Long-Run Impacts of Intergovernmental Transfers

Irineu de Carvalho Filho
Stephan Litschig

ABSTRACT

We provide regression discontinuity evidence on long-run impacts of a temporary increase in federal transfers to local governments in Brazil. Revenues and expenditures in treatment communities increased by about 20 percent during a four-year period in the early 1980s. Previously established schooling and literacy gains of school-age cohorts as well as reduced poverty in the community overall as of 1991 are generally attenuated but persist in 2000. Children and adolescents born after the funding boost show gains of about 0.06–0.10 standard deviation across the entire score distribution of two nationwide exams at the end of the 2000s.

*Irineu de Carvalho Filho is a Principal Economist at the Monetary Authority of Singapore. Stephan Litschig is Professor at the National Graduate Institute for Policy Studies (GRIPS) in Tokyo (s-litschig@grips.ac.jp).**An earlier version of this paper had the title "The Long-Run and Next-Generation Education Impact of Intergovernmental Transfers." Methodology, some data and discussion, and results in this paper's sections II.A, II.B, III, IV.A, IV.B, IV.D, V, VI.A, VI.B, VII, and IX.A, Tables 1, 2, 3 and 4, and Figures 1, 2, 3 and 4 are partially or entirely identical to corresponding sections, tables, and figures in Litschig and Morrison (2013) (© American Economic Association; reproduced with permission of the* American Economic Journal: Applied Economics*). The views expressed in this paper are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management. The authors gratefully acknowledge comments and suggestions from Lori Beaman, Antonio Ciccone, Emma Duchini, Ricardo Estrada, Gabrielle Fack, Patricia Funk, Gianmarco León, Andreas Madestam, Leonardo Monasterio, Kevin Morrison, Hannes Müller, Hessel Oosterbeek, Per Pettersson-Lidbom, Giacomo Ponzetto, Erik Plug, Alessandro Tarozzi, Björn Tyrefors Hinnerich, Daniel Wilson, and seminar participants at Hitotsubashi University, GRIPS Tokyo, FGV São Paulo, RES Manchester, Universitat Autònoma de Barcelona, IADB, Tinbergen Institute Amsterdam, Barcelona GSE Winter Workshop, Bristol University, Georgetown University, Stockholm University, the Institut d'Economia de Barcelona V Workshop on Fiscal Federalism, the Workshop on Empirical Research in Economics of Education at Universitat Rovira i Virgili, NEUDC Harvard and at Universitat Pompeu Fabra. The authors thank Reynaldo Fernandes for his help with the school census data. Litschig acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2011-0075). The authors have nothing to disclose. The data used in this article are publicly available. Please refer to the data section and the Online Appendix for details.*
[Submitted September 2017; accepted March 2020]; doi:10.3368/jhr.57.3.0917-9064R2
JEL Classification: H40, H72, I21, and O15
ISSN 0022-166X E-ISSN 1548-8004
Supplementary materials are freely available online at: http://uwpress.wisc.edu/journals/journals/jhr-supplementary.html

Stephan Litschig https://orcid.org/0000-0001-9054-6207

I. Introduction

What are the long-run impacts of federal transfers to local governments on schooling, learning, and earnings of the local population? Intergovernmental grants finance about 60 percent of decentralized public service provision in developing countries and about one-third in OECD countries (Boadway and Shah 2009). Yet, despite their importance, evidence on the short-run development impacts of such grants is scant, while evidence on their long-run effects is absent altogether, mainly due to the high data and research design requirements associated with long-run causal studies.

We investigate the impact of a temporary increase in block grants to local governments in Brazil in the early 1980s on development outcomes over the subsequent three decades. We use the same regression discontinuity design as Litschig and Morrison (2013), exploiting that a substantial part of national tax revenue was redistributed to local governments only on the basis of population, via a formula based on cutoffs. For relatively small communities, the extra funding at the cutoffs translated into public spending increases on education, transportation, and housing and urban infrastructure of about 20 percent during the four years from 1982 to the end of 1985. Importantly, the funding discontinuities between the treatment and comparison communities disappeared in 1986 and never reappeared.

The key contribution of this study is to document that the early 1980s funding boost led to gains in cognitive skills of the next generation, defined here as those born after the extra funding had expired. We examine data from two nationwide standardized tests that were administered in the late 2000s, more than 20 years after the extra funding had stopped. The Prova Brasil is a compulsory exam that measures mathematics and Portuguese language proficiency of public school students in fifth and ninth grade, approximately at ages 10 and 14, respectively. The ENEM (Exame Nacional do Ensino Médio) measures general proficiency for students in the process of completing or having completed 12th grade, and participation is voluntary. For both tests, we pool several rounds that were administered between 2007 and 2011. We find that ninth- and 12th-graders—who attended school during the mid- to late-1990s and the decade of the 2000s—show gains of about 0.06 to 0.10 standard deviation across the entire test score distribution at the end of the 2000s. In contrast, fifth-graders show no evidence of test score gains at the end of the 2000s, which is likely due to sample selection bias, as further discussed below.

Our second contribution is to provide a follow-up on the Litschig and Morrison (2013) study. Their paper shows that the temporary public spending increase generated significant improvements in completed grades and literacy rates of school-age cohorts, as well as reduced poverty in the community overall, as measured in the census of 1991. We examine whether these education gains of school-age cohorts in the early 1980s were durable or instead faded with time because completed grades and literacy in comparison communities eventually caught up through, for example, adult literacy and remedial education. We start by documenting that there are no differential migration patterns at the cutoffs in any of the census years, suggesting that the potential for sample selection bias is limited. We then show that the schooling and literacy gains of school-age cohorts during the boost period as well as reduced poverty in the community overall as of 1991 are generally attenuated but persist in 2000. By 2010 the estimated education and income gains are still positive but often statistically indistinguishable from zero.

A first potential mechanism for the gains in cognitive skills of older children and adolescents in the late 2000s are ratchet effects in public service provision. Litschig and Morrison (2013) find evidence that the student-to-teacher ratio in the local public primary school system (Grades 1–4) was reduced from about 21 by about two to three students per teacher by 1991, six years after the extra transfers had stopped. In this study we document that the class size reduction in public primary schools persisted throughout the 1990s and 2000s. The student-to-teacher ratio reduction might have persisted because dismissal of civil servants is difficult under Brazilian labor regulations, and voluntary resignation or early retirement infrequent. As a result, test score gains of ninth- and 12th-graders in the late 2000s might reflect long-run effects of reduced class size when these children were going through primary school.

How was the persistent class size reduction financed? A first possibility is that the initial funding boost triggered improvements in municipal budgets through increased own revenue collection. While point estimates are mostly positive, they are not statistically different from zero. For the total budget we lack power to rule out a revenue increase of up to about 10 percent. Even with a constant budget, it is possible that a more educated population prioritized class size reductions over other (education) spending categories, including waste. Unfortunately, however, we lack the disaggregated data on expenditure line items to further investigate how the class size reductions were financed. We also acknowledge that the test score gains we document might therefore have come at the expense of other unmeasured dimensions of local public service delivery.

A second mechanism that might account for the gains in cognitive skills of ninth- and 12th-graders is human capital transmission from their parents. To investigate the existence of intergenerational spillovers we rely on a socioeconomic questionnaire that was administered jointly with the ENEM and Prova Brasil tests, allowing us to measure education levels of test-takers' parents. Our results are consistent with some role for parental education, as we find that students from communities that benefited from extra federal transfers in the early 1980s are more likely to have a parent with some middle or high school and less likely to have parents with only a primary school education.

A third potential mechanism for the test score gains is increased parental income. As additional federal funding led to increased public spending not only on education but also on transportation and housing and urban infrastructure, impacts on test score performance in the late 2000s may arise through higher incomes even among those parents whose schooling levels were not affected by the funding boost. Indeed, Litschig and Morrison (2013) calculate that only about two percentage points of the four percentage point poverty reduction they find in 1991 is plausibly accounted for by the education channel alone, leaving the remaining two percentage points to unmeasured improved public service provision overall. Since the poverty reduction of about four percentage points persists in 2000, unmeasured and persistent public service improvements might have played a role in raising children's cognitive skills by increasing household income. Consistent with this interpretation, we find that ENEM test-takers in beneficiary communities have a lower likelihood of coming from poor households, although part of this effect might also be due to higher parental schooling.

