between score alteration and predicted poverty, even if there is negative self-targeting (ie. households who both manipulate their survey answers and alter their scores are richer). However, I do not find evidence that manipulation of survey answers is positively correlated with altering the poverty score. Given that in 1998-2003 (i) the rate of poverty score alteration was *lower* (Section 3.4.1) and (ii) households likely manipulated their survey answers at higher rates, it is more likely that this correlation is negative. To investigate this further, I estimate the correlation between altering the poverty score and the difference in machine learning predictions when using the full covariate list and the limited covariate list. Since predictions using the full covariate list are more likely to be upward biased due to manipulated survey answers, a positive difference in the machine learning predictions would suggest that the household has manipulated their survey answers. However, I find a *negative* correlation between the difference in machine learning predictions and household score alteration (Table C.6). As a result, it is unlikely that the positive correlation between score alteration and predicted poverty is due to households both manipulating their survey answers and altering their poverty scores.

3.7 Conclusion

As governments spend more on social programs to assist the poor, governments have focused more attention on determining how benefits are to be allocated. In developing countries, targeting accuracy and agency problems among officials determining eligibility are particularly acute issues. This paper studies the implementation of SISBÉN, a national proxy-means testing system in Colombia, which uses a poverty score to quantify households' economic well-being and determine eligibility for social programs. I first study how households respond to the public release of the poverty score algorithm, by examining patterns of manipulation of household poverty scores. I next evaluate the targeting accuracy of the poverty score, by training a machine learning algorithm to predict household poverty and comparing the performance of the machine learning predictions to the government poverty score in an external data set. Applying the machine learning predictions to households in SISBÉN, I test whether households self-target by manipulating their poverty scores.

I find that households responded to the public release of the poverty score algorithm by manipulating their poverty scores. First, there is a significant discontinuity in the density of households at the eligibility threshold which is present only after the release of the algorithm. Second, I document that one method which households use to manipulate their eligibility involves directly altering the poverty score.

I evaluate the targeting accuracy of the government poverty score relative to machine learning predictions of household poverty. I find that the poverty score performs worse than the machine learning predictions. In particular, the poverty score underpredicts poverty *more* for households far below the poverty line and on average does not categorize these households as eligible for social programs.

Finally, I provide suggestive evidence that households self-target by altering their poverty scores. Prior to the release of the poverty score algorithm, households *negatively* self-targeted: households with altered scores are predicted as richer than other households. However, after the release of the poverty score algorithm, households *positively* self-targeted: households with altered scores are predicted as poorer than other households.

These findings suggest that manipulation of targeting programs does occur. However, whether manipulation leads to better or worse targeting accuracy depends on the context. In this study, targeting accuracy increased when households had greater access to information on program rules and targeting mechanisms. This does not necessarily mean that manipulation leads to higher welfare, as manipulating scores may be costly to households. However, it does suggest that targeting systems that allow the use of private information can target more accurately. Future studies that evaluate either the targeting accuracy or the private costs of targeting mechanisms that allow the use of private information could help improve social program targeting and poverty alleviation goals.

References

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey Wooldridge (2017). "When Should You Adjust Standard Errors for Clustering?" *Working Paper*.

Acemoglu, Daron and James A. Robinson (2008). "Persistence of Power, Elites, and Institutions." *American Economic Review* 98.1, pp. 267–293.

Acemoglu, Daron, Tristan Reed, and James A. Robinson (2014). "Chiefs: Economic Development and Elite Control of Civil Society in Sierra Leone." *Journal of Political Economy* 122.2, pp. 319–368.

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Julia Tobias (2012). "Targeting the Poor: Evidence from a Field Experiment in Indonesia." *American Economic Review* 102.4, pp. 1206–1240.

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi (2016). "Self-targeting: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy* 124.2, pp. 371–427.

Amengual, Matthew and Greg Distelhorst (2019). "Can Sourcing Help Enforce Global Labor Standards? Evidence from the Gap Inc Supply Chain." *Working Paper*.

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin (1996). "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91.434, pp. 444–455.

Ashraf, Nava, Natalie Bau, Nathan Nunn, and Alessandra Voena (2020). "Bride Price and Female Education." *Journal of Political Economy* 128.2, pp. 591–641.

Atkin, David (2009). "Working for the Future: Female Factory Work and Child Health in Mexico." *Working Paper*.

— (2015). "Endogenous Skill Acquisition and Export Manufacturing in Mexico." *Working Paper*.

Bardhan, Pranab and Dilip Mookherjee (2006). "Decentralisation and Accountability in Infrastructure Delivery in Developing Countries." *The Economic Journal* 116, pp. 101–127.

Besley, Timothy and Robin Burgess (2004). "Can Labor Regulation Hinder Economic Performance? Evidence from India." *Quarterly Journal of Economics* 119.1, pp. 91–134.

Besley, Timothy and Stephen Coate (1992). "Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs." *American Economic Review* 82.1, pp. 249–261.

Besley, Timothy and Maitreesh Ghatak (2007). "Retailing public goods: The economics of corporate social responsibility." *Journal of Public Economics* 91, pp. 1645–1663.

Blumenstock, Joshua (2016). "Fighting poverty with data." *Science* 353.6301, pp. 753–754.

Blumenstock, Joshua, Gabriel Cadamuro, and Robert On (2015). "Predicting poverty and wealth from mobile phone metadata." *Science* 350.6264, pp. 1073–1076.

Bordignon, Massimo, Tommaso Nannicini, and Guido Tabellini (2016). "Moderating Political Extremism: Single Round versus Runoff Elections under Plurality Rule." *American Economic Review* 106.8, pp. 2349–2370.

Bormann, Nils-christian and Matt Golder (2013). "Democratic Electoral Systems around the world, 1946-2011." *Electoral Studies* 32.2, pp. 360–369.

Botero, Juan C., Simeon Djankov, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer (2004). "The Regulation of Labor." *Quarterly Journal of Economics* 119.4, pp. 1339–1382.

Boudreau, Laura (2020). "Multinational enforcement of labor law: Experimental evidence from Bangladesh's apparel sector." *Working Paper*.

Bouton, Laurent (2013). "A Theory of Strategic Voting in Runoff Elections." *American Economic Review* 103.4, pp. 1248–1288.

Bouton, Laurent and Gabriele Gratton (2015). "Majority Runoff Elections: Strategic Voting and Duverger's Hypothesis." *Theoretical Economics* 10.2, pp. 283–314.

Bouton, Laurent, Jorge Gallego, Aniol Llorente-Saguer, and Rebecca Morton (2019). "Runoff Elections in the Laboratory." *NBER Working Paper*.

Brambilla, Irene, Daniel Lederman, and Guido Porto (2012). "Exports, Export Destinations, and Skills." *American Economic Review* 102.7, pp. 3406–3438.

Breierova, Lucia and Esther Duflo (2004). "The Impact of Education on Fertility and Child Mortality: Do Fathers Really Matter Less than Mothers?" *NBER Working Paper*.

Brollo, Fernanda and Tommaso Nannicini (2012). "Tying Your Enemy's Hands in Close Races: The Politics of Federal Transfers in Brazil." *American Political Science Review* 106.4, pp. 742–761.

Burden, Barry C. (1997). “Deterministic and Probabilistic Voting Models.” *American Journal of Political Science* 41.4, pp. 1150–1169.

Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró-i Miquel (2015). “The Value of Democracy: Evidence from Road Building in Kenya.” *American Economic Review* 105.6, pp. 1817–1851.

Bustos, Paula (2011). “The Impact of Trade Liberalization on Skill Upgrading: Evidence from Argentina.” *Working Paper.*

Callander, Steven (2005). “Duverger’s Hypothesis, the Run-off Rule, and Electoral Competition.” *Political Analysis* 13.3, pp. 209–232.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik (2014). “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs.” *Econometrica* 82.6, pp. 2295–2326.

Camacho, Adriana and Emily Conover (2011). “Manipulation of Social Program Eligibility.” *American Economic Journal: Economic Policy* 3.2, pp. 41–65.

Carey, John M. and Matthew Soberg Shugart (1993). “Incentives to Cultivate a Personal Vote: A Rank Ordering of Electoral Formulas.” *Electoral Studies* 14.4, pp. 417–439.

Castañeda, Tarsicio (2005). “Targeting social spending to the poor with proxy-means testing: Colombia’s SISBEN system.” *World Bank Social Protection Dscussion Paper.*

Chamon, Marcos, Joao De Mello, and Sergio Firpo (2018). “Electoral Rules, Political Competition and Fiscal Expenditures: Regression Discontinuity Evidence from Brazilian Municipalities.” *Journal of Development Studies.*

Cipullo, Davide (2019). “Voting Systems, Incentives and Fiscal Policy: Evidence from Runoff and Plurality.” *Working Paper.*

Coady, David, Margaret Grosh, and John Hoddinott (2004a). *Targeting of Transfers in Developing Countries: Review of Lessons and Experience.* Washington DC: The World Bank.

— (2004b). “Targeting Outcomes Redux.” *World Bank Research Observer* 19.1, pp. 61–85.

Cox, Gary W. (1997). *Making Votes Count: Strategic Coordination in the World’s Electoral Systems.* Cambridge University Press.

Dragusanu, Raluca and Nathan Nunn (2018). “The Effects of Fair Trade Certification: Evidence From Coffee Producers in Costa Rica.” *Working Paper.*

Duflo, Esther (2003). “Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa.” *World Bank Economic Review* 17.1, pp. 1–25.

Eggers, Andrew C. (2013). “Proportionality and Turnout: Evidence from French Municipalities.” *Comparative Political Studies* 48.2, pp. 135–167.

Eggers, Andrew C., Ronny Freier, Veronica Grembi, and Tommaso Nannicini (2015). “Regression Discontinuity Designs Based on Population Thresholds: Pitfalls and Solutions.” *American Journal of Political Science*.

Fujiwara, Thomas (2011). “A Regression Discontinuity Test of Strategic Voting and Duverger’s Law.” *Quarterly Journal of Political Science* 6, pp. 197–233.

— (2015). “Voting Technology, Political Responsiveness and Infant Health: Evidence from Brazil.” *Econometrica* 83.2, pp. 423–464.

Fujiwara, Thomas and Leonard Wantchekon (2013). “Can Informed Public Deliberation Overcome Clientelism? Experimental Evidence from Benin.” *American Economic Journal: Applied Economics* 5.4, pp. 241–255.

Genicot, Garance, Laurent Bouton, and Micael Castanheira (2018). “Electoral Systems and Inequalities in Government Interventions.” *NBER Working Paper*.

Glaeser, Edward L., Scott Duke Kominers, Michael Luca, and Nikhil Naik (2018). “Big Data and Big Cities: The Promises and Limitations of Improved Measures of Urban Life.” *Economic Inquiry* 56.1, pp. 114–137.

Golden, Miriam and Brian Min (2013). “Distributive Politics Around the World.” *Annual Review of Political Science* 16, pp. 73–99.

Gonzales, Mariella, Gianmarco León-Ciliotta, and Luis R. Martínez (2019). “How Effective Are Monetary Incentives to Vote? Evidence from a Nationwide Policy.” *Working Paper*.

Harrison, Ann and Jason Scorse (2010). “Multinationals and Anti-Sweatshop Activism.” *American Economic Review* 100.1, pp. 247–273.

Heath, Rachel and A. Mushfiq Mobarak (2015). “Manufacturing growth and the lives of Bangladeshi women.” *Journal of Development Economics* 115, pp. 1–15.

Heckman, James J. and Jeffrey A. Smith (2004). “The Determinants of Participation in a Social Program: Evidence from a Prototypical Job Training Program.” *Journal of Labor Economics* 22.2, pp. 243–298.

Hodler, Roland and Paul A. Raschky (2014). “Regional Favoritism.” *Quarterly Journal of Economics* 129.2, pp. 995–1033.

Huber, Evelyne and John D. Stephens (2012). *Democracy and the left: Social policy and inequality in Latin America*. University of Chicago Press.

Imbens, Guido and Karthik Kalyanaraman (2012). "Optimal bandwidth choice for the regression discontinuity estimator." *Review of Economic Studies* 79.3, pp. 933–959.

Imbens, Guido W. and Thomas Lemieux (2008). "Regression discontinuity designs: A guide to practice." *Journal of Econometrics* 142.2, pp. 615–635.

Jensen, Robert (2012). "Do labor market opportunities affect young women's work and family decisions? Experimental evidence from India." *Quarterly Journal of Economics* 127.2, pp. 753–792.

Kidd, Stephen and Emily Wylde (2011). "Targeting the Poorest: An assessment of the proxy means test methodology." *AusAID Working Paper*.

Klasen, Stephan and Simon Lange (2016). "How Narrowly Should Anti-Poverty Programs be Targeted? Simulation Evidence from Bolivia and Indonesia." *Courant Research Centre Discussion Paper*.

Kleven, Henrik Jacobsen and Wojciech Kopczuk (2011). "Transfer Program Complexity and the Take-Up of Social Benefits." *American Economic Journal: Economic Policy* 3.1, pp. 54–90.

Lindbeck, Assar and Jörgen W. Weibull (1987). "Balanced-Budget Redistribution as the Outcome of Political Competition." *Public Choice* 52.3, pp. 273–297.

Lizzeri, Alessandro and Nicola Persico (2001). "The Provision of Public Goods under Alternative Electoral Incentives." *American Economic Review* 91.1, pp. 225–239.

— (2005). "A Drawback of Electoral Competition." *Journal of the European Economic Association* 3.6, pp. 1318–1348.

Locke, Richard, Matthew Amengual, and Akshay Mangla (2009). "Virtue out of Necessity? Compliance, Commitment, and the Improvement of Labor Conditions in Global Supply Chains." *Politics and Society* 37.3, pp. 319–351.

Locke, Richard M. and Monica Romis (2010). "The promise and perils of private voluntary regulation: Labor standards and work organization in two Mexican garment factories." *Review of International Political Economy* 17.1, pp. 45–74.

Locke, Richard M., Fei Qin, and Alberto Brause (2007). "Does Monitoring Improve Labor Standards? Lessons from Nike." *Industrial and Labor Relations Review* 61.1, pp. 3–31.

McCrary, Justin (2008). “Manipulation of the running variable in the regression discontinuity design: A density test.” *Journal of Econometrics* 142.2, pp. 698–714.

Milesi-Ferretti, Gian Maria, Roberto Perotti, and Massimo Rostagno (2002). “Electoral systems and public spending.” *Quarterly Journal of Economics* 117.2, pp. 609–657.

Mullainathan, Sendhil and Jann Spiess (2017). “Machine Learning: An Applied Econometric Approach.” *Journal of Economic Perspectives* 31.2, pp. 87–106.

Munshi, Kaivan and Mark Rosenzweig (2006). “Traditional Institutions Meet the Modern World: Caste, Gender, and Schooling Choice in a Globalizing Economy.” *American Economic Review* 96.4, pp. 1225–1260.

Murphy, David F. and David Mathew (2001). “Nike and Global Labour Practices.” *Case Study Prepared for the New Academy of Business Innovation Network for Socially Responsible Business.*

Myerson, Roger B. (1993). “Incentives to Cultivate Favored Minorities Under Alternative Electoral Systems.” *American Political Science Review* 87.4, pp. 856–869.

Niehaus, Paul and Sandip Sukhtankar (2013). “Corruption Dynamics: The Golden Goose Effect.” *American Economic Journal: Economic Policy* 5.4, pp. 230–269.

Olken, Benjamin A. (2006). “Corruption and the Costs of Redistribution: Micro Evidence from Indonesia.” *Journal of Public Economics* 90.4-5, pp. 853–870.

Olken, Benjamin A. and Rohini Pande (2012). “Corruption in Developing Countries.” *Annual Review of Economics* 4, pp. 479–509.

Osborne, Martin J. and Al Slivinski (1996). “A Model of Political Competition with Citizens-Candidates.” *Quarterly Journal of Economics* 111.1, pp. 65–96.

Pellicer, Miquel and Eva Wegner (2013). “Electoral Rules and Clientelistic Parties: A Regression Discontinuity Approach.” *Quarterly Journal of Political Science* 8, pp. 1–33.

Persson, Torsten and Guido Tabellini (1999). “The size and scope of government: Comparative politics with rational politicians.” *European Economic Review* 43, pp. 699–735.

— (2000). *Political Economics: Explaining Economic Policy.* Cambridge, MA: MIT Press.

— (2004). “Constitutions and Economic Policy.” *Journal of Economic Perspectives* 18.1, pp. 75–98.

Persson, Torsten, Guido Tabellini, and Francesco Trebbi (2003). “Electoral Rules and Corruption.” *Journal of the European Economic Association* 1.4, pp. 958–989.

Persson, Torsten, Gerard Roland, and Guido Tabellini (2004). “How do Electoral Rules Shape Party Structures, Government Coalitions, and Economic Policies?” *CESifo Working Paper Series 1115.*

Piketty, Thomas (2000). “Voting as Communicating.” *Review of Economic Studies* 67, pp. 169–191.

Pons, Vincent and Clemence Tricaud (2018). “Expressive Voting and Its Cost: Evidence from Runoffs with Two or Three Candidates.” *Econometrica* 86.5, pp. 1621–1649.

Qian, Nancy (2008). “Missing Women and the Price of Tea in China: The Effect of Sex-Specific Income on Sex Im-balance.” *Quarterly Journal of Economics* 123.3, pp. 1251–1285.

Rama, Martin (2001). “The Consequences of Doubling the Minimum Wage: The Case of Indonesia.” *Industrial and Labor Relations Review* 54.4, pp. 864–881.

Reardon, Sean F. and David O’Sullivan (2004). “Measures of Spatial Segregation.” *Sociological Methodology* 34.1, pp. 121–162.

Reinikka, Ritva and Jakob Svensson (2004). “Local Capture: Evidence from a Central Government Transfer Program in Uganda.” *Quarterly Journal of Economics* 119.2, pp. 679–705.

Schaner, Simone and Smita Das (2016). “Female Labor Force Participation in Asia: Indonesia Country Study.” *ADB Economics Working Paper Series* 474.

Short, Jodi L., Michael W. Toffel, and Andrea R. Hugill (2015). “Monitoring Global Supply Chains.” *Strategic Management Journal* 37.9, pp. 1878–1897.

Toffel, Michael W., Jodi L. Short, and Melissa Ouellet (2015). “Codes in Context: How States, Markets, and Civil Society Shape Adherence to Global Labor Standards.” *Regulation and Governance* 9.3, pp. 205–223.

Vélez, Carlos Eduardo, Elkin Castaño, and Ruthanne Deutsch (1998). “An Economic Interpretation of Colombia’s SISBEN: A Composite Welfare Index Derived from the Optimal Scaling Algorithm.” *Inter-American Development Bank Working Paper.*

Verhoogen, Eric A. (2008). “Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector.” *Quarterly Journal of Economics* 123.2, pp. 489–530.

White, Michael J. (1986). “Segregation and Diversity Measures in Population Distribution.” *Population Index* 52.2, pp. 198–221.