To summarize, this work shows that a temporary increase in federal transfers to local governments led to long-lasting schooling, literacy, and income gains of directly exposed cohorts, as well as gains in cognitive skills of children and youth born after the extra transfers had expired. Available evidence on mechanisms indicates that the magnitude of these cognitive gains is plausibly accounted for by reduced class size in primary school, intergenerational spillovers, and household income gains as further discussed below. Together, these results provide the first evidence on how additional resource transfers to local governments can impact human development outcomes in a typical developing country setting in the long run.

Existing studies on impacts of community-level public revenue windfalls in Brazil, such as Monteiro and Ferraz (2010), Brollo et al. (2013), Caselli and Michaels (2013), or Gadenne (2017), look at short-run effects and do not consider education outcomes. While we find income gains for the poor, Caselli and Michaels (2013) find negligible impacts on the poverty rate (the other studies do not look at poverty). The positive effects on schooling and income reported here are quantitatively similar to those found in Olsson and Valsecchi (2015) for Indonesia and are qualitatively consistent with older studies that look at the links between school resources, educational attainment, and earnings.[1] More recent field-experimental work has focused on governance or incentive reforms but sometimes has also documented short-run education gains from additional school resources alone. Examples include Muralidharan and Sundararaman (2011), who find a positive impact of school-level block grants on test scores in India, and Olken, Onishi, and Wong (2014), who find a positive impact of village-level block grants on enrollment in Indonesia. Pradhan et al. (2014) estimate a marginally significant positive impact of school-level block grants on test scores in Indonesia. Blimpo and Evans (2011) and Duflo, Dupas, and Kremer (2015) find no effect of school-level grants on test scores in Gambia and Kenya, respectively.

Also closely related to our study are a number of papers that investigate long-run impacts on cognitive skills due to reduced class size in kindergarten and early primary school (Krueger and Whitmore 2001) or late primary school (Fredriksson, Öckert, and Oosterbeek 2013). There is also a relevant literature on compulsory schooling reforms—typically accompanied by increased school resources—showing that schooling increased not only in the present generation, but the reforms also affected education outcomes of the next generation, even if the estimated intergenerational spillover effects tend to be small in magnitude (Oreopoulos, Page, and Stevens 2006; Black, Devereux, and Salvanes 2005; Holmlund, Lindahl, and Plug 2011). Our study builds on the class size and intergenerational effects identified in these papers as two potential mechanisms that could account for the long-run impacts of intergovernmental grants on cognitive skills we document. The parameters we identify in this study, however, should be thought of as policy effects that incorporate these observed mechanisms, as well as potentially other unobserved public service improvements, such as improved local roads, as further discussed in Section III below.

The paper proceeds as follows. Section II describes the role of local governments in public service provision in Brazil and gives institutional background on revenue sharing. Section III provides a conceptual framework and discusses identifying assumptions. Section IV describes the data. Section V details the estimation approach and evaluates the internal validity of the study. Sections VI–IX present main results and

1. See for example Behrman and Birdsall (1983), Birdsall (1985), and Behrman, Birdsall, and Kaplan (1996) for Brazil; Case and Deaton (1999) for South Africa; and Duflo (2001) for Indonesia.

evidence on mechanisms. Section X evaluates the plausibility of the test score gains, and Section XI discusses alternative mechanisms. We conclude by assessing the external validity of our findings.

II. Background

A. Local Public Services and Their Financing

Local (municipal) government responsibilities at the beginning of the 1980s were mostly to provide elementary education, housing and urban infrastructure, as well as local transportation services. The responsibility for elementary education—consisting of primary school (Grades 1–5) and middle school (Grades 6–9)—was shared with state governments, while the federal government was primarily involved in financing and standard setting. In 1980, 55 percent of all elementary school students in Brazil were enrolled in state-administered schools, 31 percent in municipality schools, and the remaining 14 percent in private schools. In small and rural municipalities, such as those considered here, the proportion of students in schools managed by local governments was 74 percent, while the proportions for state-run and private schools were 24 percent and 2 percent, respectively (World Bank 1985).

In the 1980s local governments managed about 17 percent of public resources in Brazil (Shah 1991), about 4 percent of GDP, with 20 percent of local budgets going to education and similar shares to housing and urban infrastructure, and transportation spending, as shown in Litschig and Morrison (2013, their Table 1). Most of these resources accrued to local governments through intergovernmental transfers, since most municipalities have never collected much in the way of taxes. The most important among these transfers was the federal Fundo de Participação dos Municípios (FPM), a largely unconditional revenue sharing grant funded by federal income and industrial products taxes.[2] The FPM transfers were the most important source of revenue for the relatively small local governments considered here, amounting to about 50 percent on average and 56 percent in rural areas, defined as those with below-median percentage of residents living in urban areas.

B. Mechanics of Revenue Sharing in Brazil

In order to estimate the impact of intergovernmental transfers on outcomes, we exploit variation in FPM funding at several population cutoffs using regression discontinuity analysis. The critical feature of the FPM revenue-sharing mechanism for the purposes of our analysis is Decree 1881/81, which stipulates that transfer amounts depend on municipality population in a discontinuous fashion. More specifically, based on municipality population estimates, pop^e, municipalities are assigned a coefficient $k = k(pop^e)$, where $k(.)$ is the step function shown in Table 1. For municipalities with up to 10,188 inhabitants, the coefficient is 0.6; from 10,189 to 13,584 inhabitants, the coefficient is

2. This constraint is usually considered nonbinding, in that municipalities typically spend about 20 percent of their total revenue on education. It is not clear how this provision was enforced in practice, since there is no clear definition of education expenditure,s and accounting information provided by local governments was not systematically verified.

Table 1
Brackets and Coefficients for the FPM Transfer

Population Bracket	Coefficient
Up to 10,188	0.6
10,189–13,584	0.8
13,585–16,980	1
16,981–23,772	1.2
23,773–30,564	1.4
30,565–37,356	1.6
37,357–44,148	1.8
44,149–50,940	2
50,941–61,128	2.2
61,129–71,316	2.4
71,317–81,504	2.6
81,505–91,692	2.8
91,693–101,880	3
101,881–115,464	3.2
115,465–129,048	3.4
129,049–142,632	3.6
142,633–156,216	3.8
Above 156,216	4

Source: Decree 1881/81.

0.8; and so forth. The coefficient $k(pop^e)$ determines the share of total FPM resources, rev_t, distributed to municipality m in year t according to the following formula:

$$FPM_{mt} = \frac{k\left(pop_m^e\right)}{\sum_m k_m} rev_t$$

This equation makes it clear that estimates of local population should be the only determinant of cross-municipality variation in FPM funding in a given year. Exact municipality population estimates are only available for census years or years when a national population count is conducted. Transfers were allocated based on 1980 census population from 1982 (the first year the 1980 census figures were used) until 1985.[3] Previously, from 1976 to 1981, the transfers had been based on extrapolations from the 1960 and 1970 censuses, produced by the national statistical agency, IBGE.[4] Likewise, from 1986 to 1988, the transfers were also based on such extrapolations, this time based

3. The 1985 official estimates were already based on extrapolations that resulted in minor changes compared to the 1980 census numbers.
4. The methodology used by the statistical agency in principle ensures that population estimates are consistent between municipalities, states, and the updated population estimate for the country as a whole (Instituto Brasileiro de Geografia e Estatística 2002).