Appendix A

Appendix to Chapter 1

A.1 Theory Appendix

A.1.1 Swingable voters

For $0 < \pi_{j1}^A < 1$, we need that:

$$u_j(q_j^B) - u_j(q_j^A) - \delta \in \left(-\frac{1}{2\psi}, \frac{1}{2\psi}\right)$$

Let $u(y)$ be the largest possible utility coming from the allocation of government resources. This assumption is satisfied if:

$$\begin{aligned}\delta \in &\left(u(y) - u(0) - \frac{1}{2\psi}, u(y) - u(0) + \frac{1}{2\psi}\right) \\ \iff &\frac{1}{2\gamma} + u(y) - u(0) < \frac{1}{2\psi}\end{aligned}$$

In other words, that swings in municipality vote shares are smaller than the variation in individual preferences. Note that this implies that $\gamma > \psi$ since $u(y) - u(0) > 0$.

A.1.2 Deriving vote shares with three candidates

Condition (1.1) corresponds to voters for whom $v_i \geq u_j(q_j^B) - u_j(q_j^A) - \delta$. Since there are swingable voters in every locality, π_{jt}^c is given by:

$$\pi_{j1}^A = \begin{cases} \frac{1}{2} + \psi\left(\Delta u_j^{AB} + \delta\right) & \text{if } j \in \{1, \ldots, J-1\} \\ \alpha\left(\frac{1}{2} + \psi\left(\Delta u_j^{AB} + \delta\right)\right) & \text{if } j = J \end{cases}$$

$$\pi_{j1}^{B}=\begin{cases}\frac{1}{2}+\psi\left(\Delta u_{j}^{BA}-\delta\right) & \text{if } j\in\{1,\ldots,J-1\}\\ \alpha\left(\frac{1}{2}+\psi\left(\Delta u_{j}^{BA}-\delta\right)\right) & \text{if } j=J\end{cases}$$

$$\pi_{j1}^{C}=\begin{cases}0 & \text{if } j\in\{1,\ldots,J-1\}\\ 1-\alpha & \text{if } j=J\end{cases}$$

Candidates' total vote share in the municipality π_t^c is given by:

$$\pi_1^A=\left(\frac{J-1+\alpha}{J}\right)\left(\frac{1}{2}+\psi\delta\right)+\frac{\psi}{J}\left(\sum_{j=1}^{J-1}\Delta u_j^{AB}+\alpha\Delta u_J^{AB}\right)$$
$$\pi_1^B=\left(\frac{J-1+\alpha}{J}\right)\left(\frac{1}{2}-\psi\delta\right)+\frac{\psi}{J}\left(\sum_{j=1}^{J-1}\Delta u_j^{BA}+\alpha\Delta u_J^{BA}\right)$$
$$\pi_1^C=\frac{1-\alpha}{J}$$

The probability that candidates A and B attain a vote share above θ is equivalent to:

$$Pr\left(\pi_1^A\geq\theta\right)\equiv Pr\left[\delta\geq\frac{1}{\psi}\left(\frac{J}{J-1+\alpha}\theta-\frac{\psi}{J-1+\alpha}\left(\sum_{j=1}^{J-1}\Delta u_j^{AB}+\alpha\Delta u_J^{AB}\right)-\frac{1}{2}\right)\right]$$
$$Pr\left(\pi_1^B\geq\theta\right)\equiv Pr\left[\delta\leq\frac{1}{\psi}\left(\frac{1}{2}+\frac{\psi}{J-1+\alpha}\left(\sum_{j=1}^{J-1}\Delta u_j^{BA}+\alpha\Delta u_j^{BA}\right)-\frac{J}{J-1+\alpha}\theta\right)\right]$$

A.1.3 Deriving vote shares with two candidates

Candidates' vote shares in each locality and the municipality as a whole are given by:

$$\pi_{j2}^A=\frac{1}{2}+\psi\left(\Delta u_j^{AB}+\delta\right) \qquad \pi_{j2}^B=\frac{1}{2}+\psi\left(\Delta u_j^{BA}-\delta\right)$$
$$\pi_2^A=\frac{1}{2}+\psi\delta+\frac{\psi}{J}\sum_{j=1}^{J}\Delta u_j^{AB} \qquad \pi_2^B=\frac{1}{2}-\psi\delta+\frac{\psi}{J}\sum_{j=1}^{J}\Delta u_j^{BA}$$

and the probability of attaining a vote share above θ is equivalent to:

$$Pr\left(\pi_2^A \geq \theta\right) \equiv Pr\left[\delta \geq \frac{1}{\psi}\left(\theta - \frac{1}{2} - \frac{\psi}{J}\sum_{j=1}^{J}\Delta u_j^{AB}\right)\right]$$
$$Pr\left(\pi_2^B \geq \theta\right) \equiv Pr\left[\delta \leq \frac{1}{\psi}\left(\frac{1}{2} - \theta + \frac{\psi}{J}\sum_{j=1}^{J}\Delta u_j^{BA}\right)\right]$$

A.1.4 Contestability of localities

For $0 < Pr(\pi_1^A \geq \theta) < 1$ and $0 < Pr(\pi_1^B \geq \theta) < 1$, we need that:

$$\frac{1}{\psi}\left[\left(\frac{J}{J-1+\alpha}\right)\theta - \frac{1}{2} - \left(\frac{\psi}{J-1+\alpha}\right)\left(\sum_{j=1}^{J-1}\Delta u_j^{AB} + \alpha\Delta u_J^{AB}\right)\right] \in \left(-\frac{1}{2\gamma}, \frac{1}{2\gamma}\right)$$
$$\frac{1}{\psi}\left[\frac{1}{2} - \left(\frac{J}{J-1+\alpha}\right)\theta + \left(\frac{\psi}{J-1+\alpha}\right)\left(\sum_{j=1}^{J-1}\Delta u_j^{BA} + \alpha\Delta u_J^{BA}\right)\right] \in \left(-\frac{1}{2\gamma}, \frac{1}{2\gamma}\right)$$

which corresponds to the following condition for the first round:

$$\theta \in \left(\left(\frac{J-1+\alpha}{J}\right)\left(-\frac{\psi}{2\gamma} + \frac{1}{2}\right), \left(\frac{J-1+\alpha}{J}\right)\left(\frac{\psi}{2\gamma} + \frac{1}{2}\right)\right) \tag{A.1}$$

For $0 < Pr\left(\pi_2^A \geq \theta\right) < 1$ and $0 < Pr\left(\pi_2^B \geq \theta\right) < 1$, we need that:

$$\frac{1}{\psi}\left[\theta - \frac{1}{2} - \frac{\psi}{J}\sum_{j=1}^{J}\Delta u_j^{AB}\right] \in \left(-\frac{1}{2\gamma}, \frac{1}{2\gamma}\right)$$
$$\frac{1}{\psi}\left[\frac{1}{2} - \theta + \frac{\psi}{J}\sum_{j=1}^{J}\Delta u_j^{BA}\right] \in \left(-\frac{1}{2\gamma}, \frac{1}{2\gamma}\right)$$

which corresponds to the following condition for the second round:

$$\theta \in \left(-\frac{\psi}{2\gamma} + \frac{1}{2}, \frac{\psi}{2\gamma} + \frac{1}{2}\right) \tag{A.2}$$

Claim A.1. *$\theta = \frac{1}{2}\left(1 - \frac{1-\alpha}{J}\right)$ satisfies condition (A.1).*

Both the upper and lower inequalities are satisfied because:

$$-\frac{\psi}{2\gamma} + \frac{1}{2} < \frac{1}{2} < \frac{\psi}{2\gamma} + \frac{1}{2}$$

Claim A.2. *$\theta = \frac{1}{2}$ satisfies condition (A.1).*

The lower inequality is satisfied:

$$\left(\frac{J-1+\alpha}{J}\right)\left(-\frac{\psi}{2\gamma} + \frac{1}{2}\right) < \frac{1}{2}$$

because $\frac{J-1+\alpha}{J} < 1$ and $-\frac{\psi}{2\gamma} + \frac{1}{2} < \frac{1}{2}$.

The upper inequality is equivalent to:

$$\frac{1}{2} < \left(\frac{J-1+\alpha}{J}\right)\left(\frac{\psi}{2\gamma} + \frac{1}{2}\right)$$
$$\iff J > (1-\alpha)\left(\frac{\gamma+\psi}{\psi}\right)$$

which is true so long as J is large enough and γ/ψ is not too large.

Claim A.3. *$\theta = \frac{1}{2}$ satisfies condition (A.2).*

Both the upper and lower inequalities are satisfied because:

$$-\frac{\psi}{2\gamma} + \frac{1}{2} < \frac{1}{2} < \frac{\psi}{2\gamma} + \frac{1}{2}$$

A.1.5 C never makes it to the second round

For C to never make it to the second round, the probability that candidates A and B attain vote shares above candidate C's must be 1, or π_1^C does not satisfy condition (A.1):

$$\frac{1-\alpha}{J} \leq \left(\frac{J-1+\alpha}{J}\right)\left(-\frac{\psi}{2\gamma} + \frac{1}{2}\right) \qquad \text{or} \qquad \frac{1-\alpha}{J} \geq \left(\frac{J-1+\alpha}{J}\right)\left(\frac{\psi}{2\gamma} + \frac{1}{2}\right)$$

The first inequality (left equation) and second inequality (right equation) are equivalent to:

$$J \geq (1-\alpha)\left(\frac{2\gamma}{\gamma-\psi}+1\right) \qquad\qquad J \leq (1-\alpha)\left(\frac{2\gamma}{\gamma+\psi}+1\right)$$

The first inequality is much more likely to be satisfied, which is true so long as J is large enough and $2\gamma/(\gamma-\psi)$ is not too large.

A.1.6 Prediction (1.1)

The first order conditions of the single-round maximization are:

$$\begin{aligned}
&\left(\frac{\gamma}{J-1+\alpha}\right) u_j'(q_j^A) = \lambda_{1R} && \text{for } j \in \{1,\ldots,J-1\} \\
&\left(\frac{\gamma}{J-1+\alpha}\right) \alpha u_j'(q_j^A) = \lambda_{1R} && \text{for } j = J \\
&\kappa G^A = \lambda_{1R}
\end{aligned}$$

where λ_{1R} is the Lagrange multiplier of the budget constraint in a single-round system.

The ratio in marginal utilities between localities is:

$$\begin{aligned}
&\text{between } j \text{ and } j'\text{:} \quad \frac{u_j'(q_j^A)}{u_{j'}'(q_{j'}^A)} = 1 && \forall \quad j, j' \in \{1,\ldots,J-1\} \\
&\text{between } j \text{ and } J\text{:} \quad \frac{u_j'(q_j^A)}{u_J'(q_J^A)} = \alpha && \forall \quad j \in \{1,\ldots,J-1\}
\end{aligned} \tag{A.3}$$

Equation (A.3) implies that $u_j'(q_j^A) < u_J'(q_J^A)$. Since $u_j(\cdot)$ is strictly increasing and strictly concave and $u_J'(q)/u_j'(q)$ is not too large, this implies that $q_j^A > q_J^A$.

A.1.7 Prediction (1.2)

The first order conditions of the two-round maximization are:

$$\begin{aligned}
&\left(\frac{\gamma}{J-1+\alpha}\right)\left(1+\frac{(1-\alpha)\gamma}{\psi J}\right)u'(q_j^A)=\lambda_{2R} && \text{for } j\in\{1,\ldots,J-1\}\\
&\left(\frac{\gamma}{J-1+\alpha}\right)\left(\alpha+\frac{(1-\alpha)\gamma}{\psi J}\right)u'(q_j^A)=\lambda_{2R} && \text{for } j=J\\
&\kappa G^A=\lambda_{2R}
\end{aligned}$$

where λ_{2R} is the Lagrange multiplier of the budget constraint in a two-round system.

The ratio in marginal utilities between localities is:

$$\begin{aligned}
&\text{between } j \text{ and } j'\text{:}\quad \frac{u_j'(q_j^A)}{u_{j'}'(q_{j'}^A)}=1 && \forall\quad j,j'\in\{1,\ldots,J-1\}\\
&\text{between } j \text{ and } J\text{:}\quad \frac{u_j'(q_j^A)}{u_J'(q_J^A)}=\frac{\alpha+\frac{(1-\alpha)\gamma}{\psi J}}{1+\frac{(1-\alpha)\gamma}{\psi J}} && \forall\quad j\in\{1,\ldots,J-1\}
\end{aligned}\tag{A.4}$$

Equation (A.4) implies that $u_j'(q_j^A)<u_J'(q_J^A)$. Since $u_j(\cdot)$ is strictly increasing and strictly concave and $u_J'(q)/u_j'(q)$ is not too large, this implies that $q_j^A>q_J^A$.

A.1.8 Comparing single- to two-round elections

I first establish three lemmas.

Lemma A.4. $\dfrac{u_j'(q_j^{1R})}{u_j'(q_j^{2R})}\dfrac{G^{2R}}{G^{1R}}>1$ *for all* $j\in\{1,\ldots,J\}$.

Proof. For $j\in\{1,\ldots,J-1\}$, combining the first round first order conditions in Appendix A.1.6:

$$\frac{u_j'(q_j^{1R})}{u_j'(q_j^{2R})}\frac{G^{2R}}{G^{1R}}=1+\frac{(1-\alpha)\gamma}{\psi J}>1$$

For $j = J$, combining the second round first order conditions in Appendix A.1.7:

$$\frac{u'_J(q_J^{1R})}{u'_J(q_J^{2R})}\frac{G^{2R}}{G^{1R}} = 1 + \frac{(1-\alpha)\gamma}{\alpha\psi J} > 1$$

□

LEMMA A.5. $\dfrac{u'_j(q_j^{1R})}{u'_j(q_j^{2R})} < \dfrac{u'_J(q_J^{1R})}{u'_J(q_J^{2R})}$

Proof. Comparing the ratio of marginal utilities in equations (A.3) and (A.4), the ratio is smaller in the single-round system compared to the two-round system:

$$\frac{u'_j(q_j^{1R})}{u'_J(q_J^{1R})} < \frac{u'_j(q_j^{2R})}{u'_J(q_J^{2R})} \iff \alpha < \frac{\alpha + \frac{(1-\alpha)\gamma}{\psi J}}{1 + \frac{(1-\alpha)\gamma}{\psi J}}$$

which is true because $\alpha < 1$. □

LEMMA A.6. *If $q_j^{1R} > q_j^{2R}$ for one $j \neq J$ then $q_{j'}^{1R} > q_{j'}^{2R}$ for all other $j' \in \{1, \ldots, J-1\}$.*

Proof. If $q_j^{1R} > q_j^{2R}$, then $u'_j(q_j^{1R}) < u'_j(q_j^{2R})$ because $u_j(\cdot)$ is strictly concave. The first order conditions in Appendices A.1.6 and A.1.7 establish that the marginal utilities between all $j, j' \in \{1, \ldots, J-1\}$ are equal. Then we must have that $u_{j'}(q_{j'}^{1R}) < u_{j'}(q_{j'}^{2R})$ and that $q_{j'}^{1R} > q_{j'}^{2R}$. □

Proof of allocations in locality J.– I prove that $q_J^{1R} < q_J^{2R}$.

Proof. I prove by contradiction. Assume that $q_J^{1R} \geq q_J^{2R}$. Then $u'_J(q_J^{1R}) \leq u'_J(q_J^{2R})$. By lemma A.4, we must have that $G^{2R} > G^{1R}$. To satisfy the budget constraint, we must have that $q_j^{1R} < q_j^{2R}$ for some $j \neq J$ and, by lemma A.6, for all $j \neq J$. Then $u'_j(q_j^{1R}) > u'_j(q_j^{2R})$. However, this violates lemma A.5, and so we must have $q_J^{1R} < q_J^{2R}$. □

Proof of overall budget.– I prove that $G^{1R} < G^{2R}$.

Proof. I prove by contradiction. Assume that $G^{1R} \geq G^{2R}$. Since $q_J^{1R} < q_J^{2R}$, to satisfy the budget constraint, we must have that $q_j^{1R} > q_j^{2R}$ for some $j \neq J$ and, by lemma A.6, for all $j \neq J$. Then $u_j'(q_j^{1R}) < u_j'(q_j^{2R})$. However, this violates lemma A.4, and so we must have $G^{1R} < G^{2R}$. □

Proof of allocations in other localities..– I show that $q_j^{1R} \lesseqgtr q_j^{2R}$.

Proof. From the first order conditions in Appendices A.1.6 and A.1.7 and since $G^{1R} < G^{2R}$, we have that $\lambda_{1R} < \lambda_{2R}$. Then, for $j \in \{1, \ldots, J-1\}$:

$$\frac{u_j'(q_j^{1R})}{u_j'(q_j^{2R})} < 1 + \frac{(1-\alpha)\gamma}{\psi J} \implies \frac{u_j'(q_j^{1R})}{u_j'(q_j^{2R})} \gtreqless 1$$

which implies that $q_j^{1R} \lesseqgtr q_j^{2R}$. □

A.1.9 Candidate C's budget

In general, for every set of utility functions $(u_1(\cdot), \cdots, u_J(\cdot))$, there exists a G^C such that G^C is the highest offer in locality J. I show this for the case where $u_j(\cdot) = \beta_j \ln(\cdot)$ and for the two-round election (since q_J is higher in a two-round election).

The first order conditions with respect to q_J and G are:

$$\frac{\gamma}{J-1+\alpha}\left(\alpha + \frac{(1-\alpha)\gamma}{\psi J}\right)\frac{\beta_J}{q_J} = \lambda_{2R}$$
$$\kappa G = \lambda_{2R}$$

We can write G as a function of q_J:

$$G = \frac{1}{\kappa}\frac{\gamma}{J-1+\alpha}\left(\alpha + \frac{(1-\alpha)\gamma}{\psi J}\right)\frac{\beta_J}{q_J}$$

Since $q_J < q_j$ for all $j \neq J$ (prediction (1.2)), then $q_J < G/J$, implying:

$$q_J < \left(\frac{1}{\kappa J}\frac{\gamma}{J-1+\alpha}\left(\alpha + \frac{(1-\alpha)\gamma}{\psi J}\right)\beta_J\right)^{1/2} \equiv \Gamma$$

So long as $G^C > \Gamma$, then candidate C's allocation to locality J will be the highest offer there. This will be true so long as κ or J is large enough and γ/ψ is not too large.

A.1.10 Relaxing the assumption that $\alpha > 0$

Assume $\alpha = 0$, so that candidate C receives all the votes in locality J. Then vote shares in the first round, or when all 3 candidates are present, are given by:

$$\begin{aligned}
\pi_1^A &= \left(\frac{J-1}{J}\right)\left(\frac{1}{2}+\psi\delta\right)+\frac{\psi}{J}\sum_{j=1}^{J-1}\Delta u_j^{AB} \\
\pi_1^B &= \left(\frac{J-1}{J}\right)\left(\frac{1}{2}-\psi\delta\right)+\frac{\psi}{J}\sum_{j=1}^{J-1}\Delta u_j^{BA} \\
\pi_1^C &= \frac{1}{J}
\end{aligned}$$

and the probability of attaining a vote share above θ is given by:

$$\begin{aligned}
Pr\left(\pi_1^A \geq \theta\right) &= \frac{1}{2}+\frac{\gamma}{\psi}\left[\frac{1}{2}-\left(\frac{J}{J-1}\right)\theta+\left(\frac{\psi}{J-1}\right)\sum_{j=1}^{J-1}\Delta u_j^{AB}\right] \\
Pr\left(\pi_1^B \geq \theta\right) &= \frac{1}{2}+\frac{\gamma}{\psi}\left[\frac{1}{2}-\left(\frac{J}{J-1}\right)\theta+\left(\frac{\psi}{J-1}\right)\sum_{j=1}^{J-1}\Delta u_j^{BA}\right]
\end{aligned}$$

Vote shares for the second round with candidates A and B are the same as when $\alpha > 0$.

Equilibrium strategies in a single-round election.– The maximization is:

$$\max_{G^A,\mathbf{q}^A=\left(q_1^A,\ldots,q_J^A\right)} \frac{1}{2}+\frac{\gamma}{J-1}\sum_{j=1}^{J-1}\Delta u_j^{AB}-\frac{1}{2}\kappa\left(G^A\right)^2 \qquad \text{s.t.} \quad \sum_j q_j^A \leq G^A$$

The first order conditions are:

$$\begin{aligned}
&\frac{\gamma}{J-1}u'(q_j^A) = \lambda_{1R}' && \forall\ j \in \{1,\ldots,J-1\} \\
&\kappa G^A = \lambda_{1R}'
\end{aligned}$$

where λ'_{1R} is the Lagrange multiplier of the budget constraint in a single-round system when $\alpha = 0$. I show that in equilibrium, the optimal strategy is to allocate $q_J^A = 0$.

Say $q_J^A > 0$. Consider the following deviation: $\left(q_J^A\right)' = q_J^A - \epsilon$ and $\left(q_k^A\right)' = q_k^A + \epsilon$ for some $k \neq J$ and $\epsilon > 0$. Candidate C's vote share is unchanged, so the threshold for winning remains $\frac{1}{2}\left(1 - \frac{1}{J}\right)$. The net change in the probability of winning is given by:

$$\left[\frac{1}{2} + \frac{\gamma}{J-1}\left(\Delta\left(u_k^{AB}\right)' + \sum_{j\neq k,J} \Delta u_j^{AB}\right)\right] - \left[\frac{1}{2} + \frac{\gamma}{J-1}\left(\Delta u_k^{AB} + \sum_{j\neq k,J} \Delta u_j^{AB}\right)\right]$$
$$= \frac{\gamma}{J-1}\left[u_k\left(\left(q_k^A\right)'\right) - u_k\left(q_k^A\right)\right] > 0$$

where the last line follows because $\epsilon > 0$ and $u_k(\cdot)$ is strictly increasing. There is a deviation that strictly increases the probability of winning. As a result, any $q_j^A > 0$ cannot be optimal so the optimal strategy is to allocate $q_J^A = 0$.

Equilibrium strategies in a two-round election.– The maximization is:

$$\max_{G^A, \mathbf{q}^A = \left(q_1^A, \ldots, q_J^A\right)} \left(\frac{1}{2} + \frac{\gamma}{\psi}\left[-\left(\frac{1}{J-1}\right)\frac{1}{2} + \left(\frac{\psi}{J-1}\right)\sum_{j=1}^{J-1} \Delta u_j^{AB}\right]\right)$$
$$+ \frac{\gamma}{\psi}\left(\frac{1}{J-1}\right)\left[\frac{1}{2} + \frac{\gamma}{J}\sum_{j=1}^{J} \Delta u_j^{AB}\right] - \frac{1}{2}\kappa\left(G^A\right)^2 \quad \text{s.t.} \quad \sum_j q_j^A \leq G^A$$

The first order conditions are:

$$\left(\frac{\gamma}{J-1}\right)\left(1 + \frac{\gamma}{\psi J}\right) u'(q_j^A) = \lambda'_{2R} \qquad \text{for } j \in \{1, \ldots, J-1\}$$
$$\left(\frac{\gamma}{J-1}\right)\left(\frac{\gamma}{\psi J}\right) u'(q_j^A) = \lambda'_{2R} \qquad \text{for } j = J$$
$$\kappa G^A = \lambda'_{2R}$$

where λ'_{2R} is the Lagrange multiplier of the budget constraint in a two-round system when $\alpha = 0$.

From this, we can see that $q_J^A > 0$. Thus, $q_J^{1R} < q_J^{2R}$, and prediction (1.3) holds.

Similarly, $G^{1R} < G^{2R}$. Assume not. To satisfy the budget constraint, we need that $q_j^{1R} > q_j^{2R}$ for some $j \neq J$ and by extension all $j \neq J$. Then $u_j'(q_j^{1R}) < u_j'(q_j^{2R})$. However, then $\frac{u_j'(q_j^{1R})}{u_j'(q_j^{2R})} \frac{G^{2R}}{G^{1R}} < 1$, which violates the first order conditions. Thus, $G^{1R} < G^{2R}$ and prediction (1.4) holds.

We also have that $q_j^{1R} \lesseqgtr q_j^2$. From the first order conditions, we have that $\frac{u_j'(q_j^{1R})}{u_j'(q_j^{2R})} < 1 + \frac{\gamma}{\psi J}$ and so $q_j^{1R} \lesseqgtr q_j^2$. Predictions (1.3) and (1.5) follow.

A.2 Data Appendix

Table A.1 Variables used to construct the equipment index

Variable	2000	2001-2003	2004	2005-2006	2007	2008	2009-2011	2012-2016
VCR/DVD	x	x	x	x	x	x	x	x
TV	x	x	x	x	x	x	x	x
satellite dish	x	x	x	x	x	x	x	x
satellite dish with internet			x	x				
overhead projector	x	x	x	x	x	x	x	x
projector			x	x				x
fax		x	x	x				x
copier		x	x	x	x	x	x	x
stereo/sound system	x	x	x	x				x
camera/camcorder			x	x				x
drinking fountain			x	x				
special needs accom.	x	x	x	x	x	x		
fan		x	x	x				
air conditioning		x	x	x				
computers	x	x	x	x	x	x	x	x
printer	x	x	x	x	x	x	x	x
local network	x	x	x					
internet	x	x	x		x	x	x	x
broadband						x	x	x

Table A.2 Variables used to construct the infrastructure index

Variable	1997	1998	1999-2000	2001-2003	2004-2006	2007-2008	2009-2011	2012-2016
principal office	x	x	x	x	x	x	x	x
secretary office	x	x	x	x	x			x
teacher lounge		x	x	x	x	x	x	x
teacher housing								x
library	x	x	x	x	x	x	x	x
reading room				x	x		x	x
video library / room			x	x	x			
toy library					x			
auditorium					x			x
solarium					x			
science lab	x	x	x	x	x	x	x	x
computer lab	x	x	x	x	x	x	x	x
other lab		x	x	x	x			
kitchen	x	x	x	x	x	x	x	x
food pantry	x	x	x	x	x			x
cafeteria	x	x	x	x	x			x
warehouse				x	x			x
schoolyard	x	x	x	x	x			x
green area								x
sports field	x	x	x	x	x	x	x	x
pool				x	x			
gymnasium					x			
playground	x	x	x	x	x	x	x	x
laundry					x			x
sanitation	x	x	x	x	x	x	x	x
special needs accomm.	x	x	x	x	x	x	x	x
classrooms	x	x	x	x	x	x	x	x

A.3 Additional Figures and Tables

Figure A.3 Relationship between municipality population and number of registered voters

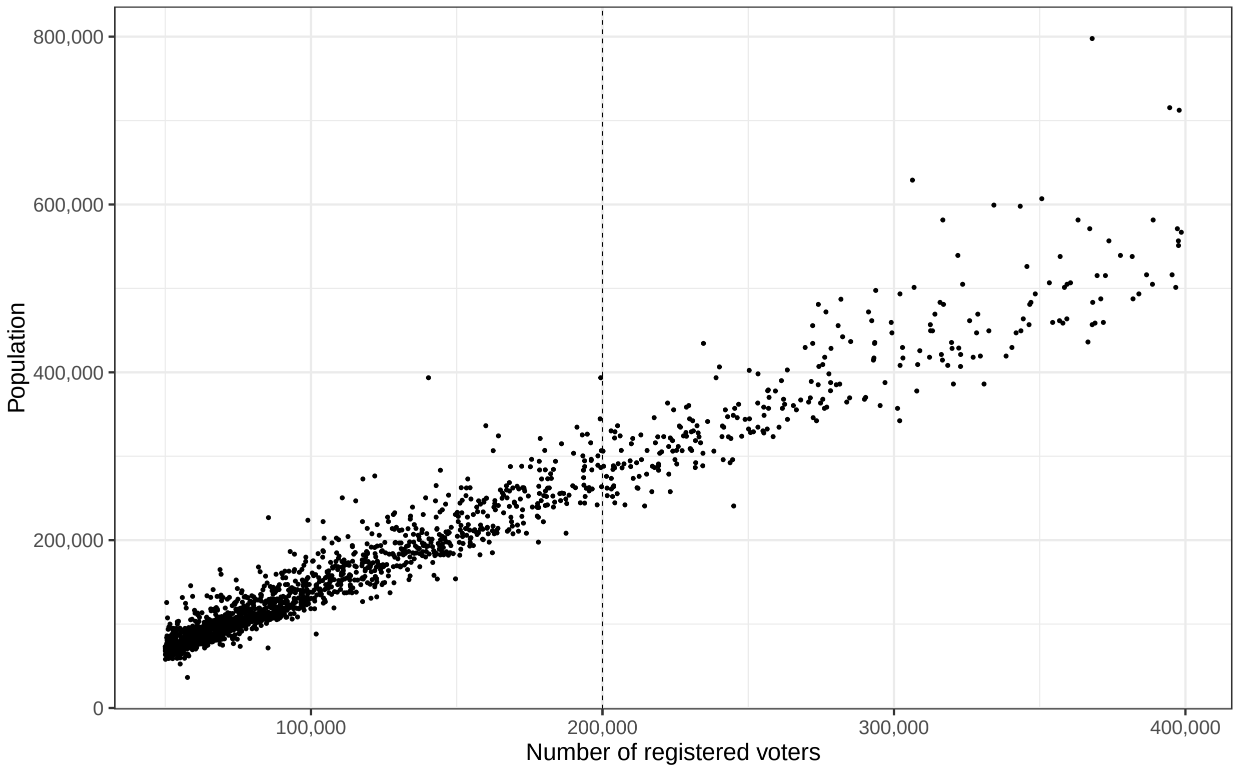

The vertical axis is municipality population in the most recent census prior to the election. Plot includes only elections with between 50,000 and 400,000 registered voters (6.0% of the universe of elections).

Figure A.4 Regression discontinuity plots of the probability of falling above/below other policy thresholds

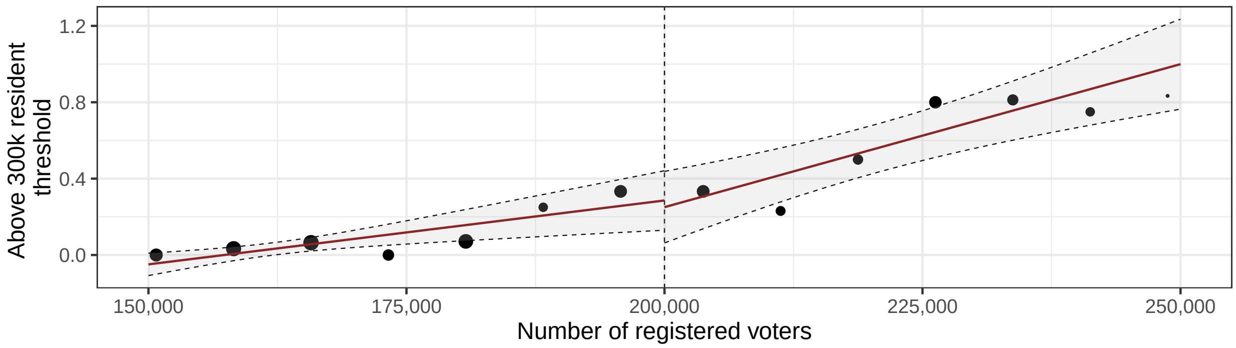

(a) Threshold: 300,000 inhabitants

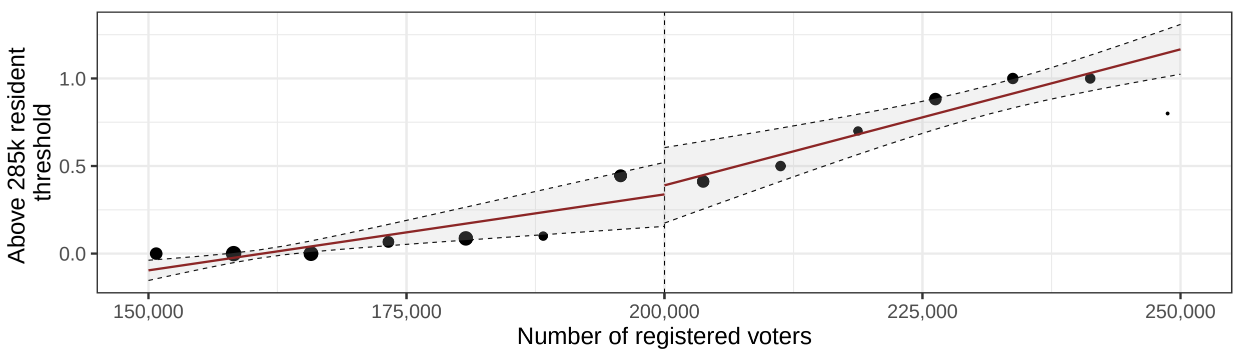

(b) Threshold: 285,714 inhabitants

The vertical axis is the fraction of elections above (a) the 300,000 resident threshold and (b) the 285,714 resident threshold. At 300,000 residents, a salary cap for municipal legislators comes into effect. At 285,714 residents, the size of the legislature changes. In each panel, each point plots an average value within a 7,500 voter bin. Diameter of the points is proportional to the number of observations. Confidence intervals (dashed lines) represent the 95% confidence intervals of a local linear regression (solid red line) with standard errors clustered at the municipality level.

Figure A.5 Regression discontinuity plots of pre-treatment population density

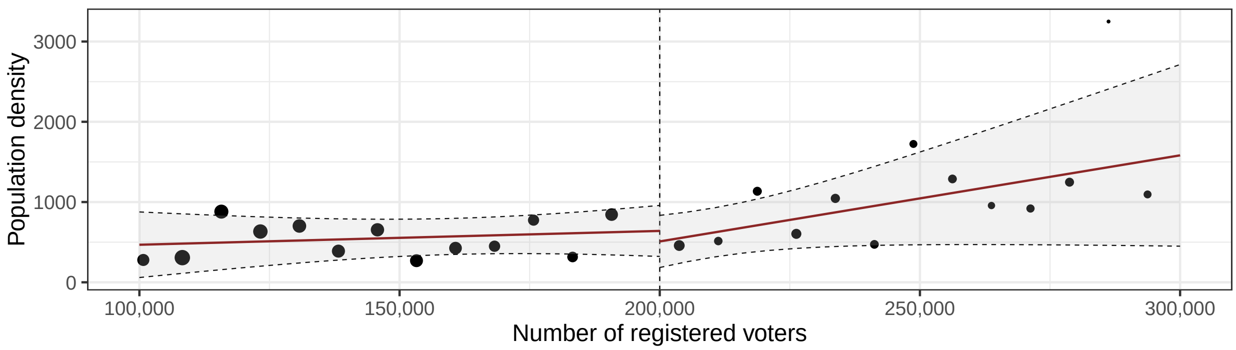

(a) Measured prior to the 1988 Constitution

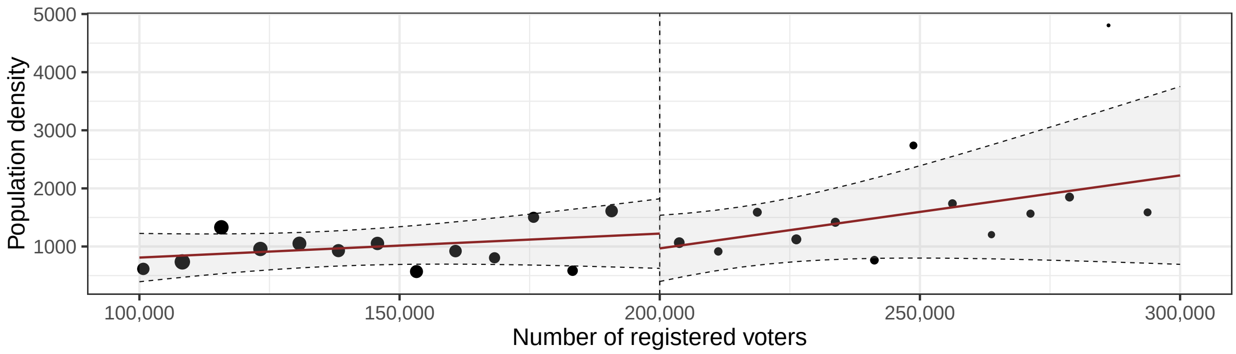

(b) Measured prior to the most recent single-round

Population density measured (a) in the 1980 census or (b) in the census prior to the most recent year in a single-round system or in the 1991 census (b). In each panel, each point plots an average value within a 7,500 voter bin. Diameter of the points is proportional to the number of observations. Confidence intervals (dashed lines) represent the 95% confidence intervals of a local linear regression (solid red line) with standard errors clustered at the municipality level.