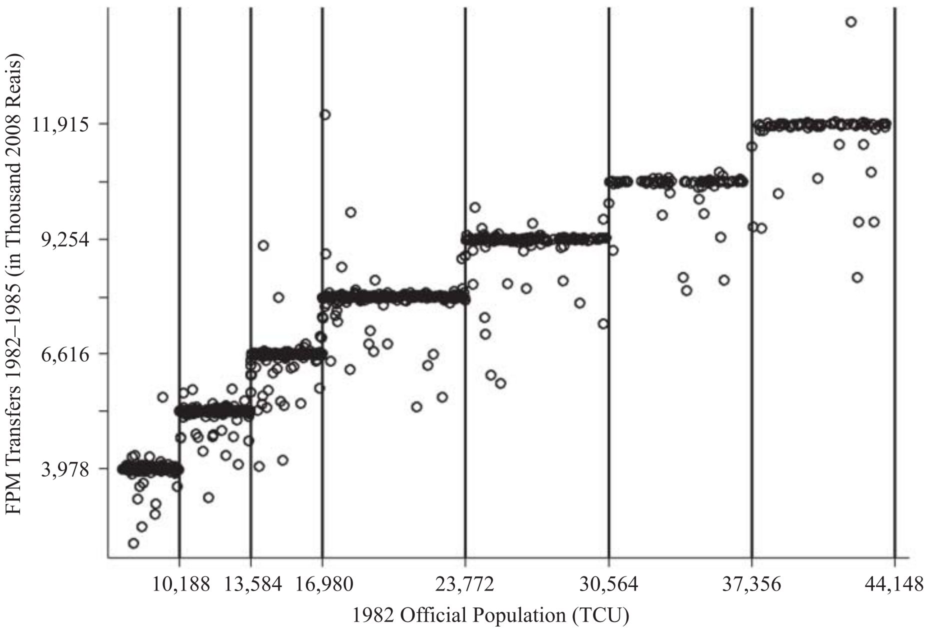

Figure 1
FPM Transfers 1982–1985 (in Thousand 2008 Reais)

Source: From Litschig and Morrison (2013, Figure 1).
Notes: Each dot represents a municipality. FPM transfers are self-reported by municipalities. 1982 official population is based on the 1980 census conducted by the national statistical agency, IBGE.

on 1970 and 1980 census population figures. Beginning in 1989, these extrapolations were updated on a yearly basis, which is still the practice at the time of this writing.

Figure 1 plots cumulative FPM transfers over the period 1982–1985 against 1982 official population. The ticks on the vertical axis correspond to the modal levels of cumulative transfers for each bracket in our data. The figure shows that funding jumps by about 1,320,000 reais (2008 prices) or about 1,000,000 international US$ at each threshold over this period.[5] Observations that appear above or below the horizontal lines are most likely due to measurement error because transfer data in this figure are self-reported by municipalities, rather than based on administrative records of the Ministry of Finance, which are not available for the period considered.[6] For the municipalities in our estimation sample, the cumulative transfer differential over the 1982–1985 period corresponds to about 2.5 percent of annual GDP in rural areas of the country and about 1.4 percent of annual GDP in urban areas.

5. The 2005 Real/$ PPP exchange rate was about 1.36.
6. For later periods the data are available from the Ministry of Finance, and in these data there is essentially no variation in FPM transfers for a given state and population bracket.

As is apparent from Figure 1, the transfer jump is the same in absolute terms at each cutoff, but the jumps decline in per capita terms the higher the cutoff. Transfers per capita jump by about R$130 (US$95) at the first threshold, R$97 (US$70) at the second, R$78 (US$57) at the third and decline monotonically for the following cutoffs. Immediately to the left of the first three cutoffs, per capita FPM funding is about R$390 (US$286), and this amount declines monotonically for the following cutoffs. For the first three cutoffs, the funding increase per capita is therefore from the same baseline level and represents about 33 percent at the first, 25 percent at the second, and 20 percent at the third cutoff. Though the differences are not great, this means that the treatment in terms of additional per capita funding is not exactly the same across these cutoffs. However, since there are likely to be economies of scale in the provision of local public services—that is, unit costs decline with scale—the differences in treatment across cutoffs might be even smaller than what the per capita funding jumps would suggest. It thus seems reasonable as a first approximation to expect similar treatment effects around these cutoffs.

Online Appendix Figure 1 plots cumulative FPM transfers over the period 1986–1989 against 1982 official population. As in Figure 1, each dot represents a municipality. In contrast to the 1982–1985 period, however, the conditional variance of FPM funding is large, and there are no visible funding discontinuities at any of the six cutoffs. It is important to note that the allocation formula has remained essentially unchanged over this period.[7] Funding discontinuities that were present from 1982 to 1985 disappeared in 1986 because many municipalities changed brackets due to decreases or, more often, increases in their population relative to 1980. Online Appendix Table 1 shows that there are no economically or statistically significant differences in FPM transfers between the treatment and comparison group (those around the first three cutoffs based on the 1980 census) in each year from 1986 onwards until 1989, as well as cumulatively for the entire 1986–1989 period. The "treatment" therefore consists of a (presumably) unexpected temporary funding windfall to the municipal budget, lasting for four years from the beginning of 1982 through the end of 1985.

Even a temporary funding boost might affect future budgets through, for example, increased own revenue collection based on higher local economic activity. Unfortunately, our analysis lacks statistical power when it comes to budgetary outcomes. Online Appendix Table 2 shows impact estimates on municipal own revenue during the post-windfall period. While point estimates are mostly positive, they are never statistically different from zero. Online Appendix Table 3 shows that the estimates for total revenue bounce around zero and are invariably small. But because the confidence intervals are rather wide, we cannot rule out substantive total revenue increases up to about 10 percent.

III. Conceptual Framework and Identification

A. Conceptual Framework

Because the additional FPM transfers provide unrestricted budget support, effects on schooling, learning, and earnings may arise through a variety of channels in addition to

7. From 1989 onwards the allocation formula mechanically increases the variation of funding because it incorporates a state-level share so that conditional on population there is variation in FPM funding across states.

education spending, such as improved local roads, for example. The following presents a framework for thinking about the long-run effects estimated here and for comparing them to those identified in existing studies. We consider four points in time $t=0, 1, 2, 3$ that are separated by a decade, which in our setting corresponds to the Brazilian census years 1980, 1991, 2000, and 2010. FPM transfers F temporarily increase for some communities between $t=0$ and $t=1$. For school-age cohorts in the early 1980s, we measure schooling S at $t=1, 2, 3$. We also measure income I in the community overall at $t=1, 2, 3$. For cohorts born after the funding boost expired, we measure academic achievement A, parental schooling S^P, and parental income I^P at $t=3$.

Assume that schooling in subsequent periods depends on initial period public spending on education E and on other local public services O, such as local roads, which both depend on FPM transfers. A part of education spending is used to reduce class size C in primary school. Education spending, FPM transfers, and class size are all observed in all periods, while other local public services are never observed. Also assume that household income depends on schooling—which in turn depends on other local public services—and on other local public services directly. Academic achievement of children at $t=3$ depends on class size in primary school, parental schooling, parental income, and other local public services. These relations can be summarized as follows:

$$S_t = S\left\{C[E(F_0)], E(F_0), O(F_0)\right\}$$

$$I_t = I[S(.), O(F_0)]$$

$$A_t = A\left[C(F_0), S^P(.), I^P(.), O(F_0)\right]$$

Litschig and Morrison (2013) present estimates of $\partial S_1/\partial F_0$ and $\partial I_1/\partial F_0$, which are schooling and income effects of block grants in the medium run, as well as estimates of the effect on average class size $\partial C/\partial F_0$. This paper presents estimates of long-run impacts of intergovernmental transfers. For schooling and income gains of directly exposed cohorts, the parameters can be expressed as $\partial S_2/\partial F_0$, $\partial S_3/\partial F_0$, $\partial I_2/\partial F_0$, and $\partial I_3/\partial F_0$. For test score gains of the next generation, the parameter we can identify is $\partial A_3/\partial F_0$, which we decompose as follows:

$$\frac{dA_3}{dF_0} = \frac{\partial A_3}{\partial C}\frac{\partial C}{\partial F_0} + \frac{\partial A_3}{\partial S^P}\frac{\partial S^P}{\partial F_0} + \frac{\partial A_3}{\partial I^P}\frac{\partial I^P}{\partial F_0} + \frac{\partial A_3}{\partial O}\frac{\partial O}{\partial F_0}$$

This parameter should be thought of as a policy effect that incorporates both observed and unobserved mechanisms. Regarding the class size channel $\partial C/\partial F_0$, we provide evidence of a class size reduction in public primary school that lasts for more than 20 years after the initial extra funding had stopped. We also provide evidence on schooling and income gains of test-takers' parents as given by $\partial S^P/\partial F_0$ and $\partial I^P/\partial F_0$, respectively, which are consistent with intergenerational human capital spillovers. Combined with estimates of $\partial A_3/\partial C$, $\partial A_3/\partial S^P$, and $\partial A_3/\partial I^P$, both from Brazil and from prior studies, the class size and intergenerational channels account for most if not the entire next-generation test score gain $\partial A_3/\partial F_0$ we estimate in this study, as further discussed in Section X below. We also investigate but find no evidence of persistent improvements in other measures of education supply.[8] Nonetheless, there might be other unmeasured persistent public service improvements $\partial O/\partial F_0$ that explain part of the gains in cognitive skills. We note in particular that we cannot identify intergenerational spillover effects on schooling $\partial S_3/\partial S^P$, as in Black, Devereux, and Salvanes (2005) and Holmlund, Lindahl, and Plug (2011), or intergenerational effects on student test performance $\partial A_3/\partial S^P$, as in Oreopoulos, Page, and Stevens (2006) and Carneiro, Costas, and Parey (2013).