Figure A.6 Regression discontinuity coefficients on pre-treatment population density at different bandwidths

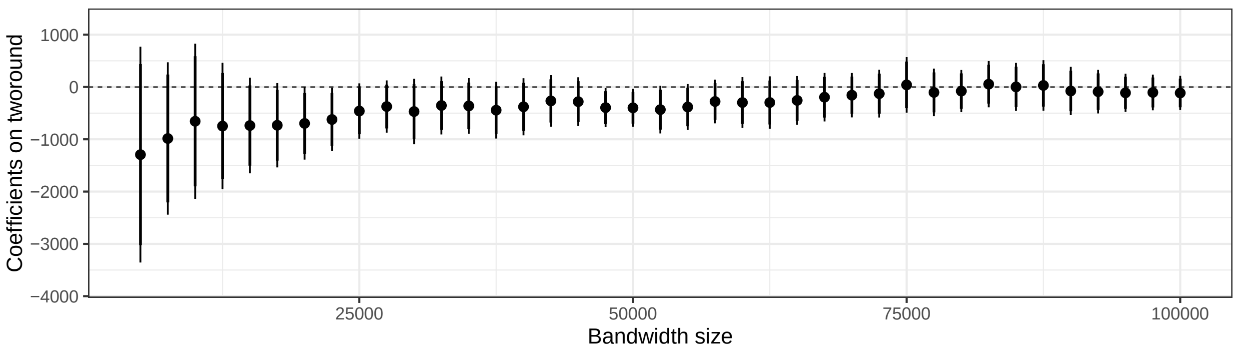

(a) Measured prior to the 1988 Constitution

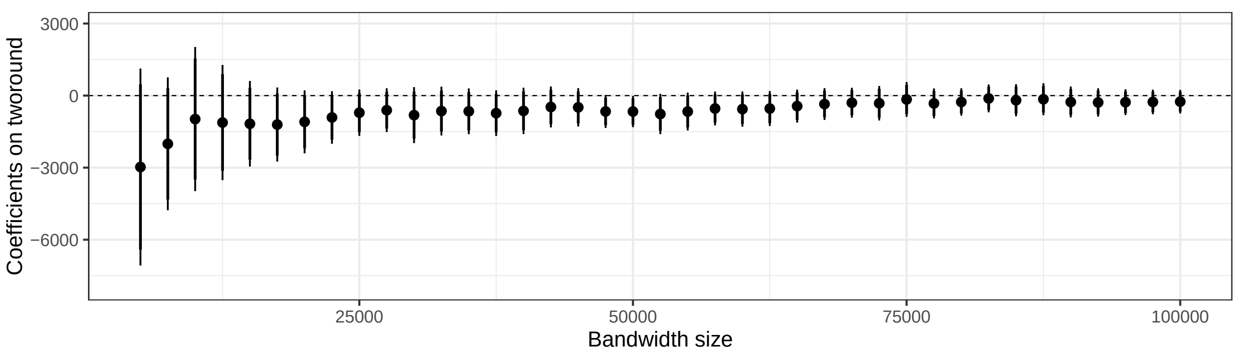

(b) Measured prior to the most recent single-round

Population density measured (a) in the 1980 census or (b) in the census prior to the most recent year in a single-round system or in the 1991 census. The thicker vertical lines represent the 90% confidence interval and the thinner vertical lines represent the 95% confidence interval. *Estimation method:* Local linear regression with the specified voter bandwidth and election-year fixed effects. Standard errors clustered at the municipality level. *Source:* 1980, 1991, 2000, and 2010 Demographic Census.

Figure A.7 Compliance with treatment assignment

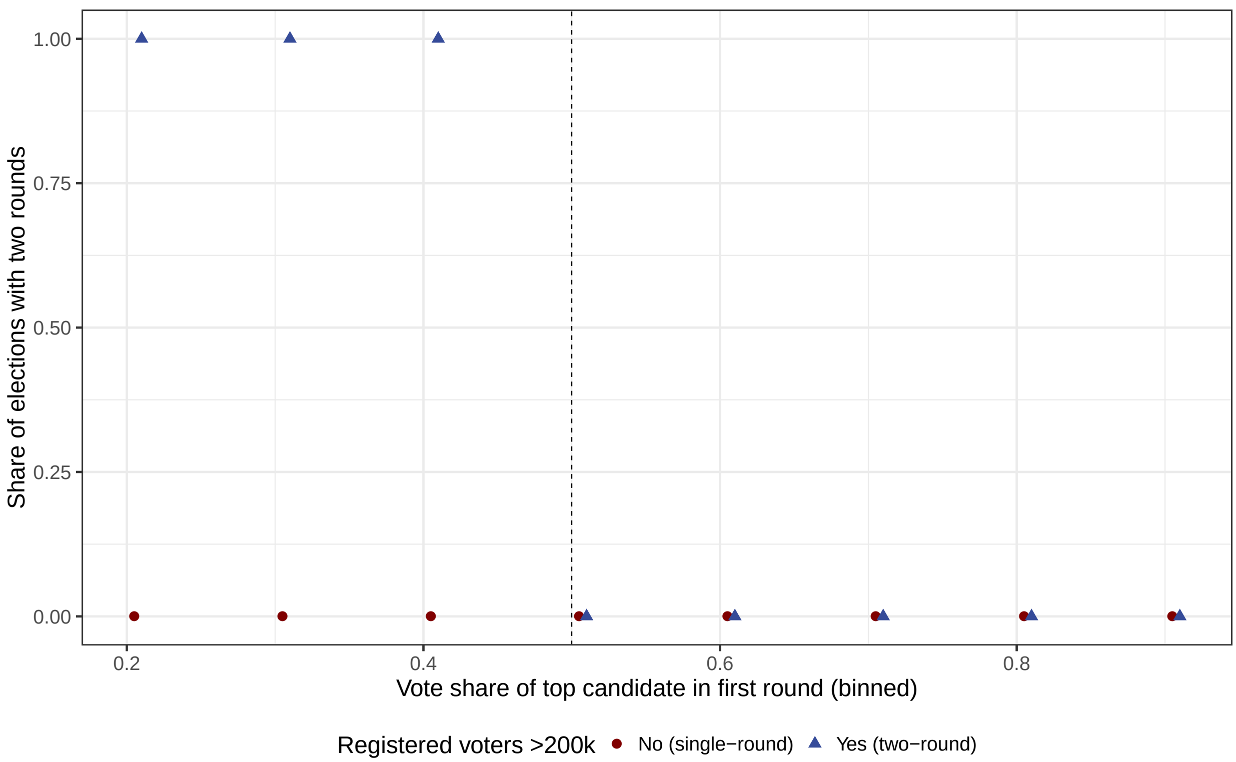

The vertical axis is the fraction of elections that held two rounds of elections. The horizontal axis is the vote share of the top candidate in the first round. Municipalities below the 200,000 registered voter threshold (and thus should always hold one round) are denoted by red circles. Municipalities above the 200,000 registered voter threshold (and thus should hold two rounds if no candidate receives 50% in the first round) are denoted by blue triangles. Bin sizes are 10% vote share.

Figure A.8 Regression discontinuity plot of probability of treatment in the previous election

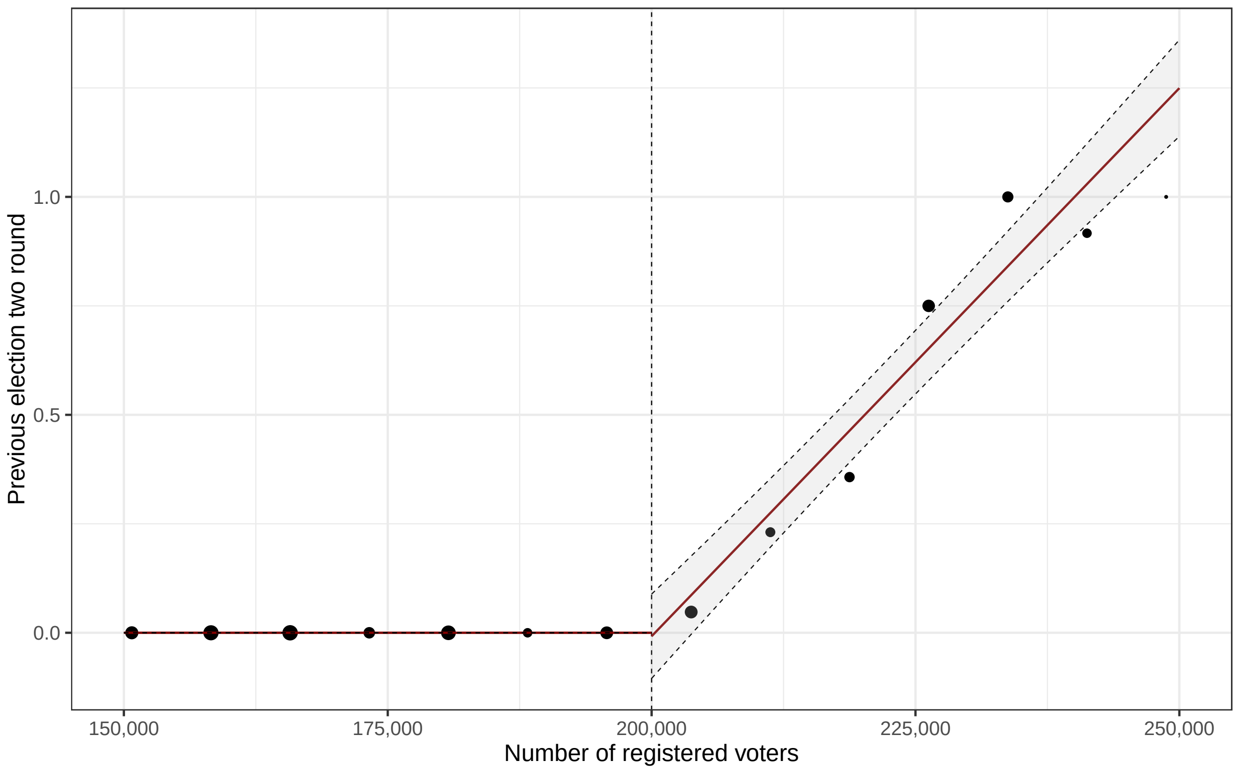

The vertical axis is the probability that the *previous* election was a two-round election. The horizontal axis is the number of registered voters in the *current* election. In each panel, each point plots an average value within a 7,500 voter bin. Diameter of the points is proportional to the number of observations. Confidence intervals (dashed lines) represent the 95% confidence intervals of a local linear regression (solid red line) with standard errors clustered at the municipality level. Note that because all observations to the left of the threshold are 0, there are no standard errors.

Table A.9 Regression discontinuity estimates on resources in municipal schools, by year in the term

	Mean level of resources		*Standard deviation in resources*	
	Equipment	Infrastructure	Equipment	Infrastructure
Panel A: First year in the electoral term				
TwoRound	0.077^{**}	0.052	-0.013	-0.021
	(0.039)	(0.035)	(0.011)	(0.018)
Observations	197	227	197	227
Panel B: Second year in the electoral term				
TwoRound	0.091^{**}	0.059^{*}	-0.018^{*}	-0.022
	(0.039)	(0.033)	(0.011)	(0.017)
Observations	197	228	197	228
Panel C: Third year in the electoral term				
TwoRound	0.087^{**}	0.059^{*}	-0.026^{**}	-0.022
	(0.037)	(0.034)	(0.011)	(0.016)
Observations	197	228	197	228
Panel D: Fourth year in the electoral term				
TwoRound	0.070^{**}	0.059^{*}	-0.015	-0.020
	(0.033)	(0.034)	(0.010)	(0.017)
Bandwidth size	50,000	50,000	50,000	50,000
Observations	229	229	229	229

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Mayoral terms are for four years. Each panel displays the estimate separately for the 1st year (*Panel A*), 2nd year (*Panel B*), 3rd year (*Panel C*), and 4th year (*Panel D*) of the term. *Equipment* and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. The first two columns (*Mean level of resources*) have as the dependent variable the mean index level across schools in the municipality. The last two columns (*Standard deviation in resources*) have as the dependent variable the standard deviation in the index across schools in the municipality. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Table A.10 Regression discontinuity estimates on the geographic concentration of voters, with number of candidates as a control

Panel A: Concentration indices of voters for specific candidates				
	Coefficient of variation	Fractionalization	Entropy	
TwoRound	-0.005	-0.010^{**}	-0.009^{*}	
	(0.003)	(0.005)	(0.005)	
Single-round mean	0.019	0.027	0.030	
Observations	264	264	264	
Panel B: Standard deviation in vote shares for each candidate				
	1st place candidate	2nd place candidate	3rd place candidate	4th place candidate
TwoRound	-0.016^{**}	-0.012	-0.010	-0.002
	(0.007)	(0.008)	(0.007)	(0.004)
Single-round mean	0.080	0.075	0.042	0.023
Observations	264	264	251	216

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Panel A: overall concentration of voters for specific candidates, as measured by concentration indices (coefficient of variation, fractionalization, and entropy) of vote counts in polling stations. *Panel B:* candidate-level concentration of voters, measured by standard deviation in a candidate's vote shares (for the 1st-4th place candidate) across polling stations. Vote shares are from the first round. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Number of candidates included as a control. Standard errors clustered at the municipality level.

Table A.11 Regression discontinuity estimates on resources in municipal schools, using z-scores

	Mean level of resources		*Standard deviation in resources*	
	Equipment	Infrastructure	Equipment	Infrastructure
TwoRound	0.079**	0.069*	−0.014	−0.007
	(0.033)	(0.037)	(0.009)	(0.017)
Single-round mean	0.724	0.739	0.120	0.146
Observations	820	912	820	912

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Equipment and *Infrastructure* are indices constructed by taking the z-score of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. The first two columns (*Mean level of resources*) have as the dependent variable the mean index level across schools in the municipality. The last two columns (*Standard deviation in resources*) have as the dependent variable the standard deviation in the index across schools in the municipality. *Estimation method:* Local linear regression with a 50,000 voter bandwidth. Standard errors clustered at the municipality level.

Figure A.12 Regression discontinuity coefficients on overall concentration of voters for specific candidates at different bandwidths

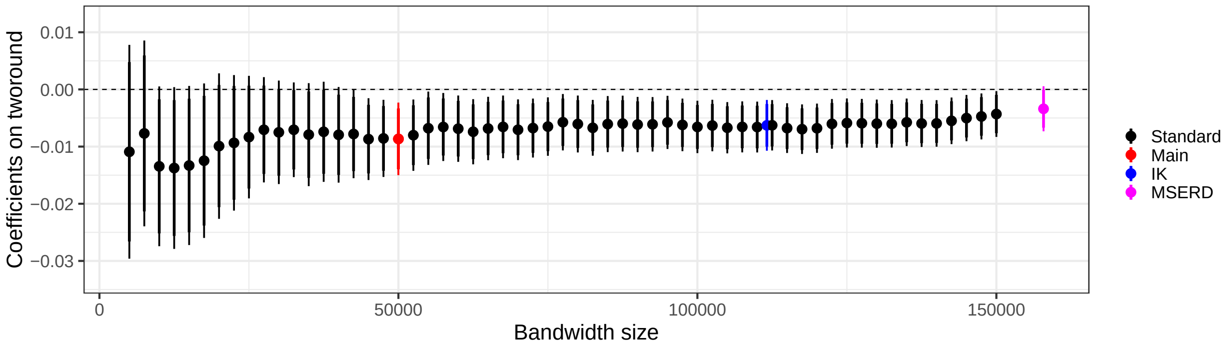

(a) Coefficient of variation

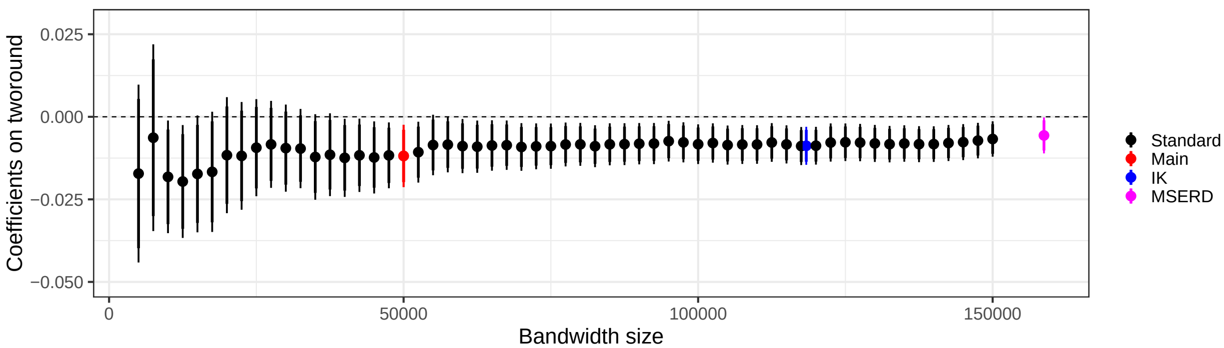

(b) Fractionalization

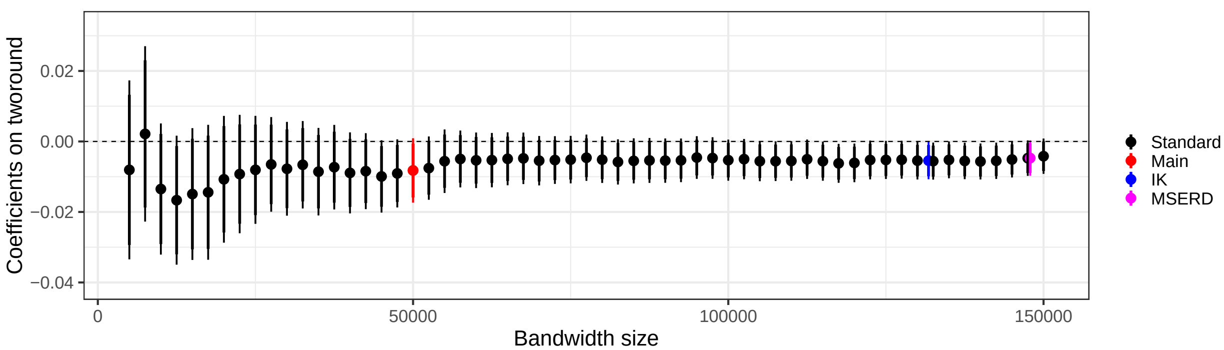

(c) Entropy

Overall concentration of voters for specific candidates, as measured by (a) *Coefficient of variation*, (b) *Fractionalization*, and (c) *Entropy*, using vote counts in polling stations. Vote shares are from the first round. The thicker vertical lines represent the 90% confidence interval and the thinner vertical lines represent the 95% confidence interval. *Estimation method:* Local linear regression with election-year fixed effects and with the specified voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Figure A.13 Regression discontinuity coefficients on the candidate-level concentration in voters at different bandwidths

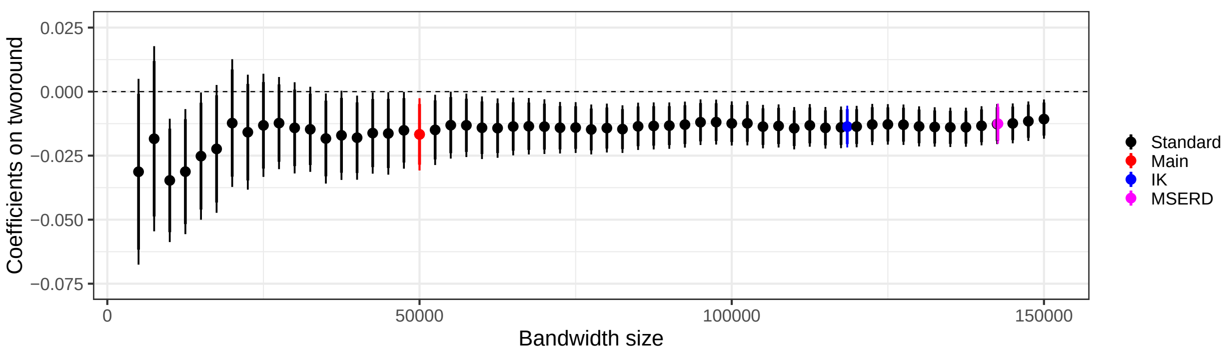

(a) Standard deviation in votes for 1st place candidate

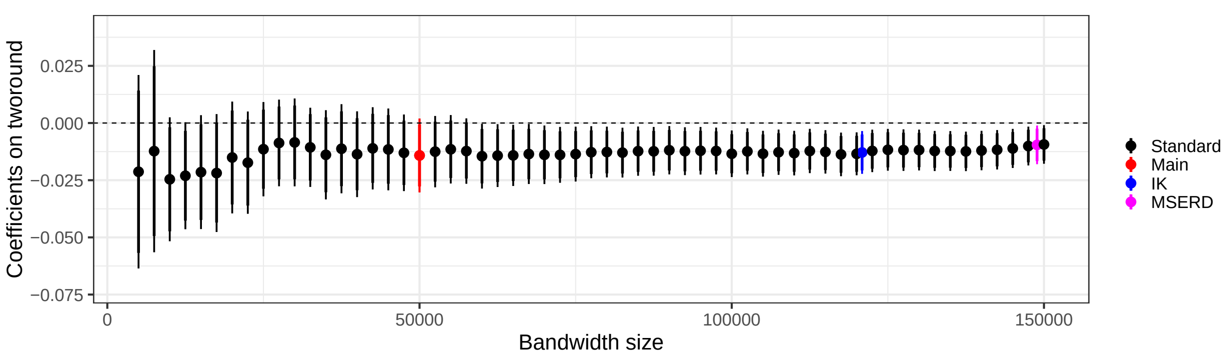

(b) Standard deviation in votes for 2nd place candidate

Standard deviation in a candidate's vote counts across polling stations, for the (a) 1st place and (b) 2nd place candidate. Vote shares are from the first round. The thicker vertical lines represent the 90% confidence interval and the thinner vertical lines represent the 95% confidence interval. *Estimation method:* Local linear regression with election-year fixed effects and with the specified voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Figure A.14 Regression discontinuity coefficients on overall level of resources in municipal schools at different bandwidths

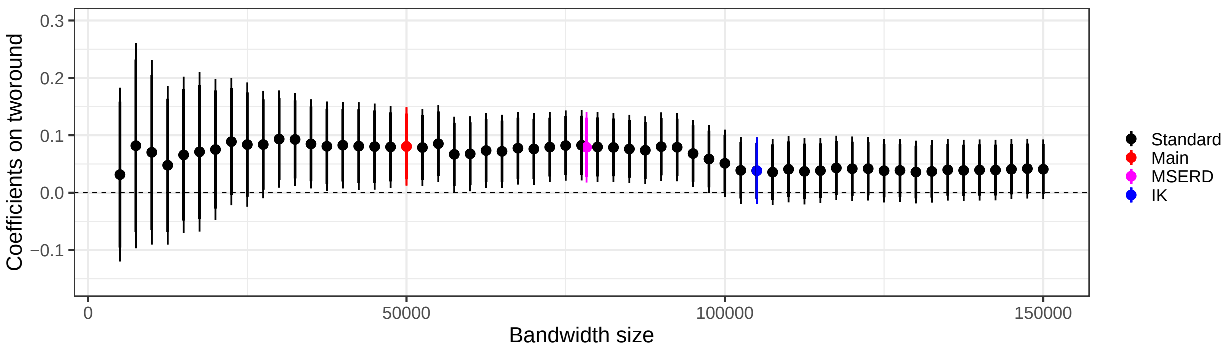

(a) Equipment, mean level of resources

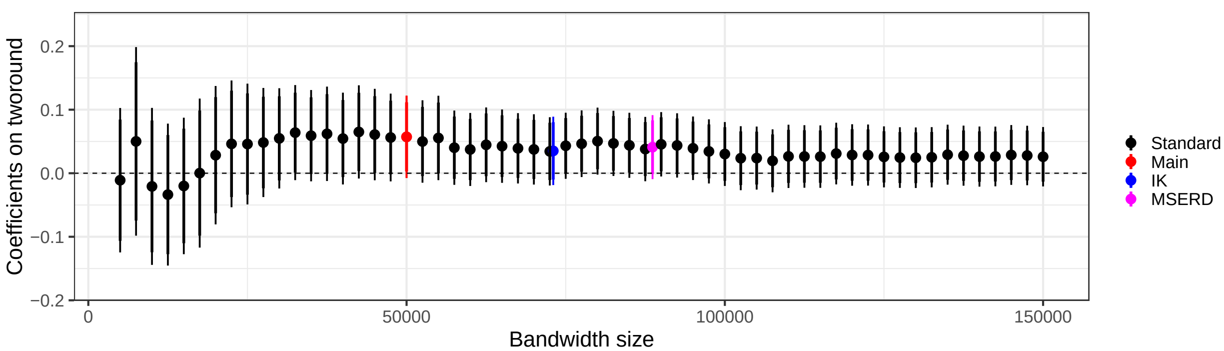

(b) Infrastructure, mean level of resources

Equipment and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. *Mean level of resources* is the mean index level across schools in the municipality for (a) equipment and (b) infrastructure. The thicker vertical lines represent the 90% confidence interval and the thinner vertical lines represent the 95% confidence interval. *Estimation method:* Local linear regression with election-year fixed effects and with the specified voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Figure A.15 Regression discontinuity coefficients on distribution of resources in municipal schools at different bandwidths

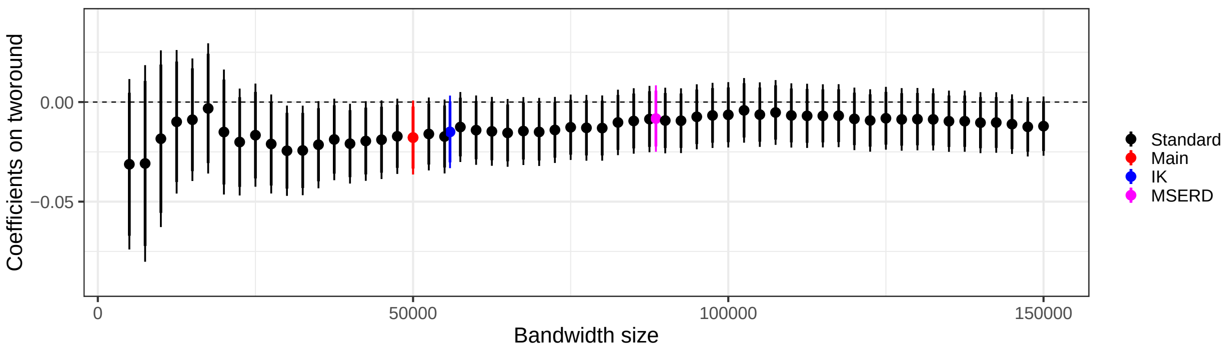

(a) Equipment, standard deviation in resources

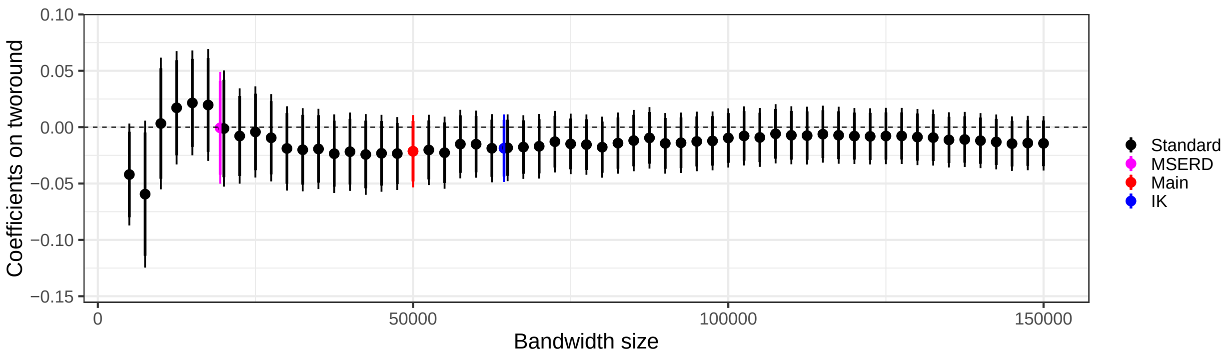

(b) Infrastructure, standard deviation in resources

Equipment and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. *Standard deviation in resources* is the standard deviation in the index across schools in the municipality for (a) equipment and (b) infrastructure. The thicker vertical lines represent the 90% confidence interval and the thinner vertical lines represent the 95% confidence interval. *Estimation method:* Local linear regression with election-year fixed effects and with the specified voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Figure A.16 Regression discontinuity coefficients on municipal education outcomes at different bandwidths

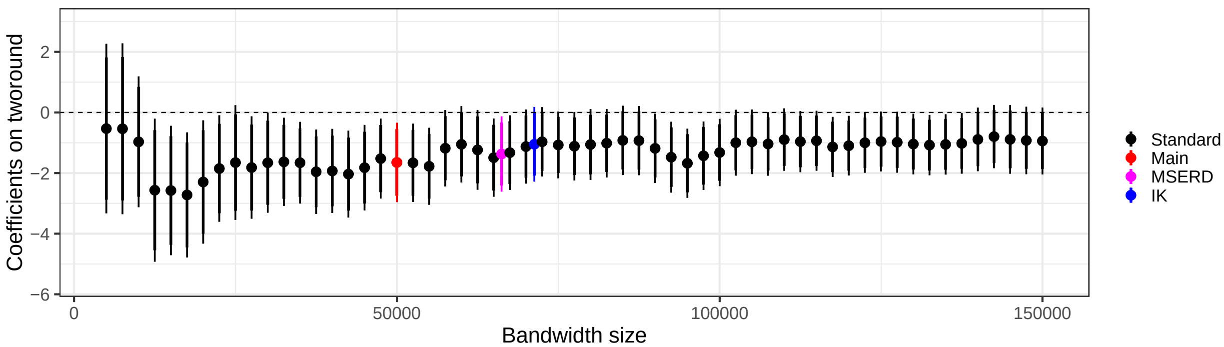

(a) Drop-out rate

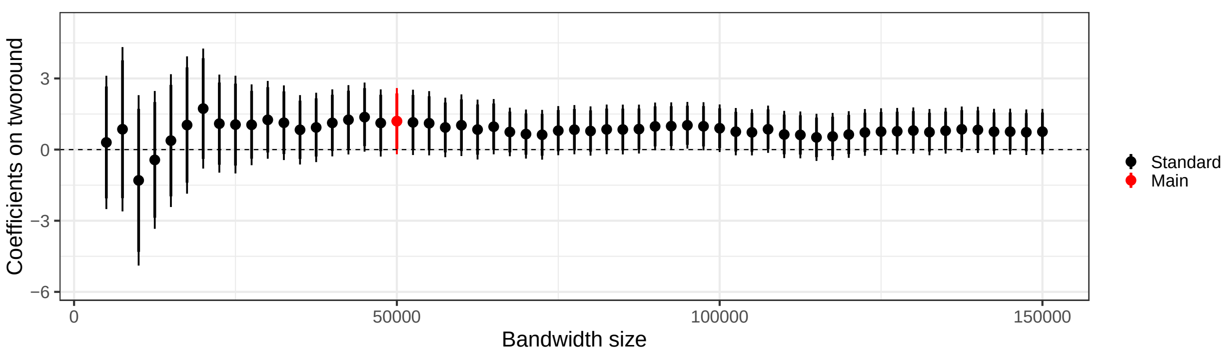

(b) Elementary literacy rate

Drop-out rate is from the School Census. It is the mean rate across schools in the municipality. *Elementary literacy rate* is from the 2000 and 2010 Demographic Census. It is the literacy rate of cohorts who are of elementary school age during the mayoral term. The thicker vertical lines represent the 90% confidence interval and the thinner vertical lines represent the 95% confidence interval. IK and MSERD bandwidths not shown for *Elementary literacy rate*, as the bandwidth chosen was larger than the support. *Estimation method:* Local linear regression with election-year fixed effects and with the specified voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Table A.17 Regression discontinuity estimates on the geographic concentration of voters, without controls

Panel A: Concentration indices of voters for specific candidates				
	Coefficient of variation	Fractionalization	Entropy	
TwoRound	-0.008^{**}	-0.010^{**}	-0.007	
	(0.003)	(0.004)	(0.005)	
Single-round mean	0.019	0.027	0.030	
Observations	264	264	264	
Panel B: Standard deviation in vote shares for each candidate				
	1st place candidate	2nd place candidate	3rd place candidate	4th place candidate
TwoRound	-0.012^{*}	-0.011	-0.003	0.005
	(0.006)	(0.008)	(0.006)	(0.005)
Single-round mean	0.080	0.075	0.042	0.023
Observations	264	264	251	216

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Panel A: overall concentration of voters for specific candidates, as measured by concentration indices (coefficient of variation, fractionalization, and entropy) of vote counts in polling stations. *Panel B:* candidate-level concentration of voters, measured by standard deviation in a candidate's vote shares (for the 1st-4th place candidate) across polling stations. Vote shares are from the first round. *Estimation method:* Local linear regression with a 50,000 voter bandwidth. Standard errors clustered at the municipality level.

Table A.18 Regression discontinuity estimates on resources in municipal schools, without controls

	Mean level of resources		*Standard deviation in resources*	
	Equipment	Infrastructure	Equipment	Infrastructure
TwoRound	0.068^{*}	0.036	-0.019^{*}	-0.014
	(0.035)	(0.029)	(0.011)	(0.015)
Single-round mean	0.738	0.731	0.121	0.157
Observations	821	916	821	916

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Equipment and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. The first two columns (*Mean level of resources*) have as the dependent variable the mean index level across schools in the municipality. The last two columns (*Standard deviation in resources*) have as the dependent variable the standard deviation in the index across schools in the municipality. *Estimation method:* Local linear regression with a 50,000 voter bandwidth. Standard errors clustered at the municipality level.

Table A.19 Regression discontinuity estimates on municipal education outcomes, without controls

	Drop-out rate	Failing rate	Passing rate	Elem. literacy rate
TwoRound	-1.340^{*}	-0.051	1.291	2.918
	(0.755)	(1.167)	(1.686)	(2.030)
Single-round mean	3.211	8.645	88.283	91.445
Observations	913	912	913	178

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Drop-out rate, *Failing rate*, and *Passing rate* are from the School Census. They are the mean rate across schools in the municipality and should add up to 1 in each school. *Elem. literacy rate* is from the 2000 and 2010 Demographic Census. It is the literacy rate of cohorts who are of elementary school age during the mayoral term. *Estimation method:* Local linear regression with a 50,000 voter bandwidth. Standard errors clustered at the municipality level.

Table A.20 Placebo regression discontinuity estimates on the geographic concentration of voters, at 300,000 inhabitant threshold

Panel A: Concentration indices of voters for specific candidates				
	Coefficient of variation	Fractionalization	Entropy	
TwoRound	0.005	0.005	0.005	
	(0.004)	(0.005)	(0.005)	
Single-round mean	0.019	0.024	0.027	
Observations	471	471	471	
Panel B: Standard deviation in vote shares for each candidate				
	1st place candidate	2nd place candidate	3rd place candidate	4th place candidate
TwoRound	0.001	0.0001	0.001	−0.003
	(0.008)	(0.007)	(0.006)	(0.005)
Single-round mean	0.075	0.072	0.040	0.021
Observations	471	471	444	373

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

At 300,000 inhabitants, a 2000 constitutional amendment placing a cap on local legislator salaries comes into effect. *Panel A:* overall concentration of voters for specific candidates, as measured by concentration indices (coefficient of variation, fractionalization, and entropy) of vote counts in polling stations. *Panel B:* candidate-level concentration of voters, measured by standard deviation in a candidate's vote shares (for the 1st-4th place candidate) across polling stations. Vote shares are from the first round. Includes only elections after 2000. *Estimation method:* Local linear regression with election-year fixed effects and a 125,000 inhabitant bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Table A.21 Placebo regression discontinuity estimates on the geographic concentration of voters, at 285,714 inhabitant threshold

Panel A: Concentration indices of voters for specific candidates				
	Coefficient of variation	Fractionalization	Entropy	
TwoRound	0.004	0.005	0.007*	
	(0.003)	(0.004)	(0.004)	
Single-round mean	0.019	0.024	0.027	
Observations	423	423	423	
Panel B: Standard deviation in vote shares for each candidate				
	1st place candidate	2nd place candidate	3rd place candidate	4th place candidate
TwoRound	0.005	0.003	0.009*	0.006
	(0.005)	(0.006)	(0.005)	(0.004)
Single-round mean	0.075	0.071	0.040	0.022
Observations	424	423	400	331

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

At 285,714 inhabitants, a 2004 constitutional amendment changing the size of the local legislature comes into effect. *Panel A:* overall concentration of voters for specific candidates, as measured by concentration indices (coefficient of variation, fractionalization, and entropy) of vote counts in polling stations. *Panel B:* candidate-level concentration of voters, measured by standard deviation in a candidate's vote shares (for the 1st-4th place candidate) across polling stations. Vote shares are from the first round. Includes only elections after 2004. *Estimation method:* Local linear regression with election-year fixed effects and a 125,000 inhabitant bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Figure A.22 Regression discontinuity coefficients on overall concentration of voters for specific candidates at different thresholds