B. Identification

Three identifying assumptions are required to recover the policy parameters discussed above. The first is that unobserved determinants of outcomes vary smoothly as a function of population (if at all) and, in particular, do not jump at the cutoffs. As shown in Lee and Lemieux (2010), the assumption that individual densities of the treatment-determining variable are smooth is sufficient for the continuity of unobservables. In our case, this assumption does not preclude that mayors or other agents in the municipality have some control over their particular value of population. As long as this control is imprecise, treatment assignment is randomized around the cutoff. In our case, the continuity of individual population density functions also directly ensures that treatment status (extra transfers) is randomized close to the cutoff. An additional concern would be imperfect compliance with the treatment rule, but in our study period all eligible municipalities received more FPM transfers, and none of the ineligible ones did. As discussed in more detail in Litschig and Morrison (2013), the key continuity assumption is likely to hold here because mayors did not know the exact locations of the thresholds until after the release of the 1980 census results. Litschig and Morrison (2013) also provide corroborating internal validity tests and robustness checks.

The second assumption is the exclusion restriction. The potential concern here is that other government policies are also related to the cutoffs specified in Decree 1881/81. If so, we would identify the combined causal effect of extra funding and other policies. To our knowledge, however, there are no other programs that used the same cutoffs in the early 1980s. Moreover, if total spending is the only channel through which additional transfers operate, the estimates presented here identify long-run impacts of local public spending, rather than effects of intergovernmental transfers. Reductions in local taxes and corresponding increases in private consumption would violate the exclusion restriction for example. Empirically, local taxes do not seem to have responded to additional transfers as further detailed in Litschig and Morrison (2013).

The third assumption is that there is no selective attrition or sample selection at the cutoff. As discussed in Rosenzweig and Wolpin (1988), sample selection bias is a particularly important concern with site-specific programs, such as the extra funding to local governments considered here. Imperfect control over 1980 population ensures that initial distributions of unobserved determinants of outcomes are identical close to the cutoff. But our analysis compares average outcomes of resident populations 10, 20, and 30 years after the initial disbursement of funds, including both native residents who stayed, as well as in-migrants. The potential threat is thus that unobservables of both

8. For example, we find no evidence of discontinuities in the proportion of primary school teachers with some college education or availability of internet access and IT or science labs. Results are available upon request.

migrants and natives who continue to reside in the municipality are systematically different in municipalities immediately around the cutoff, although it is worth bearing in mind that treatment and comparison communities are typically quite far apart geographically. Nonetheless, even nonselective migration might mechanically attenuate impact estimates if the proportion of native residents decreases over time and the new in-migrants have the same average outcomes in both treatment and comparison communities. We provide several pieces of evidence suggesting that sample selection is unlikely to bias the results presented below. Please refer to the Online Appendix for further discussion and corresponding results.

IV. Data

A. *Official Population, FPM Transfers, and Covariates*

Our analysis draws on multiple data sources from several time periods. Population estimates determining transfer amounts over the period 1982–1989 were taken from successive reports issued by the Tribunal de Contas da União (TCU). Data on FPM transfers are self-reported by municipal officials and compiled into reports by the Secretariat of Economics and Finance inside the federal Ministry of Finance. The data from these reports were entered into spreadsheets using independent double-entry processing. Data on FPM transfers were converted into 2008 currency units using the GDP deflator for Brazil and taking account of the various monetary reforms that occurred in the country since 1980.

We include as pre-treatment covariates the 1980 levels of municipality income per capita, average years of schooling for individuals 25 years and older, the percentage of people over 14 years old who are illiterate, the infant mortality rate, the school enrollment rate of 7–14-year-olds, and the percent of the municipal population living in urban areas. Data on these 1980 municipality characteristics are based on the long-form sample of the census and have been calculated by the national statistical agency.[9] The 1980 poverty headcount ratio was calculated by the government research institute Instituto de Pesquisa Econômica Aplicada (IPEA). The poverty line in 1980 was about R$95 in August 2000 reais. Electoral data for the municipal executive elections in 1996, 2000, 2004, and 2008 are from the Supreme Electoral Tribunal.

B. *Schooling and Literacy of Directly Exposed Cohorts*

For education outcomes of cohorts directly affected by the increase in federal transfers, we use the long-form samples of the 1991, 2000, and 2010 population censuses to compute municipal-level average years of schooling (that is, grades completed, not just "years in school") and the percent literate. For 1991 and 2000 the census forms allow us to compute years of schooling directly based on completed grades. For 2010 we compute an individual's schooling based on highest grade enrollment and impute schooling using the 2000 census in case the highest-grade enrollment was not completed. Details of the

9. The 1980 census had one long enumeration form that was applied to 25 percent of the population and a shorter census form that was administered to 75 percent of the population.

imputation process are given in Online Appendix Tables 4 and 5. The resulting (likely random) measurement error in schooling would reduce the precision of impact estimates for 2010.

We focus on two cohorts, 0–9-year-olds and 10–19-year-olds in 1982 when the extra transfers started. The older cohort was aged 19–28 years in 1991, 28–37 in 2000, and 38–47 in 2010. This was the cohort most likely affected by the public spending increase from 1982 to 1985, since the 19-year-olds in 1991 were about ten years old in 1982 and hence in the middle of elementary schooling age (7–14), while the 29-year-olds were at least 19 years old (age 20 on September 1, 1982 but 19 at some point during the year 1982 for some) and hence ineligible to attend regular elementary school, which has a cutoff age at 18. By 1991, most of the 19–28-year-olds likely had completed most of their education, so we should be able to capture most of any effect on their level of schooling.

We also compute average years of schooling and the literacy rate for the cohort that was 0–9 years old in 1982, 9–18 in 1991, 18–27 in 2000, and 28–37 in 2010 because local governments in Brazil also provided preschool education and daycare services that could have benefited even the newborn cohort in 1982. One would expect this younger age group to exhibit a smaller treatment effect by 1991 (at least in absolute terms) because most of them were not of elementary schooling age when spending increased in 1982. Moreover, most of this cohort had not completed elementary school in 1991, so the 1991 census might fail to capture part of the impact on their level of schooling if the increased spending produced school supply improvements that had not faded completely by 1991. By ages 18–27 in 2000, most individuals in this cohort likely had completed most of their schooling careers.