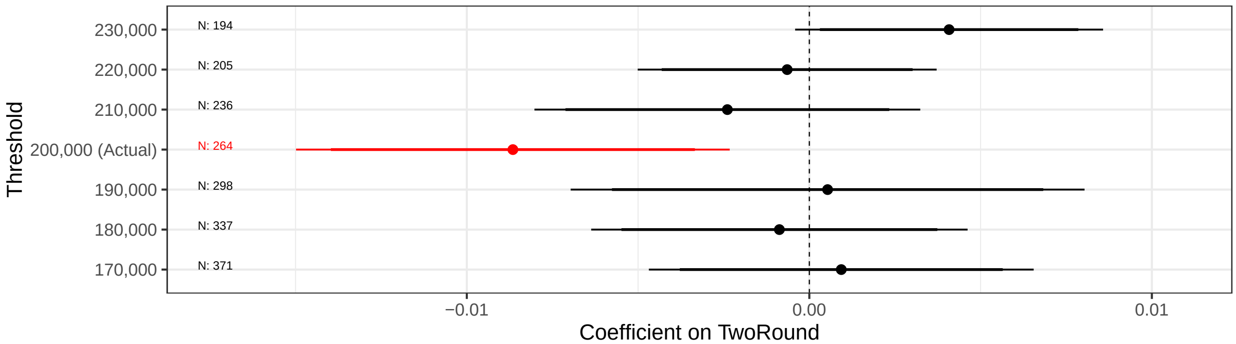

(a) Coefficient of variation

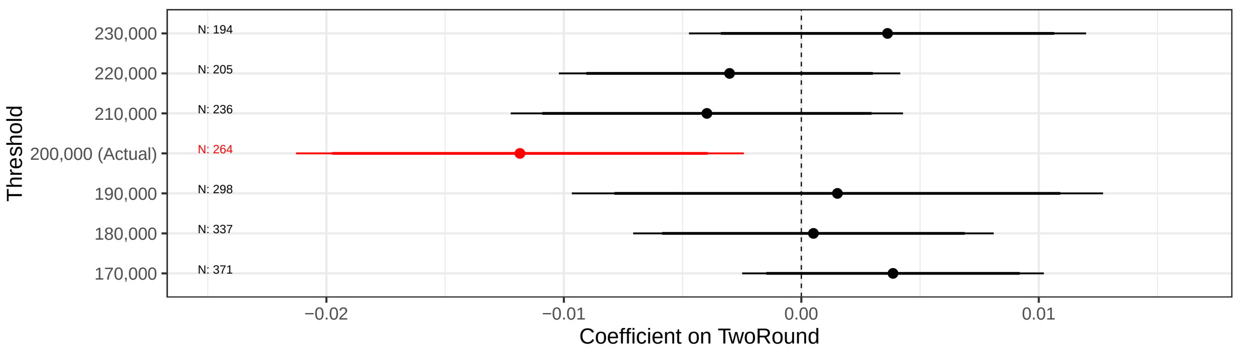

(b) Fractionalization

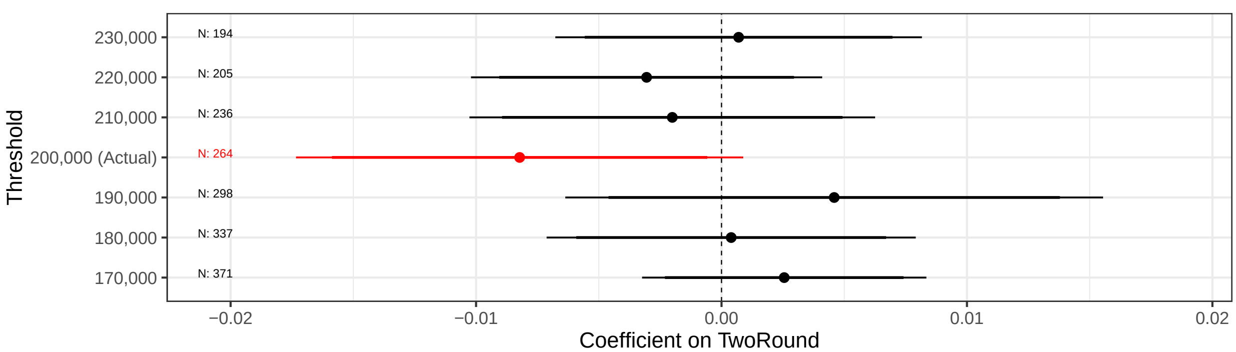

(c) Entropy

Overall concentration of voters for specific candidates, as measured by (a) *Coefficient of variation*, (b) *Fractionalization*, and (c) *Entropy*, using vote counts in polling stations. Vote shares are from the first round. The thicker horizontal lines represent the 90% confidence interval and the thinner horizontal lines represent the 95% confidence interval. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Figure A.23 Regression discontinuity coefficients on the candidate-level concentration in voters at different thresholds

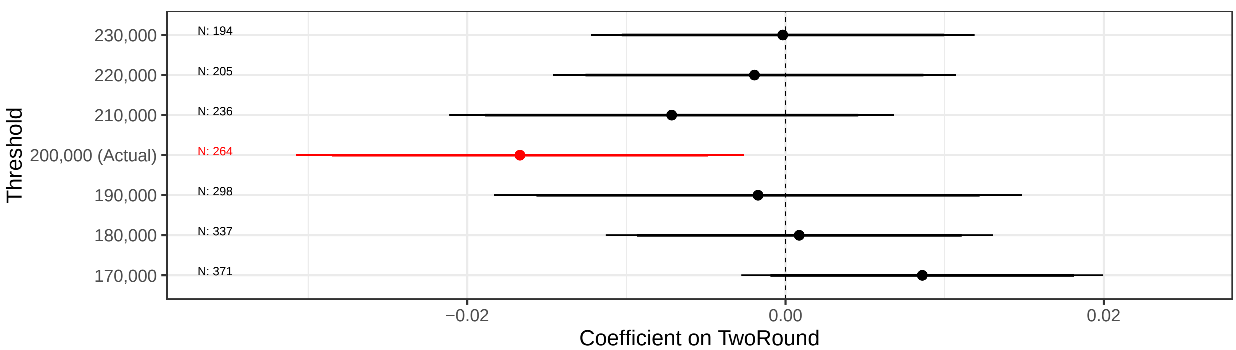

(a) Standard deviation in votes for 1st place candidate

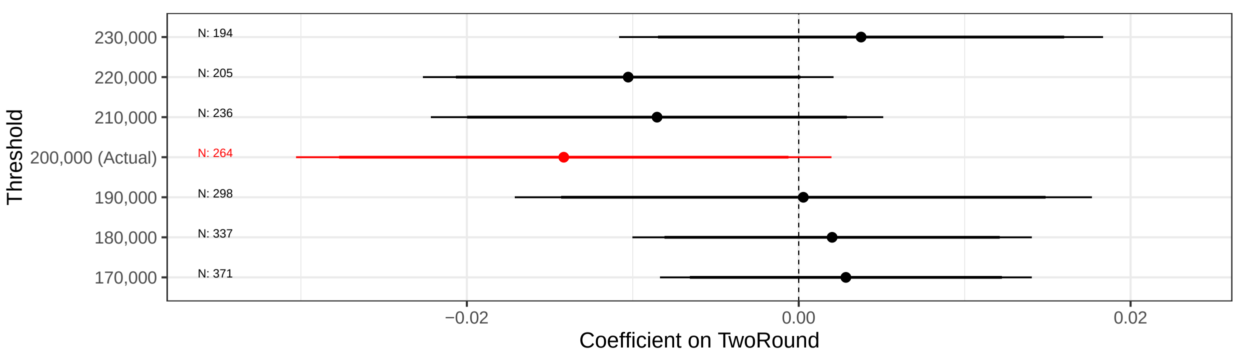

(b) Standard deviation in votes for 2nd place candidate

Standard deviation in a candidate's vote counts across polling stations, for the (a) 1st place and (b) 2nd place candidate. Vote shares are from the first round. The thicker horizontal lines represent the 90% confidence interval and the thinner horizontal lines represent the 95% confidence interval. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Figure A.24 Regression discontinuity coefficients on overall level of resources in municipal schools at different thresholds

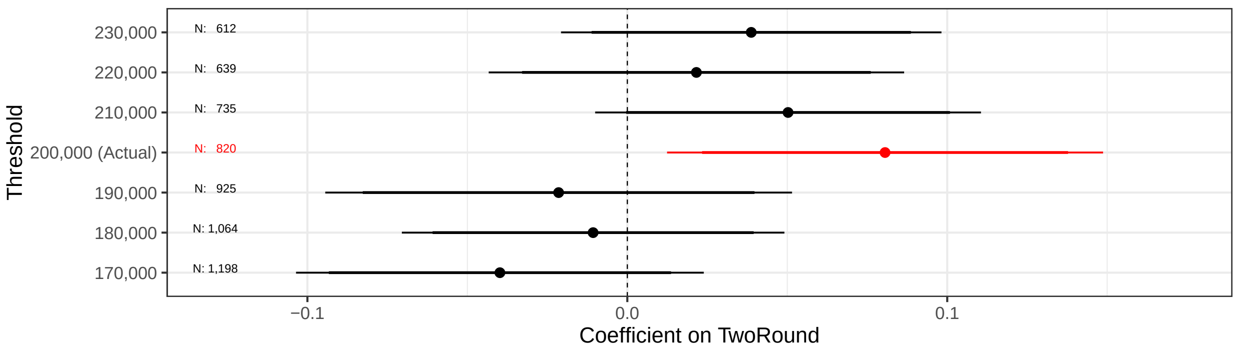

(a) Equipment, mean level of resources

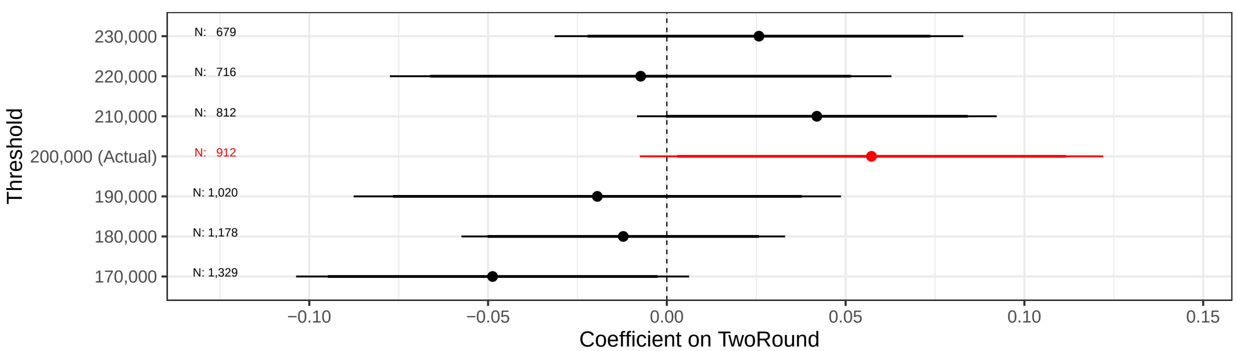

(b) Infrastructure, mean level of resources

Equipment and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. *Mean level of resources* is the mean index level across schools in the municipality for (a) equipment and (b) infrastructure. The thicker horizontal lines represent the 90% confidence interval and the thinner horizontal lines represent the 95% confidence interval. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Figure A.25 Regression discontinuity coefficients on distribution of resources in municipal schools at different thresholds

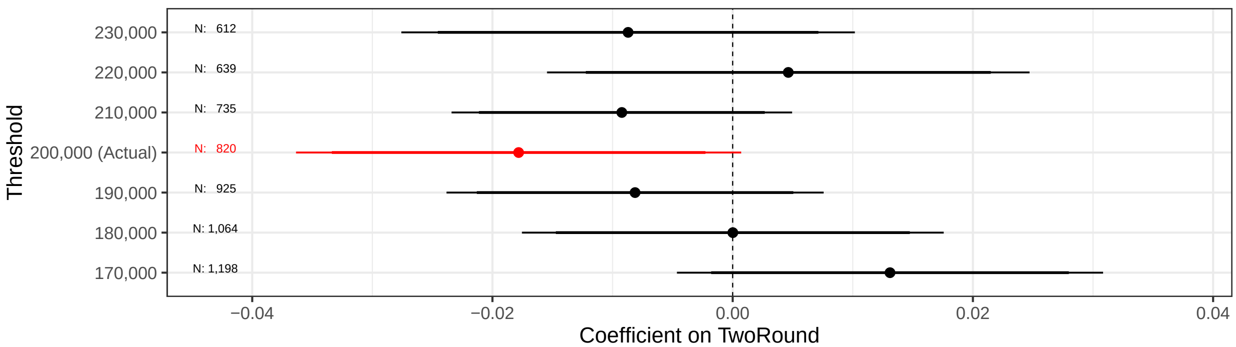

(a) Equipment, standard deviation in resources

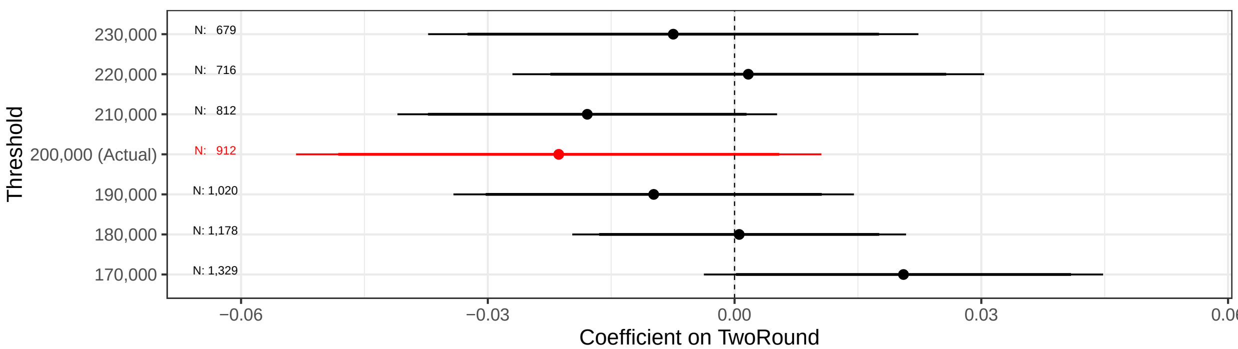

(b) Infrastructure, standard deviation in resources

Equipment and *Infrastructure* are indices constructed by taking the first principal component of a school's equipment and infrastructure elements, then calculating the school's percentile in the national distribution. *Standard deviation in resources* is the standard deviation in the index across schools in the municipality for (a) equipment and (b) infrastructure. The thicker horizontal lines represent the 90% confidence interval and the thinner horizontal lines represent the 95% confidence interval. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Figure A.26 Regression discontinuity coefficients on municipal education outcomes at different thresholds

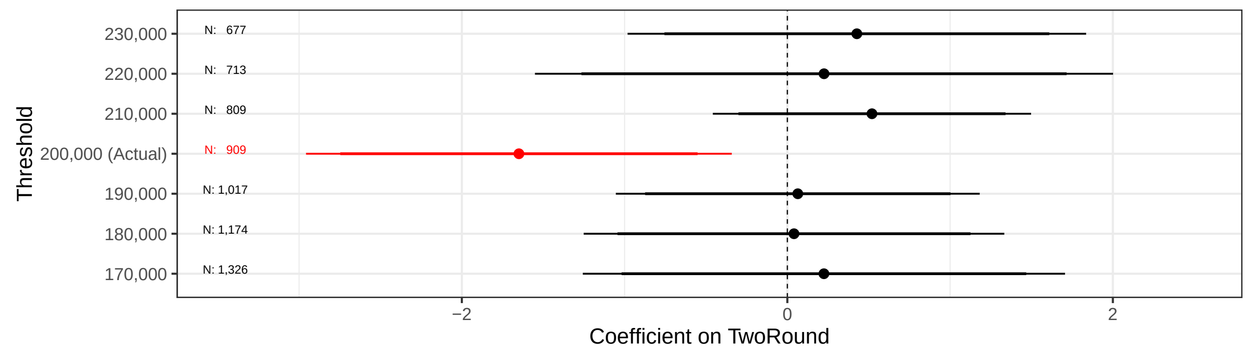

(a) Drop-out rate

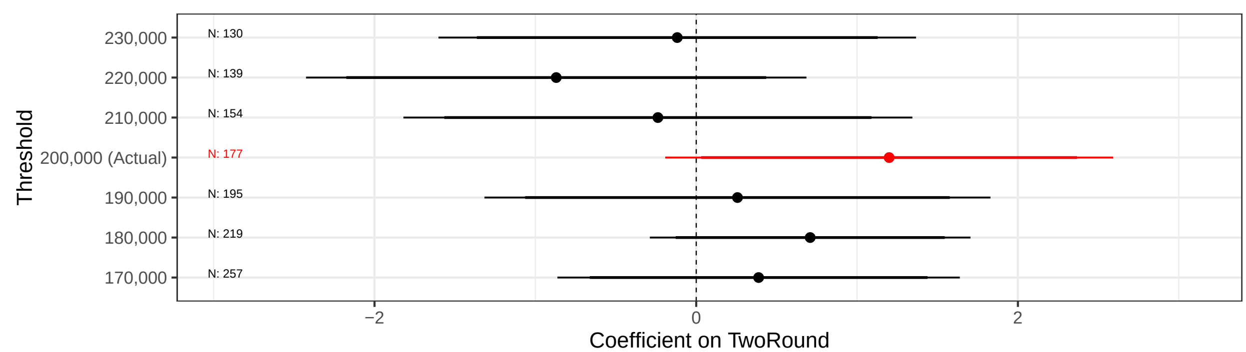

(b) Elementary literacy rate

Drop-out rate is from the School Census. It is the mean rate across schools in the municipality. *Elementary literacy rate* is from the 2000 and 2010 Demographic Census. It is the literacy rate of cohorts who are of elementary school age during the mayoral term. The thicker horizontal lines represent the 90% confidence interval and the thinner horizontal lines represent the 95% confidence interval. *Estimation method:* Local linear regression with election-year fixed effects and a 50,000 voter bandwidth. Population density included as a control separately across the cutoff. Standard errors clustered at the municipality level.

Appendix B

Appendix to Chapter 2

Table B.1 Data sources for outcomes of female textile, footwear, apparel workers

Data Source	**Year**	**Description**	**Industry code**	**Sub-district identifier**	**Variables**
1980 Population Census	1980	World Bank	2-digit	Yes	Female labor force participation; Educational attainment; Days and hours worked; Age at first marriage; Contraception use; Children ever born; Educational attainment of children
1990 Population Census	1990	World Bank	2-digit	No	Female labor force participation; Educational attainment; Days and hours worked; Age at first marriage; Children ever born; Educational attainment of children
1995 Intercensal Survey	1995	World Bank	2-digit	No	Female labor force participation; Educational attainment; Days and hours worked; Age at first marriage; Children ever born; Educational attainment of children
SUSENAS (National Socioeconomic Survey)	1992-1997	RAND	1-digit	Yes	Female labor force participation; Educational attainment; Days and hours worked; Age at first marriage; Contraception use; Children ever born; Educational attainment of children
	1998-2000	RAND	2-digit	Yes	Female labor force participation; Educational attainment; Days and hours worked; Age at first marriage; Contraception use; Children ever born; Educational attainment of children

Industry code: Industry codes at the 1-digit level only identify manufacturing versus non-manufacturing industries. Industry codes at the 2-digit level identify textile, footwear, apparel industries within the manufacturing industry.

Table B.2 Estimates on production worker wages, with province-by-year fixed effects: all manufacturing

	Treatment variable: Firms producing textile, footwear, apparel			
	Full sample	Exporters	Nike, Adidas, Reebok subdistricts	Exporters in Nike, Adidas, Reebok subdistricts
	(1)	(2)	(3)	(4)
TREAT × POST92	0.033***	0.110**	0.085***	0.155**
	(0.012)	(0.055)	(0.028)	(0.076)
Firms	40,103	2,677	3,029	486
Observations	169,575	13,537	23,259	3,065
Dep. var. mean	8.627	8.923	8.995	9.079

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variable is the log real base wage paid to production workers. Treatment is firms producing textile, footwear, apparel for all years between 1988-1991, relative to the rest of manufacturing. *Full sample* is the full sample of firms. *Exporters* is the sample of firms reporting exporting a positive percentage of output for all years between 1992-1995. *Nike, Adidas, Reebok subdistricts* is the sample of firms located in subdistricts where Nike, Adidas, Reebok factories operate. *Estimation method:* Regression includes province-by-year fixed effects, firm fixed effects, and district-specific time trends. Standard errors clustered at the firm level.