C. Test Scores, School Enrollment, and Test Participation of the Next Generation

For education outcomes of the next generation, we rely on two standardized nationwide tests, Prova Brasil and ENEM (Exame Nacional do Ensino Médio), which started to be administered in the late 2000s. Prova Brasil is a compulsory test taken by students at the end of primary and middle school. We use the microdata for the 2007, 2009, and 2011 rounds of the test, each with more than two million observations per grade. In 2007 the test was given at the end of fourth and eighth grade, while from 2009 onwards the test has been given at the end of fifth and ninth grade due to a compulsory schooling extension. Student performance is measured in two subjects: Portuguese language (reading) and mathematics (problem solving). Prova Brasil covers all public schools that enroll at least 20 students.[10] We calculate z-scores with mean zero and standard deviation one by year, grade, and discipline on the universe of test-takers. We add up the two standardized scores for each subject and again standardize it to get a total z-score for each individual. We then pool these total z-scores across years for a given grade and compute the mean, median, 10th, 25th, 75th, and 90th percentiles of the individual-level total z-score distribution for each municipality. To assess the potential for sample selection bias, we use the 2010 census to compute primary and middle school net enrollment rates for both

10. Prova Brasil 2007 was applied to children in fourth and eighth grades in urban but not rural schools.

public and private schools and for public schools only. Primary school enrollment in 2010 is about 98 percent overall and 92 percent for public schools. Middle school enrollment is about 96 percent overall and also 92 percent for public schools.

ENEM is an annual exam designed for students in their final year of high school and high school graduates. Its original goal up until 2008 was to provide a reference for self-evaluation of the student's capabilities, and it was used as an input in the selection process of a few universities. From 2009 onwards, ENEM gained in importance as it became the unified entrance exam for the federal universities system, which provides tuition-free college education. In our analysis, we select only test-takers graduating from high school the year the test was taken. These adolescents are typically 17 or 18 years old and represent about one-third of all ENEM test-takers. Focusing on high school graduating cohorts allows us to compare our results to other studies and to compute a meaningful participation rate by dividing the total number of ENEM test-takers between 2007 and 2011 by the total number of individuals aged 16–21 in the 2010 census, which corresponds to 17- or 18-year-olds at the time of the respective test-year. The ENEM participation rate among individuals aged 16–21 in 2010 is about 21 percent. We again standardize test scores by year and pool together all five years to compute the mean, median, 10th, 25th, 75th, and 90th percentiles of the individual-level test score distribution for each municipality.

D. Poverty Rates

Poverty headcount ratios for 1991 and 2000 were computed by the government research institute IPEA using census data based on a poverty line of R$75.5 per month in August 2000 prices. We computed the 2010 poverty rate ourselves based on a poverty line of R$146.5 in July 2010 prices, corresponding to IPEA's poverty line adjusted for national inflation. All poverty rates use household income per capita as the measure of individual-level income.

E. Parental Education of ENEM and Prova Brasil Test-Takers

To investigate intergenerational education spillovers we rely on a socioeconomic questionnaire that was administered jointly with the ENEM and Prova Brasil tests, allowing us to measure parental education levels in the late 2000s. For ENEM test-takers we restrict the sample to high school graduating cohorts, as we did for the computation of test score statistics. We aggregate responses into four categories, depending on the highest education level reached by the most educated parent, and compute the municipality-level proportion of respondents falling into each category. For Prova Brasil the categories are: "no more than primary school" (completed Grades 0–4, including respondents who did not know the education level of their parents), "some middle school" (completed Grades 5–7), "some high school" (completed Grades 8–11), and "completed at least high school" (completed Grades 12 or above). For ENEM the first two categories are the same, but due to differences in survey response categories, the "some high school" category includes high school graduates (completed Grades 8–12), while the highest category is "some college" (completed Grades 13 and above). For Prova Brasil there are sometimes substantial numbers of test-takers who did not fill out the socioeconomic survey (on average 16 percent for fifth-graders and 20 percent for ninth-graders). We verify that the proportion of nonrespondents exhibits no jump at the cutoffs and disregard these individuals in the computation of parental education.

F. Household Income of ENEM Test-Takers

The socioeconomic questionnaire associated with ENEM includes a question about household income (the Prova Brasil questionnaire does not). We again restrict the sample to high school graduating cohorts. Response categories are comparable only for the years 2007–2009 and are in multiples of the nominal monthly Brazilian minimum wage. We aggregate responses into four categories, corresponding to household income up to one minimum wage, between one and two minimum wages, between two and five minimum wages, and above five minimum wages, and compute the municipality-level proportion of respondents in each category. One minimum wage was R$380, R $415, and R$465 in 2007, 2008, 2009, respectively. Responses are missing for about 1 percent of test-takers, and we disregard these individuals in the computation of income categories.

G. Student-to-Teacher Ratios

From available annual school censuses for 1991–2011 we draw the student-to-teacher ratio in public primary school (Grades 1–4 up to the mid-2000s and Grades 1–5 thereafter) aggregated by municipality. We focus on student-to-teacher ratios in public primary schools because later grades are frequently managed not by the municipality but by state governments. Unfortunately, earlier rounds of the school census are not available, and even post-1991 some census years are not available. Moreover, school census information is missing for about 10–20 percent of municipalities in some years, although the probability of being missing is smooth at the cutoff (results available upon request). In order to save space without dropping many municipalities, we compute average student-to-teacher ratios for adjacent years with similar coverage, resulting in average class size measures for 1991, 1995–1996, 1997–2003, 2004–2006, and 2007–2011.

V. Estimation Approach

Following Hahn, Todd, and Van der Klaauw (2001) and Lee and Lemieux (2010), we use local linear regressions as our main estimation approach. We focus on the first three population cutoffs ($c_1 = 10{,}188$, $c_2 = 13{,}584$, and $c_3 = 16{,}980$). At subsequent cutoffs the variation in FPM transfers is too small to affect municipal overall budgets, and hence there is no "first stage" in terms of overall resources available for the municipality (see Section VI in Litschig and Morrison 2013). For our pooled analysis, we need to make observations comparable in terms of the distance from their respective cutoff. Let pop_{ms} denote population in municipality m and state s and seg_j, with $j = 0, 1, 2, 3$ the four integers (7,500, 11,800, 15,100, and 23,772) that bound and partition the population support into three segments. We rescale population to equal zero at the respective thresholds within each of the first three segments, and then use the normalized variable, X_{ms} for estimation purposes:

$$X_{ms} = \begin{cases} pop_{ms} - c_1 \text{ if } seg_0 < pop_{ms} \leq seg_1 \\ pop_{ms} - c_2 \text{ if } seg_1 < pop_{ms} \leq seg_2 \\ pop_{ms} - c_3 \text{ if } seg_2 < pop_{ms} \leq seg_3 \end{cases}$$

Let Y_{ms} denote an outcome, $\mathbf{z}_{ms}$ a set of pretreatment covariates, a_s a fixed effect for each state, and U_{ms} the influence of unobserved factors on outcomes. Neither covariates nor state fixed effects are needed for identification. We include them to guard against chance correlations with treatment status and to increase precision of the estimates. The ordinary least squares (OLS) specification we use is:

$$(1)\quad Y_{ms} = \tau 1[X_{ms} > 0] + \left\{\alpha_{10} X_{ms} + \alpha_{11} X_{ms} 1[X_{ms} > 0]\right\} 1_{1p}$$

$$+ \left\{\alpha_{20} X_{ms} + \alpha_{21} X_{ms} 1[X_{ms} > 0]\right\} 1_{2p}$$

$$+ \left[\alpha_{30} X_{ms} + \alpha_{31} X_{ms} 1[X_{ms} > 0]\right] 1_{3p}$$

$$+ \sum_{j=1}^{3} \beta_j 1\left[seg_{j-1} < pop_{ms} \leq seg_j\right] 1_{jp} + \gamma \mathbf{z}_{ms} + a_s + U_{ms}$$

$$1_{jp} = 1\left[c_j(1-p) < pop_{ms} < c_j(1+p)\right], \qquad j = 1, 2, 3; \qquad p = 2, 3, 4 \text{ percent}$$

$$1_p = 1_{1p} + 1_{2p} + 1_{3p}$$

Essentially, Equation 1 allows for six different slopes, one each on either side of the three cutoffs, but imposes a common effect τ. Under the three identifying assumptions from Section III.B above, the pooled treatment effect is given by $\lim_{\Delta \downarrow 0} E[Y_{ms}|X = \Delta] - \lim_{\Delta \uparrow 0} E[Y_{ms}|X = \Delta] = \tau$. All the tables below show results for the first three cutoffs pooled and for successively larger samples around the cutoffs ($p = 2$, 3, and 4 percent), for each sample with and without covariates. Those estimates that control for covariates are probably the most reliable because they control for chance correlations with treatment status. They are also typically the most precisely estimated because covariates absorb some of the variation in outcome measures. In order to benchmark the magnitude of impact estimates, we also report the intercept estimate from a linear spline in the normalized running variable without other covariates, which corresponds to the estimated conditional mean outcome at $X = 0$.

In addition to OLS results, we show estimates and standard errors from the Imbens and Kalyanaraman (2012) and Calonico, Cattaneo, and Titiunik (2014) optimal bandwidth choice procedures with triangular kernels based on the Stata routine "rdrobust." We report (weighted) least squares estimates and heteroskedasticity-robust standard errors based on the Imbens–Kalyanaraman optimal bandwidths and Calonico–Cattaneo–Titiunik bias-corrected estimates with standard errors that are robust to "large" bandwidths based on their optimal bandwidths. Because the "rdrobust" routine does not accommodate covariates, we run it with residual outcomes where the residual comes from a regression of the outcome variable on state fixed effects, segment dummies, and pretreatment covariates from the 1980 census: municipality income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas. As shown in Lee and Lemieux (2010), this "residualizing" approach allows for consistent estimation of the effect as long as the continuity assumption holds.

Following Lee and Lemieux (2010), our figures plot outcome residuals against normalized population. Intuitively, since outcome residuals are by construction uncorrelated with fixed effects and pretreatment covariates, any discontinuity in outcome residuals at the cutoff cannot be driven by chance correlations with these covariates. Raw data plots tend to produce larger discontinuity estimates and are invariably much noisier. These plots are available upon request.

VI. Impacts on 1980s School-Age Cohorts

A. Impacts on Schooling—1991, 2000, and 2010

Table 2 shows results for average years of schooling (completed grades) for individuals 9–18 and 19–28 years of age in 1991 and for the same two cohorts in 2000 and 2010. OLS estimates with pretreatment covariates, as well as optimal bandwidth–based estimates with residualized schooling shown in the first row, suggest that the older cohort accumulated on average about 0.3 additional years of schooling by 1991. While the inclusion of pretreatment covariates systematically attenuates impact estimates, the confidence intervals show substantial overlap. For example, the 95 percent confidence interval for the effect on the older cohort based on the 4 percent neighborhood without pretreatment covariates is about [0.19, 0.87], while with covariates the confidence interval is about [0.07, 0.48]. Results in Rows 2 and 3 show that schooling gains of about 0.2 years and 0.1 years persist in 2000 and 2010, respectively. While the estimates in 1991 are statistically significant (at 1 percent) even within a relatively small neighborhood of ±3 percent around the cutoffs, estimates in 2000 are only marginally significant at 10 percent, and in 2010 the estimates typically cannot be distinguished from zero. Since schooling outcomes for 2010 had to be imputed for some individuals, it is unsurprising that standard errors tend to be largest for the 2010 impact estimates.

Although the schooling gains in 1991 and 2000 are statistically indistinguishable, the attenuation of estimated gains would be consistent with three out of ten individuals from treatment communities completing an additional year of schooling by 1991 and one out of ten individuals from comparison communities eventually completing an additional year of schooling by 2000, for example. In fact, given that average schooling for the 1991 19–28-year-old cohort in comparison communities increased by about 0.6 years by 2000 (see comparison means), a more accurate interpretation is that average schooling in marginal treatment communities only increased by about 0.5 between 1991 and 2000, leaving them with an average 0.2 year educational advantage by 2000.

Table 2

Impacts on Years of Schooling for 1980s School-Age Cohorts in 1991, 2000, and 2010

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
Avg. schooling (19–28-year-olds in 1991)	4.26								
$I[X>0]$		0.322	0.225	0.516	0.301	0.528	0.275	0.288	0.343
		(0.260)	(0.151)	(0.198)	(0.114)	(0.171)	(0.102)	(0.097)	(0.111)
R^2/observations		0.71	0.89	0.71	0.89	0.69	0.88	457	420
Avg. schooling (28–37-year-olds in 2000)	4.88								
$I[X>0]$		0.143	0.062	0.381	0.188	0.430	0.193	0.155	0.204
		(0.225)	(0.163)	(0.182)	(0.128)	(0.161)	(0.116)	(0.099)	(0.118)
R^2/observations		0.73	0.86	0.70	0.86	0.69	0.86	617	556
Avg. schooling (38–47-year-olds in 2010)	5.15								
$I[X>0]$		0.091	−0.004	0.322	0.095	0.376	0.130	−0.017	0.047
		(0.233)	(0.182)	(0.197)	(0.149)	(0.171)	(0.127)	(0.079)	(0.112)
R^2/observations		0.63	0.81	0.59	0.79	0.59	0.80	1051	717
Avg. schooling (9–18-year-olds in 1991)	2.61								
$I[X>0]$		0.207	0.155	0.287	0.166	0.288	0.136	0.131	0.124
		(0.157)	(0.095)	(0.117)	(0.071)	(0.099)	(0.062)	(0.054)	(0.053)
R^2/observations		0.84	0.94	0.83	0.93	0.81	0.93	575	735

(*continued*)

Table 2 (*continued*)

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
Avg. schooling (18–27-year-olds in 2000)	5.78								
$I[X>0]$		0.174	0.108	0.385	0.200	0.436	0.229	0.104	0.228
		(0.245)	(0.182)	(0.186)	(0.139)	(0.158)	(0.118)	(0.096)	(0.123)
R^2/observations		0.74	0.86	0.73	0.86	0.73	0.86	714	556
Avg. schooling (28–37-year-olds in 2010)	6.39								
$I[X>0]$		0.046	–0.030	0.321	0.121	0.449	0.221	0.066	0.081
		(0.223)	(0.193)	(0.189)	(0.153)	(0.168)	(0.134)	(0.101)	(0.116)
R^2/observations		0.65	0.79	0.61	0.78	0.60	0.78	806	835
Observations		202	199	297	294	391	387		

Notes: Authors' calculations of average schooling are based on long-form samples from the 1991, 2000, and 2010 censuses. Census 1991 results from Litschig and Morrison (2013, Tables 5 and 6). All specifications pool across the first three cutoffs and include state fixed effects and segment dummies. Columns 1–6 are OLS local linear discontinuity estimates with robust standard errors in parentheses. OLS specifications allow for differential slopes by segment and on each side of the cutoff. Columns 7 and 8 give estimates and standard errors based on Imbens–Kalyanaraman (IK) and Calonico–Cattaneo–Titiunik (CCT) optimal bandwidths. Neighborhood (percent) is percent distance from respective cutoff. Pretreatment covariates (1980 census) include county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas.

Estimates for the younger cohort of 9–18-year-olds in 1991 (0–9 in 1982) shown in the fourth row of Table 2 suggest a schooling gain of about 0.15 years on average. For this younger cohort, the schooling gain tends to increase to about 0.2 years on average by 2000, which is consistent with the fact that by 1991 almost the entire cohort was still eligible for primary or middle school and that average schooling in marginal comparison municipalities more than doubled between 1991 and 2000 (from about 2.6 years of schooling on average in 1991 to about 5.8 in 2000). Estimates are statistically significant (at 5 percent) in 1991 and marginally significant at 10 percent in 2000. Between 2000 and 2010 the schooling gain attenuates and is typically indistinguishable from zero in 2010. Overall, this evidence suggests that the younger cohort had not realized the entire schooling gain by 1991 and that by 2010 impact estimates are too noisy to be informative.

Figure 2 presents graphical evidence of discontinuities in schooling at the cutoff for both cohorts in 1991 and 2000. Each dot represents the average of residual years of schooling for a given cohort, year, and bin. There are about 50 municipalities per bin. The correspondence between Panel A of Figure 2 and the results in Table 2 is that the vertical difference between the two straight lines at the cutoff illustrates the jump estimated in Row 1, Column 6 of Table 2. In addition to the linear spline, each panel shows a cubic spline fitted through individual municipalities underlying the ten dots on either side of the cutoff. With this in mind, Figure 2 shows clear evidence of discontinuities in schooling at the cutoff in both 1991 and 2000 and for both cohorts. The figure additionally shows that for neighborhoods beyond 4 percent, the linear specification might yield downward-biased estimates of the discontinuity at the cutoff because of the curvature evident in the bin averages and in the cubic approximation of the regression function.