Table B.3 Estimates on production worker wages, with province-by-year or island-by-year fixed effects: textile, footwear, apparel

	Treatment variable: Firms in Nike, Adidas, Reebok subdistricts		*Treatment variable:* Firms in Nike, Adidas, Reebok districts	
	All TFA (1)	TFA exporters (2)	All TFA (3)	TFA exporters (4)
TREAT × POST92	0.120***	0.214**		
	(0.035)	(0.104)		
TREAT × POST92			0.053	0.053
			(0.035)	(0.094)
Firms	4,966	345	4,966	345
Observations	32,555	2,802	37,898	3,092
Dep. var. mean	8.598	8.824	8.598	8.824

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variable is the log real base wage paid to production workers. For Columns 1-2, treatment is textile, footwear, apparel firms operating in Nike, Adidas, Reebok subdistricts, relative to textile, footwear, apparel firms in other subdistricts. For Columns 3-4, treatment is textile, footwear, apparel firms operating in Nike, Adidas, Reebok districts, relative to textile, footwear, apparel firms in other districts. *All TFA* is the full sample of firms producing textile, footwear, apparel for all years between 1988-1991. *TFA exporters* is the sample of textile, footwear, apparel firms reporting exporting a positive percentage of output for all years between 1992-1995. *Estimation method:* Regression includes province-by-year fixed effects (in Columns 1-2), island-by-year fixed effects (in Columns 3-4), subdistrict fixed effects (in Columns 1-2), district fixed effects (in Columns 3-4), district-specific time trends (in Columns 1-2), and province-specific time trends (in Columns 3-4). Standard errors clustered at the subdistrict level (Columns 1-2) or district level (Columns 3-4).

Table B.4 Estimates on labor standards among female textile, footwear, apparel workers, with province-by-year or island-by-year fixed effects

	Treatment variable: Employed in textile, footwear, apparel		*Treatment variable:* Location with a Nike, Adidas, Reebok factory	
Geographic level of analysis	Subdistrict (1)	District (2)	Subdistrict (3)	District (4)
Panel A *Dependent variable: Younger than 16 years of age*				
TREAT × POST92	0.022^{**}	-0.005		
	(0.011)	(0.004)		
TREAT × POST92			-0.009	-0.015^{*}
			(0.026)	(0.008)
Observations	25,689	169,129	11,564	67,397
Dep. var. mean	0.075	0.073	0.059	0.069
Panel B *Dependent variable: Hours worked in past week*				
TREAT × POST92	2.914^{***}	0.202		
	(0.837)	(0.500)		
TREAT × POST92			-0.156	-1.279^{*}
			(1.579)	(0.727)
Observations	25,113	165,019	11,293	65,768
Dep. var. mean	42.927	41.352	43.995	43.538
Pre-period	1980	1980, 1990	1980	1980, 1990
Post-period	1998-2000	1995, 1998-2000	1998-2000	1995, 1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variables are an indicator equal to 1 if the individual is younger than 16 years of age (Panel A) and number of hours worked in the past week (Panel B). For Columns 1-2, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Columns 3-4, treatment is a female textile, footwear, apparel worker residing in a subdistrict or district where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in subdistricts or districts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes province-by-year fixed effects (Columns 1 and 3), island-by-year fixed effects (Columns 2 and 4), subdistrict-industry fixed effects (Column 1), district-industry fixed effects (Column 2), subdistrict fixed effects (Column 3), district fixed effects (Column 4), district-specific time trends (Columns 1 and 3), and province-specific time trends (Columns 2 and 4). Standard errors clustered at the subdistrict-industry (Column 1), district-industry (Column 2), subdistrict (Column 3), or district (Column 4) level.

Table B.5 Estimates on female labor force participation, district-level analysis

	Treatment variable: District with a textile, footwear, apparel firm		*Treatment variable:* District with a Nike, Adidas, Reebok factory	
Age group	15-64 (1)	15-25 (2)	15-64 (3)	15-25 (4)
TREAT × POST92	-0.021^{***} (0.006)	-0.035^{***} (0.008)		
TREAT × POST92			0.002 (0.009)	0.003 (0.013)
Observations	3,587,450	1,399,550	2,596,430	1,002,411
Dep. var. mean	0.400	0.344	0.415	0.369
Pre-period	1980, 1990	1980, 1990	1980, 1990	1980, 1990
Post-period	1992-2000	1992-2000	1992-2000	1992-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variable is an indicator equal to 1 if the individual is in the labor force. For Columns 1-2, treatment is residing in a district where a textile, footwear, apparel firm operates, relative to residing in urban areas in other districts. For Columns 3-4, treatment is residing in a district where a Nike, Adidas, Reebok factory operates, relative to residing in districts where a textile, footwear, apparel firm operates. An individual is considered in the labor force if the main activity in the last week was working or looking for work. *Estimation method:* Regression includes year fixed effects, district fixed effects, and province-specific time trends. Standard errors clustered at the district level.

Figure B.6 Estimates by year of female labor force participation among females 15-25 years

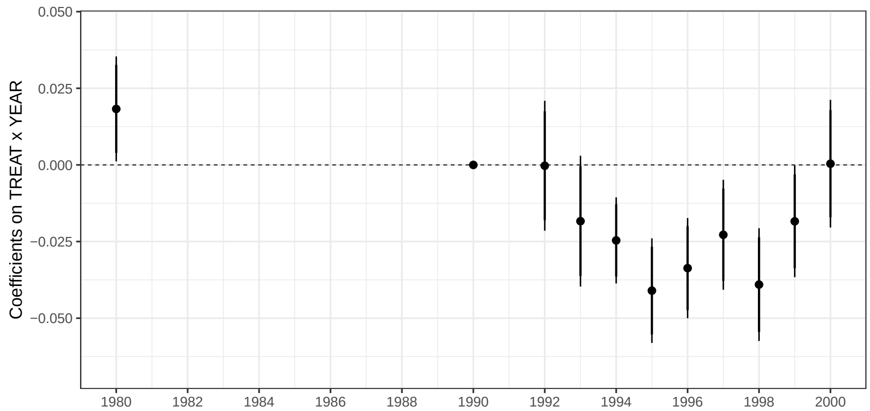

(a) District with a textile, footwear, apparel firm

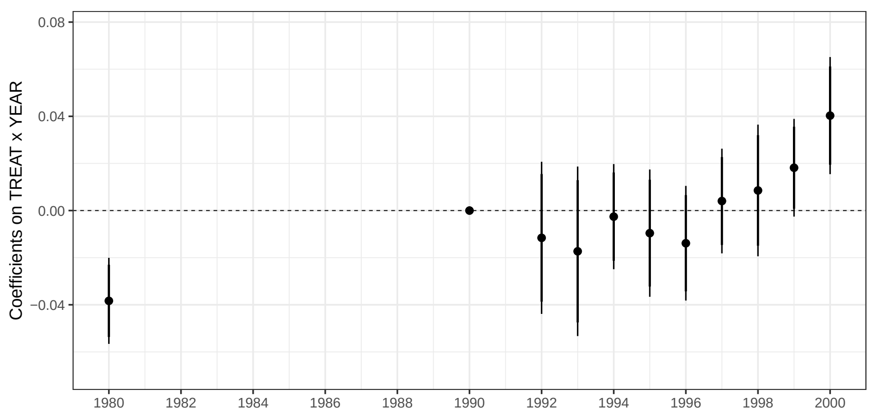

(b) District with a Nike, Adidas, Reebok factory

Dependent variable is an indicator equal to 1 if the individual is in the labor force. In Panel (a), treatment is residing in a district with a textile, footwear, apparel firm, relative to urban areas in other districts. In Panel (b), treatment is residing in a district with a Nike, Adidas, Reebok factory, relative to districts with a textile, footwear, apparel firm. The thicker vertical lines represent the 90% confidence interval and the thinner vertical lines represent the 95% confidence interval. *Estimation method:* Regression includes year fixed effects, district fixed effects, and province-specific time trends. Standard errors clustered at the district level.

Table B.7 Estimates on educational composition of female textile, footwear, apparel workers, district-level analysis

	Treatment variable: Employed in textile, footwear, apparel		*Treatment variable:* District with a Nike, Adidas, Reebok factory	
Age group	All (1)	$\leq$ 25 years (2)	All (3)	$\leq$ 25 years (4)
<u>*Panel A*</u> *Dependent variable: No schooling*				
TREAT × POST92	0.026	0.058^{***}		
	(0.017)	(0.017)		
TREAT × POST92			0.062^{**}	0.031
			(0.030)	(0.034)
Observations	169,135	91,274	67,399	41,315
Dep. var. mean	0.437	0.302	0.324	0.230
<u>*Panel B*</u> *Dependent variable: Primary schooling*				
TREAT × POST92	-0.053^{**}	-0.065^{**}		
	(0.022)	(0.027)		
TREAT × POST92			−0.060	−0.022
			(0.037)	(0.049)
Observations	169,135	91,274	67,399	41,315
Dep. var. mean	0.438	0.554	0.539	0.622
<u>*Panel C*</u> *Dependent variable: Secondary schooling*				
TREAT × POST92	0.034^{***}	0.013		
	(0.011)	(0.015)		
TREAT × POST92			−0.00002	−0.010
			(0.025)	(0.029)
Observations	169,135	91,274	67,399	41,315
Dep. var. mean	0.110	0.135	0.127	0.143
Pre-period	1980, 1990	1980, 1990	1980, 1990	1980, 1990
Post-period	1995, 1998-2000	1995, 1998-2000	1995, 1998-2000	1995, 1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variables are indicators for the highest level of schooling attained: none (Panel A), primary (Panel B), and secondary (Panel C). For Columns 1-2, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Columns 3-4, treatment is a female textile, footwear, apparel worker residing in a district where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in districts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes year fixed effects, district-industry fixed effects (Columns 1-2), district fixed effects (Columns 3-4), and province-specific time trends. Standard errors clustered at the district-industry (Columns 1-2) or district (Columns 3-4) level.

Table B.8 Estimates on educational composition of female textile, footwear, apparel workers, with province-by-year fixed effects

	Treatment variable: Employed in textile, footwear, apparel		*Treatment variable:* Subdistrict with a Nike, Adidas, Reebok factory	
Age group	All (1)	$\leq$ 25 years (2)	All (3)	$\leq$ 25 years (4)
Panel A *Dependent variable: No schooling*				
TREAT × POST92	0.016	0.117***		
	(0.023)	(0.025)		
TREAT × POST92			−0.020	−0.025
			(0.099)	(0.098)
Observations	25,690	15,052	11,565	6,949
Dep. var. mean	0.460	0.410	0.390	0.337
Panel B *Dependent variable: Primary schooling*				
TREAT × POST92	−0.036	−0.133***		
	(0.028)	(0.038)		
TREAT × POST92			−0.008	0.008
			(0.125)	(0.111)
Observations	25,690	15,052	11,565	6,949
Dep. var. mean	0.391	0.449	0.475	0.535
Panel C *Dependent variable: Secondary schooling*				
TREAT × POST92	0.055***	0.035		
	(0.017)	(0.028)		
TREAT × POST92			0.038	0.018
			(0.040)	(0.040)
Observations	25,690	15,052	11,565	6,949
Dep. var. mean	0.131	0.132	0.125	0.123
Pre-period	1980	1980	1980	1980
Post-period	1998-2000	1998-2000	1998-2000	1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variables are indicators for the highest level of schooling attained: none (Panel A), primary (Panel B), and secondary (Panel C). For Columns 1-2, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Columns 3-4, treatment is a female textile, footwear, apparel worker residing in a subdistrict where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in subdistricts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes province-by-year fixed effects, subdistrict-industry fixed effects (Columns 1-2), subdistrict fixed effects (Columns 3-4), and district-specific time trends. Standard errors clustered at the subdistrict-industry (Columns 1-2) or subdistrict (Columns 3-4) level.

Table B.9 Estimates on downstream marriage outcomes for female textile, footwear, apparel workers, district-level analysis

	Dependent variable:					
	Ever married		Age at marriage		Contraception use	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A *Treatment variable: Employed in textile, footwear, apparel*						
TREAT × POST92	−0.006		0.259***		0.002	
	(0.017)		(0.093)		(0.018)	
Panel B *Treatment variable: District with a Nike, Adidas, Reebok factory*						
TREAT × POST92		0.083**		−0.149		0.044
		(0.038)		(0.164)		(0.037)
Observations	169,135	67,399	96,319	33,116	33,293	11,599
Dep. var. mean	0.578	0.497	18.134	18.516	0.432	0.427
Pre-period	1980, 1990	1980, 1990	1980, 1990	1980, 1990	1980	1980
Post-period	1995, 1998-2000	1995, 1998-2000	1995, 1998-2000	1995, 1998-2000	1998-2000	1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Ever married is an indicator equal to 1 if the marital status is married, widowed, or divorced. *Age at marriage* is the individual's age at her first marriage. *Contraception use* is an indicator equal to 1 if the individual ever used a contraceptive method. For Panel A, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Panel B, treatment is a female textile, footwear, apparel worker residing in a district where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in districts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes year fixed effects, district-industry fixed effects (Panel A), district fixed effects (Panel B), and province-specific time trends. Standard errors clustered at the district-industry (Panel A) or district (Panel B) level.

Figure B.10 Estimates by year on downstream marriage outcomes for female textile, footwear, apparel workers, district-level analysis

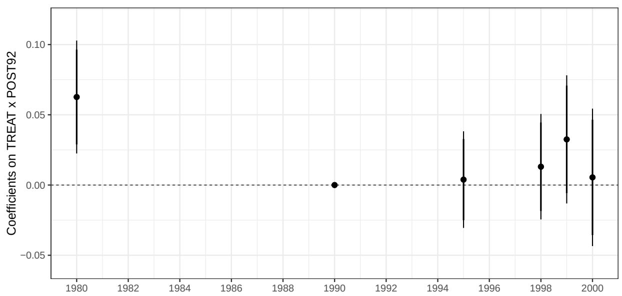

(a) Ever married
Employed in textile, footwear, apparel

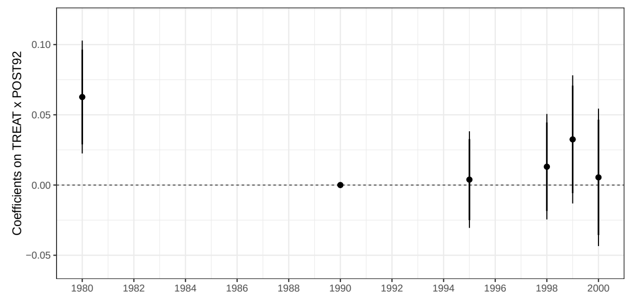

(b) Ever married
District with a Nike, Adidas, Reebok factory

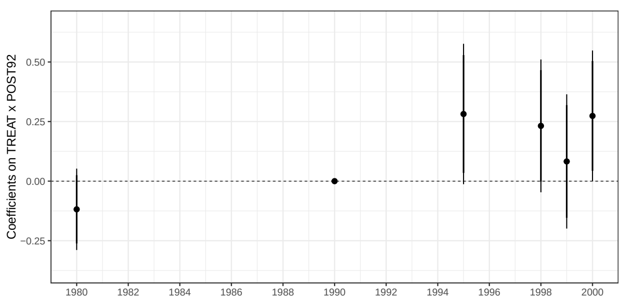

(c) Age at first marriage
Employed in textile, footwear, apparel

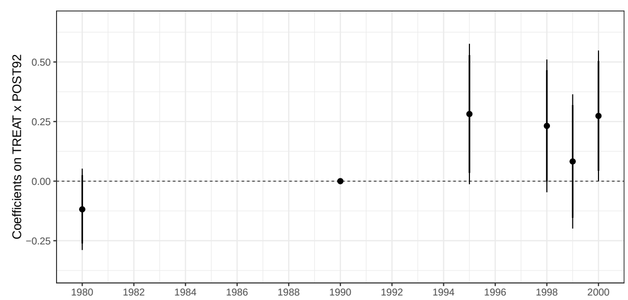

(d) Age at first marriage
District with a Nike, Adidas, Reebok factory

Ever married is an indicator equal to 1 if the marital status is married, widowed, or divorced. *Age at marriage* is the individual's age at her first marriage. For Panels (a) and (c), treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Panels (b) and (d), treatment is a female textile, footwear, apparel worker residing in a district where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in districts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes year fixed effects, district-industry fixed effects (Panels (a) and (c)), district fixed effects (Panels (b) and (d)), and province-specific time trends. Standard errors clustered at the district-industry (Panels (a) and (c)) or district (Panels (b) and (d)) level.

Table B.11 Estimates on downstream marriage outcomes for female textile, footwear, apparel workers, with province-by-year fixed effects

	Dependent variable:					
	Ever married		Age at marriage		Contraception use	
	(1)	(2)	(3)	(4)	(5)	(6)
<u>*Panel A*</u> *Treatment variable: Employed in textile, footwear, apparel*						
TREAT × POST92	-0.050^{**}		-0.033		-0.003	
	(0.024)		(0.229)		(0.031)	
<u>*Panel B*</u> *Treatment variable: Subdistrict with a Nike, Adidas, Reebok factory*						
TREAT × POST92		-0.030		0.104		0.009
		(0.047)		(0.491)		(0.057)
Observations	25,690	11,565	13,782	6,002	10,464	4,610
Dep. var. mean	0.548	0.529	18.430	18.548	0.426	0.441
Pre-period	1980	1980	1980	1980	1980	1980
Post-period	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Ever married is an indicator equal to 1 if the marital status is married, widowed, or divorced. *Age at marriage* is the individual's age at her first marriage. *Contraception use* is an indicator equal to 1 if the individual ever used a contraceptive method. For Panel A, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Panel B, treatment is a female textile, footwear, apparel worker residing in a subdistrict where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in subdistricts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes province-by-year fixed effects, subdistrict-industry fixed effects (Panel A), subdistrict fixed effects (Panel B), and district-specific time trends. Standard errors clustered at the subdistrict-industry (Panel A) or subdistrict (Panel B) level.

Table B.12 Estimates on downstream fertility outcomes for female textile, footwear, apparel workers, district-level analysis

	Dependent variable:							
	Ever had children		Age at first child		Number children born		Fraction children in school	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A *Treatment variable: Employed in textile, footwear, apparel*								
TREAT × POST92	−0.001		0.021		0.080		-0.036^{***}	
	(0.017)		(0.156)		(0.061)		(0.011)	
Panel B *Treatment variable: District with a Nike, Adidas, Reebok factory*								
TREAT × POST92		0.074^{**}		−0.005		0.335^{***}		−0.025
		(0.031)		(0.259)		(0.097)		(0.017)
Observations	169,135	67,399	71,928	24,131	169,135	67,399	48,435	16,228
Dep. var. mean	0.485	0.405	22.515	22.049	1.622	1.230	0.582	0.599
Pre-period	1980, 1990	1980, 1990	1980, 1990	1980, 1990	1980, 1990	1980, 1990	1980, 1990	1980, 1990
Post-period	1995, 1998-2000	1995, 1998-2000	1995, 1998-2000	1995, 1998-2000	1995, 1998-2000	1995, 1998-2000	1995, 1998-2000	1995, 1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Ever had children is an indicator equal to 1 if the individual ever had a child. *Age at first child* is the individual's age at which she had her first child. *Number children born* is the number of children ever born. *Fraction children in school* is the fraction of the individual's children who are between 5-18 years who are currently attending school. For Panel A, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Panel B, treatment is a female textile, footwear, apparel worker residing in a district where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in districts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes year fixed effects, district-industry fixed effects (Panel A), district fixed effects (Panel B), and province-specific time trends. Standard errors clustered at the district-industry (Panel A) or district (Panel B) level.