B. Impacts on Literacy—1991, 2000, and 2010

Table 3 shows that students not only completed more grades in municipalities that received extra funds but that for some of them it made the difference between being able to read and write or not. Results are broadly similar across estimation approaches. For the older cohort the effect on literacy amounts to about four to five percentage points in 1991, compared to an average literacy rate of about 76 percent in the comparison group. The literacy gains in 2000 and 2010 are reduced to about two to three percentage points. For 1991 the estimates are highly significant (at 1 percent) and most estimates remain significant at 5 percent in 2000 and 2010. For the younger cohort, the literacy gain in 1991 is about three percentage points, compared to an average literacy rate of about 73 percent in the comparison group, as shown in Row 4. This gain is reduced to about 1.5 percentage points in 2000 and 2010. Most estimates are statistically significant at 5 percent in 1991 and marginally significant at 10 percent in 2000 and 2010. Figure 3 shows the literacy gains for both cohorts in 1991 and 2000 graphically (2010 results are available upon request). As with schooling above, the figure shows clear evidence of persistent discontinuities in literacy rates at the cutoff in both 1991 and 2000 and for both cohorts. Overall, these results suggest that the literacy gains of school-age cohorts first measured in 1991 are generally attenuated but persist in 2000 and even in 2010.

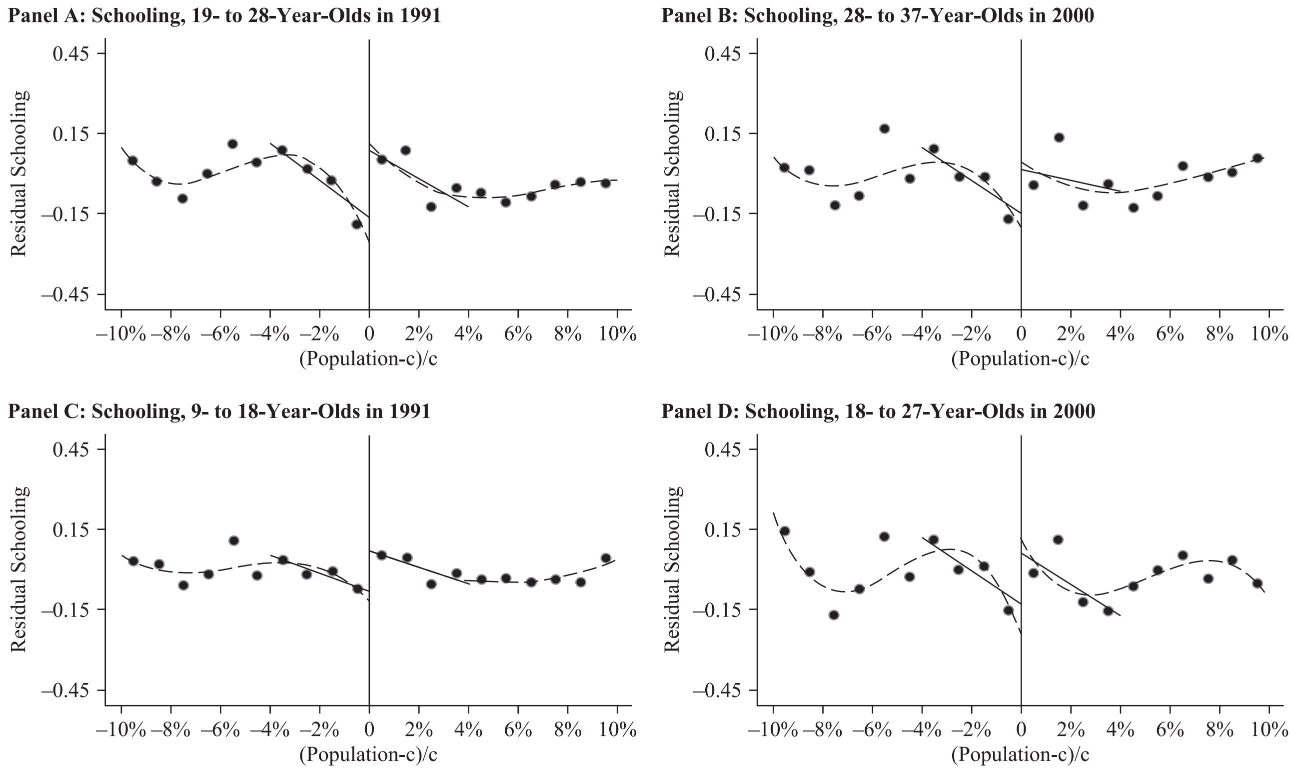

Figure 2

Impacts on Years of Schooling for 1980s School-Age Cohorts, 1991 and 2000

Source: Authors' calculations of average schooling are based on long-form samples of the 1991 and 2000 censuses. Census 1991 results from Litschig and Morrison (2013, Figure 5).

Notes: Each dot represents the sample average of residual schooling in a given bin. The residual comes from a regression of average schooling on state fixed effects, segment dummies, and pretreatment covariates from the 1980 census: county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas. The bin width is one percentage point of the respective threshold, $c = 10{,}188$; 13,584; 16,980.

Table 3
Impacts on Literacy for 1980s School-Age Cohorts in 1991, 2000, and 2010

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
Literacy rate (19–28-year-olds in 1991)	0.76								
$I[X>0]$		0.057	0.047	0.062	0.049	0.059	0.041	0.036	0.043
		(0.027)	(0.016)	(0.019)	(0.012)	(0.016)	(0.011)	(0.009)	(0.011)
R^2/observations		0.78	0.91	0.80	0.91	0.80	0.91	493	410
Literacy rate (28–37-year-olds in 2000)	0.80								
$I[X>0]$		0.023	0.013	0.041	0.029	0.039	0.022	0.018	0.022
		(0.022)	(0.014)	(0.016)	(0.011)	(0.014)	(0.010)	(0.008)	(0.010)
R^2/observations		0.78	0.91	0.79	0.90	0.79	0.90	617	564
Literacy rate (38–47-year-olds in 2010)	0.80								
$I[X>0]$		0.023	0.015	0.032	0.021	0.034	0.019	0.011	0.016
		(0.019)	(0.012)	(0.014)	(0.009)	(0.012)	(0.008)	(0.007)	(0.009)
R^2/observations		0.81	0.92	0.81	0.91	0.81	0.92	595	507
Literacy rate (9–18-year-olds in 1991)	0.73								
$I[X>0]$		0.037	0.028	0.043	0.027	0.046	0.024	0.022	0.022
		(0.028)	(0.019)	(0.020)	(0.014)	(0.017)	(0.012)	(0.010)	(0.010)
R^2/observations		0.82	0.93	0.82	0.91	0.82	0.91	566	806

(*continued*)

Table 3 (*continued*)

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
Literacy rate (18–27-year-olds in 2000)	0.88								
$I[X>0]$		0.021	0.015	0.026	0.018	0.025	0.014	0.007	0.015
		(0.016)	(0.012)	(0.011)	(0.009)	(0.010)	(0.008)	(0.006)	(0.008)
R^2/observations		0.76	0.88	0.76	0.86	0.76	0.86	703	472
Literacy rate (28–37-year-olds in 2010)	0.87								
$I[X>0]$		0.012	0.004	0.023	0.014	0.025	0.013	0.010	0.008
		(0.016)	(0.012)	(0.012)	(0.009)	(0.010)	(0.008)	(0.006)	(0.005)
R^2/observations		0.78	0.88	0.78	0.87	0.79	0.87	702	922
Observations		202	199	297	294	391	387		

Notes: Authors' calculations of literacy rates are based on long-form samples from the 1991, 2000, and 2010 censuses. Census 1991 results from Litschig and Morrison (2013, Tables 5 and 6). All specifications pool across the first three cutoffs and include state fixed effects and segment dummies. Columns 1–6 are OLS local linear discontinuity estimates with robust standard errors in parentheses. OLS specifications allow for differential slopes by segment and on each side of the cutoff. Columns 7 and 8 give estimates and standard errors based on Imbens–Kalyanaraman (IK) and Calonico–Cattaneo–Titiunik (CCT) optimal bandwidths. Neighborhood (percent) is percent distance from respective cutoff. Pretreatment covariates (1980 census) include county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas.