Figure B.13 Estimates by year on downstream fertility outcomes for female textile, footwear, apparel workers, district-level analysis

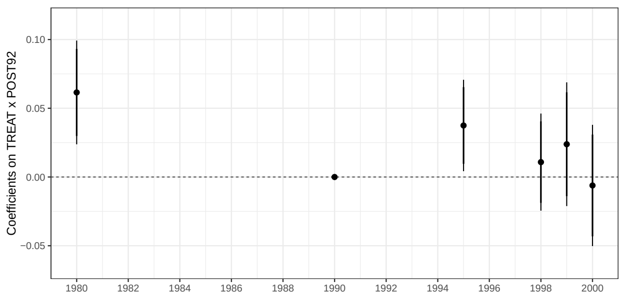

(a) Ever had children
Employed in textile, footwear, apparel

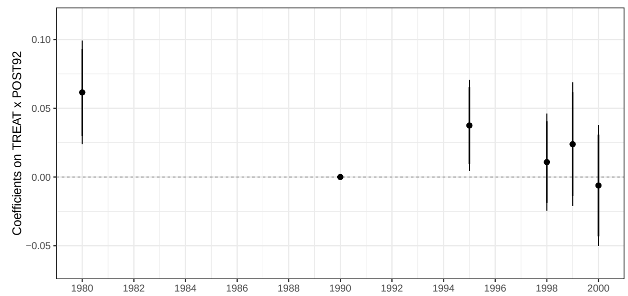

(b) Ever had children
District with a Nike, Adidas, Reebok factory

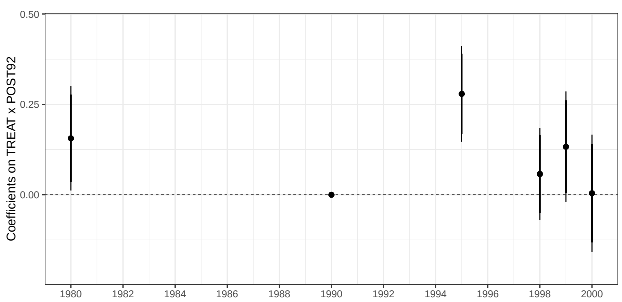

(c) Number children born
Employed in textile, footwear, apparel

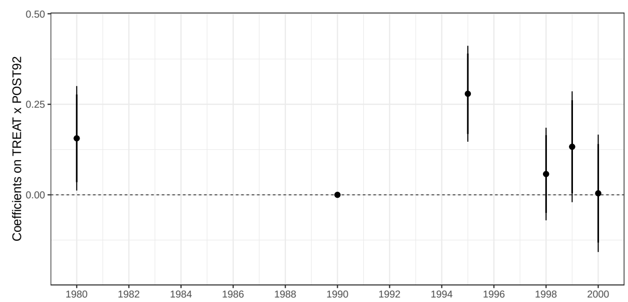

(d) Number children born
District with a Nike, Adidas, Reebok factory

Ever had children is an indicator equal to 1 if the individual ever had a child. *Number children born* is the number of children ever born. For Panels (a) and (c), treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Panels (b) and (d), treatment is a female textile, footwear, apparel worker residing in a district where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in districts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes year fixed effects, district-industry fixed effects (Panels (a) and (c)), district fixed effects (Panels (b) and (d)), and province-specific time trends. Standard errors clustered at the district-industry (Panels (a) and (c)) or district (Panels (b) and (d)) level.

Table B.14 Estimates on downstream fertility outcomes for female textile, footwear, apparel workers, with province-by-year fixed effects

	Dependent variable:							
	Ever had children		Age at first child		Number children born		Fraction children in school	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A *Treatment variable: Employed in textile, footwear, apparel*								
TREAT × POST92	-0.059^{**}		-0.592^{*}		-0.158^{*}		-0.038	
	(0.024)		(0.335)		(0.096)		(0.034)	
Panel B *Treatment variable: Subdistrict with a Nike, Adidas, Reebok factory*								
TREAT × POST92		-0.045		0.181		0.038		-0.027
		(0.054)		(0.755)		(0.214)		(0.076)
Observations	25,690	11,565	9,869	4,280	25,690	11,565	6,292	2,743
Dep. var. mean	0.442	0.422	22.280	22.215	1.466	1.327	0.584	0.593
Pre-period	1980	1980	1980	1980	1980	1980	1980	1980
Post-period	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000	1998-2000

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Ever had children is an indicator equal to 1 if the individual ever had a child. *Age at first child* is the individual's age at which she had her first child. *Number children born* is the number of children ever born. *Fraction children in school* is the fraction of the individual's children who are between 5-18 years who are currently attending school. For Panel A, treatment is employment in the textile, footwear, apparel sector, relative to employment in all other manufacturing sectors. For Panel B, treatment is a female textile, footwear, apparel worker residing in a subdistrict where a Nike, Adidas, Reebok factory operates, relative to female textile, footwear, apparel workers residing in subdistricts where a textile, footwear, apparel firm operates. *Estimation method:* Regression includes province-by-year fixed effects, subdistrict-industry fixed effects (Panel A), subdistrict fixed effects (Panel B), and district-specific time trends. Standard errors clustered at the subdistrict-industry (Panel A) or subdistrict (Panel B) level.

Appendix C

Appendix to Chapter 3

Table C.1 Description of variables in poverty score calculation and machine learning covariates

Variable	Poverty score subcomponent	Machine learning covariate list	
		Full	**Limited**
Neighborhood strata	–	X	X
Household variables			
Material of housing exterior	Housing	X	X
Material of floor	Housing	X	X
Material of roof	Housing	X	X
Main source of lighting in home	–	X	X
Sanitation system	Utilities	X	X
Main water source	Utilities	X	X
Trash disposal system	Utilities	X	X
Household owns: refrigerator	Housing	X	X
Household owns: television	Housing	X	
Household owns: fan	Housing	X	
Household owns: blender	Housing	X	
Household owns: washing machine	Housing	X	X
Household size	Demographic	X	X
Total rooms in house	Demographic	X	X
Total bedrooms in house	–	X	X
Individual variables			
Relationship to head of household	–	X	X
Age	Demographic	X	X
Sex	–	X	X
Marital status	–	X	X
School enrollment status	–	X	X
Highest educational attainment	Education	X	X
Main economic activity in past week	Education, Demographic	X	X
Normal salary or income from main job	Education, Demographic	X	
Number of employees at main job	Education	X	
Social security health coverage	Education	X	

Table C.2 Balance in household characteristics in Quality of Life training and holdout sets

	Training set (1)	Holdout set (2)	p-value of difference (3)
Number of households	40,938	4,039	
Below the poverty line	0.561	0.556	0.602
Characteristics of household head			
Female	0.339	0.351	0.173
Age	42.391	42.295	0.866
Years of schooling	8.582	8.629	0.732
Married or in union	0.623	0.608	0.093^{*}
Characteristics of household			
Household size	4.053	4.062	0.509
Number of children under 4 years	0.391	0.405	0.256
Total rooms	3.587	3.584	0.767
Owns refrigerator	0.513	0.506	0.073^{*}
Wall material: mud, adobe, brick	0.976	0.973	0.457
Floor material: cement, tile, rug, marble	0.928	0.921	0.088^{*}
Water connected	0.970	0.971	0.675
Toilet connected	0.921	0.925	0.252
Electric lighting	0.996	0.995	0.529

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Summary characteristics of households in Quality of Life training set (Column 1) and holdout set (Column 2). p-value on the difference in means between the training and holdout set (Column 3) controls for dummies for survey year and neighborhood strata. F-stat for all differences jointly significant: 1.00 ($p = 0.146$). *Below the poverty line* is defined as having a consumption level below the national poverty line. *Water connected* is an indicator equal to 1 if a household's main water source is piped water. *Toilet connected* is an indicator equal to 1 if a household's toilet is connected to the sewage system. *Electric lighting* is an indicator equal to 1 if a household's main source of lighting is electricity.

Figure C.3 Estimated size of discontinuity in density of assigned poverty scores at placebo eligibility thresholds, for households surveyed in 1998-2003

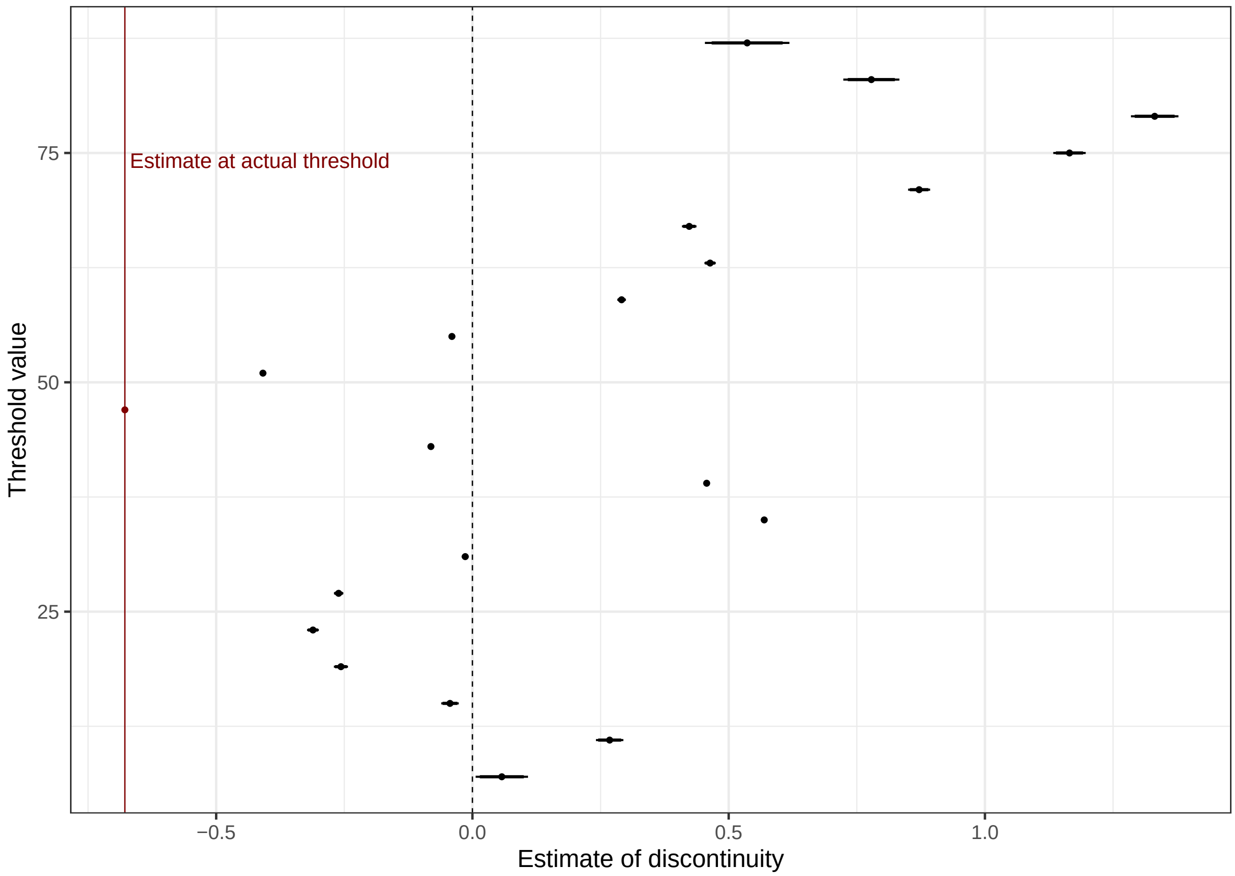

Estimates of the size of the discontinuity in the density of assigned poverty scores at placebo eligibility thresholds, using a McCrary test. The discontinuity estimate is the estimated log difference in heights of the density function. The thicker horizontal lines represent the 90% confidence interval and the thinner horizontal lines represent the 95% confidence interval. The poverty score algorithm was released to local officials in late 1997. Data from SISBÉN survey.

Figure C.4 Density of reconstructed poverty scores

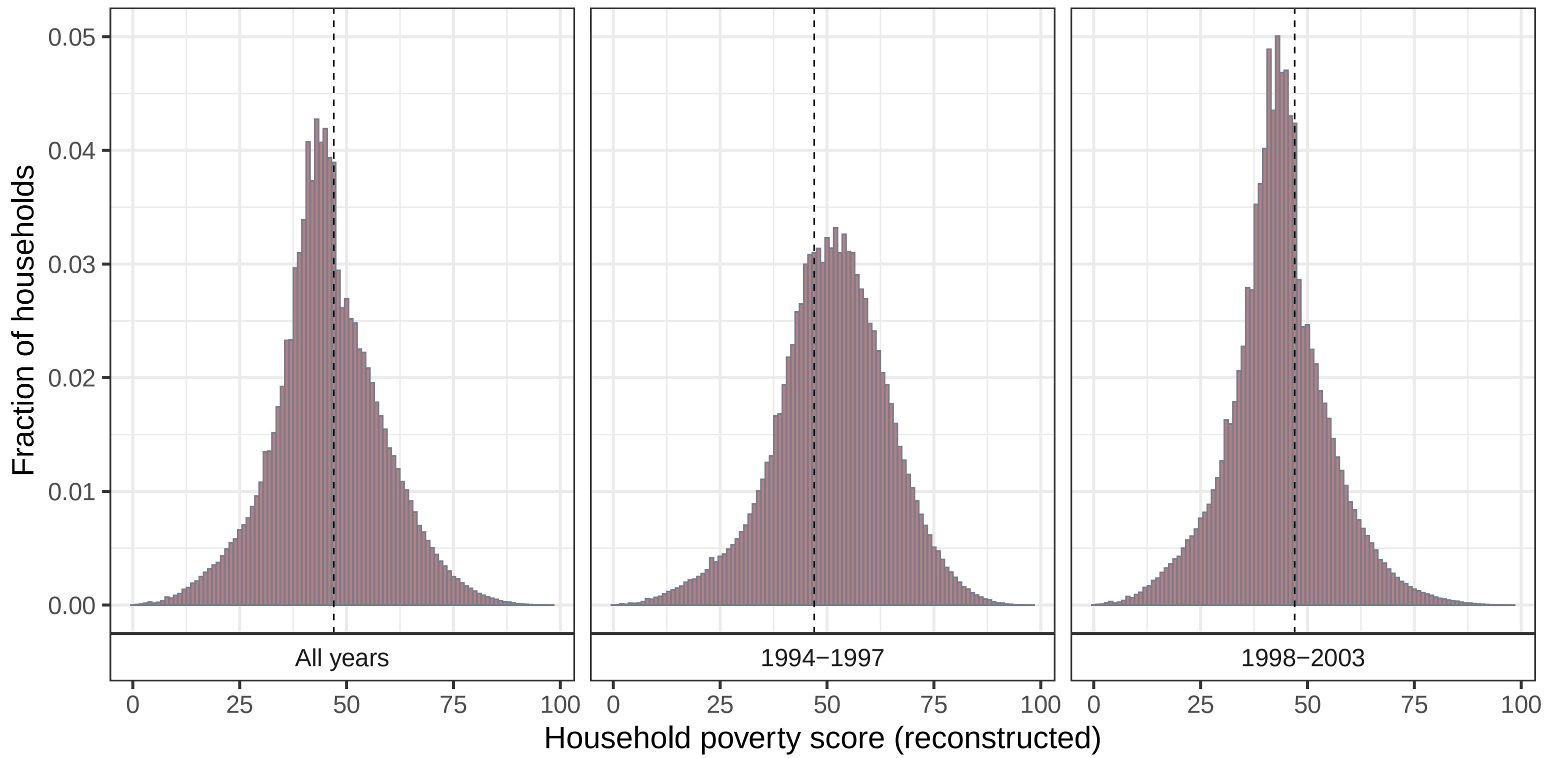

Figures plot the density of reconstructed poverty scores in SISBÉN. Distribution plotted separately for households surveyed across all years (left panel), 1994-1997 (middle panel), and 1998-2003 (right panel). The poverty score is reconstructed from the household's survey answers or the sum of the poverty score subcomponents. The poverty score algorithm was released to local officials in late 1997. Households with a poverty score below 47 (vertical dotted line) are eligible for a variety of social programs. Data from SISBÉN survey.

Table C.5 Estimated size of discontinuity in density of reconstructed poverty scores at eligibility threshold

Years	Estimate	Standard error	p-value	Bandwidth
Panel A: Households surveyed in 1994-1997				
1994	0.066^{***}	[0.008]	< 0.001	26.961
1995	0.137^{***}	[0.004]	< 0.001	24.015
1996	0.050^{***}	[0.006]	< 0.001	25.986
1997	0.121^{***}	[0.005]	< 0.001	24.593
1994-1997	0.115^{***}	[0.003]	< 0.001	25.029
Panel B: Households surveyed in 1998-2003				
1998	-0.285^{***}	[0.005]	< 0.001	27.909
1999	-0.510^{***}	[0.004]	< 0.001	24.449
2000	-0.515^{***}	[0.003]	< 0.001	26.314
2001	-0.731^{***}	[0.004]	< 0.001	24.128
2002	-0.936^{***}	[0.005]	< 0.001	19.346
2003	-0.908^{***}	[0.009]	< 0.001	18.878
1998-2003	-0.668^{***}	[0.002]	< 0.001	24.396
Panel C: Households across all years				
All	-0.450^{***}	[0.001]	< 0.001	29.037

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Estimates of the size of the discontinuity in the density of reconstructed poverty scores at the eligibility threshold, using a McCrary test. The discontinuity estimate is the estimated log difference in heights of the density function. The poverty score is reconstructed from the household's survey answers or the sum of the poverty score subcomponents. The poverty score algorithm was released to local officials in late 1997. Households with a poverty score below 47 are eligible for a variety of social programs. Data from SISBÉN survey.

Table C.6 Correlation between manipulating survey answers and altering the poverty score

	Dependent variable			
	Altering poverty score		Altering poverty score to be eligible	
	(1)	(2)	(3)	(4)
Full ML - Limited ML	-0.072^{***}	0.020^{***}	-0.017^{***}	-0.022^{***}
	(0.002)	(0.003)	(0.001)	(0.001)
Full ML - Limited ML × Post-1998		-0.070^{***}		0.008^{***}
		(0.003)		(0.001)
N	2,098,162	2,098,162	2,098,162	2,098,162

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Dependent variable is either an indicator equal to 1 if the household's reconstructed poverty score does not match the assigned poverty score (Columns 1-2) or an indicator equal to 1 if the household's altered score changes their eligibility status (Columns 3-4). *Full ML - Limited ML* is the difference in a household's poverty prediction when using predictions trained on the full covariate list and when using predictions trained on the limited covariate list. *Post-1998* is an indicator equal to 1 if the household was surveyed between 1998-2003 and 0 if the household was surveyed between 1994-1997. The poverty score is reconstructed from the household's survey answers or the sum of the poverty score subcomponents. The poverty score algorithm was released to local officials in late 1997. Sample restricted to households with a reconstructed poverty score above the eligibility threshold (47). Data from SISBÉN.