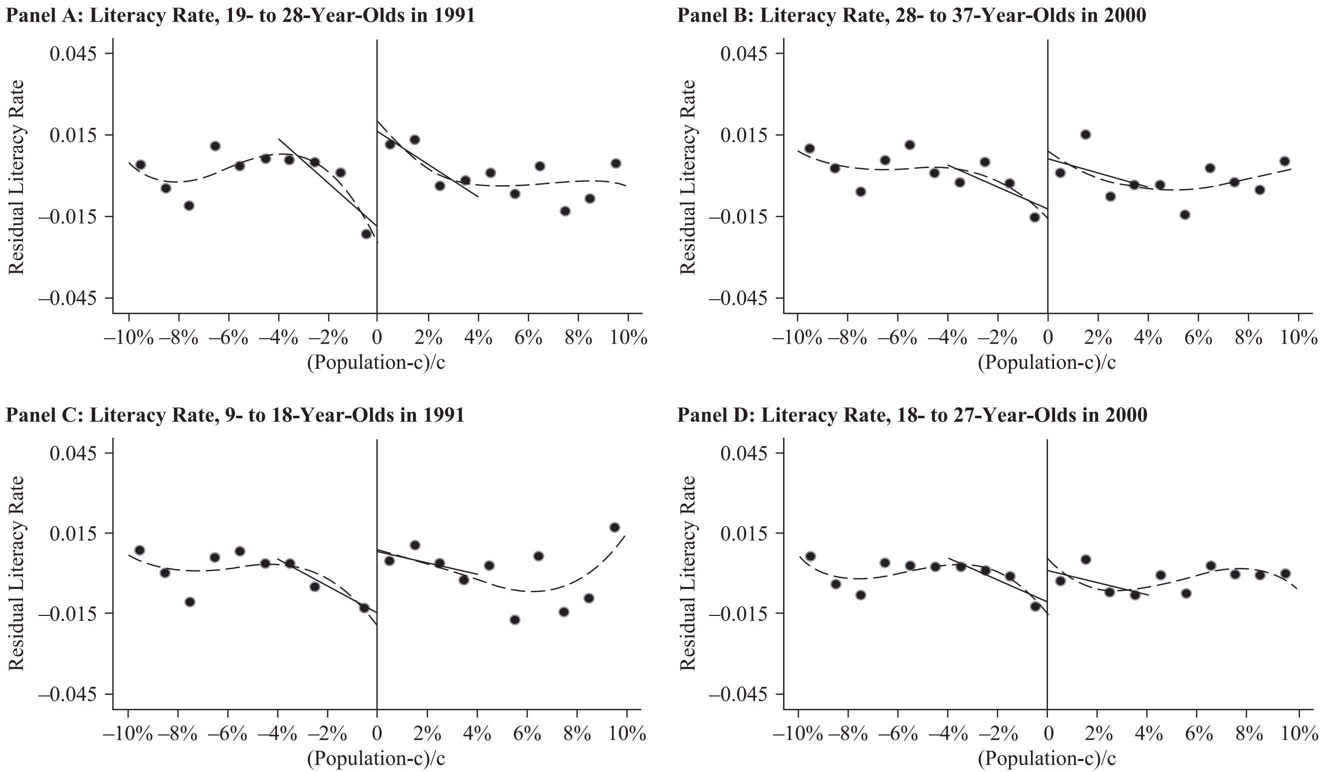

Figure 3

Impacts on Literacy Rates for 1980s School-Age Cohorts, 1991 and 2000

Source: Authors' calculations of literacy rates are based on long-form samples of the 1991 and 2000 censuses. Census 1991 results from Litschig and Morrison (2013, Figure 5).
Notes: Each dot represents the sample average of residual literacy in a given bin. The residual comes from a regression of the literacy rate on state fixed effects, segment dummies, and pretreatment covariates from the 1980 census: county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas. The bin width is one percentage point of the respective threshold, $c = 10{,}188$; $13{,}584$; $16{,}980$.

VII. Impacts on Poverty Reduction in 1980, 1991, 2000, and 2010

In this section we include all residents irrespective of age and show that the poverty reduction found in prior work using the 1991 census persists in 2000 but largely disappears by 2010. Table 4 shows estimates of impact on the poverty rate, measured relative to the national income poverty line. The first row shows that estimates for the pretreatment year 1980 are all close to zero and statistically insignificant. Estimates for 1991 and 2000 shown in Rows 2 and 3 are all negative, ranging mostly from −3 to −5 percentage points, and are typically significant at least at 5 percent. Estimates for 2010 in the bottom row mostly fall in the range of −1 to −2 percentage points and are typically not statistically significant. Panel A in Figure 4 shows that the poverty rate in 1980 is smooth at the cutoff. Panels B and C provide clear graphical evidence of a reduction in the poverty rate in 1991 and 2000, while Panel D suggests that by 2010 the discontinuity is much attenuated if not gone completely. In sum, these results suggest that the poverty reduction previously established for 1991 persisted in 2000 and largely disappeared by 2010.

In order to interpret the results on poverty reduction in 1991 and 2000, it is useful to do some back-of-the-envelope calculations. Impacts on poverty are likely to arise through better and more widespread education, as well as through better local public service provision overall. Regarding the education channel, the estimates discussed above suggest schooling gains for the 19–27-year-olds and 9–18-year-olds in 1991 of 0.3 years and 0.15 years, respectively. By 2000, the then 28–37-year-olds and 18–27-year-olds both showed schooling gains of about 0.2 on average. A likely distribution of individual-level gains that would lead to this average impact is that 30 out of 100 individuals in the older cohort and 15 out of 100 in the younger cohort completed another year of schooling by 1991, and that by 2000, 20 out 100 in both cohorts had a one-year education advantage over comparison cohorts. Given the shares of these cohorts in the total population—23 percent and 27 percent, respectively, according to De Carvalho (1997)—we can thus estimate what percent of the overall population got an additional year of schooling, namely about $23\% \times 30\% + 27\% \times 15\% = 11\%$ in 1991 and about $23\% \times 20\% + 27\% \times 20\% = 10\%$ in 2000.

Now suppose that an extra year of schooling raises wages by 12 percent (Behrman and Birdsall 1983), that labor supply is constant, and that about 10 percent of the population earn per capita income that falls within a 12 percent range below the poverty line. Suppose further that in 1991 about 65 percent of the total population would have been poor in the absence of the extra funding (this corresponds to the comparison group average poverty rate shown in Table 4) and that schooling only increased among the poor, so that $0.11/0.65 = 17\%$ of the poor got an additional year of education. If the schooling gains are independent from the distance to the poverty line, then $10\% \times 17\% = 1.7\%$ of the total population escaped poverty through the schooling channel alone. This number will be higher the larger the (average) returns to schooling, the larger the share of the population within range to cross the poverty line given returns to schooling, and the higher the share of the poor within that range that do get an additional year of schooling (those closer to the poverty line might be more likely to get more schooling than those that are extremely poor). The education

Table 4
Impacts on the Poverty Rate in 1980, 1991, 2000, and 2010

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
Poverty headcount ratio in 1980	0.62								
$I[X>0]$		0.039		–0.005		–0.013		–0.010	–0.002
		(0.037)		(0.028)		(0.024)		(0.014)	(0.021)
R^2/observations		0.77		0.77		0.76		1190	802
Poverty headcount ratio in 1991	0.65								
$I[X>0]$		–0.037	–0.064	–0.060	–0.051	–0.054	–0.037	–0.018	–0.019
		(0.039)	(0.022)	(0.029)	(0.017)	(0.024)	(0.015)	(0.009)	(0.012)
R^2/observations		0.79	0.93	0.78	0.92	0.76	0.91	954	764
Poverty headcount ratio in 2000	0.53								
$I[X>0]$		–0.053	–0.064	–0.069	–0.051	–0.055	–0.028	–0.022	–0.028
		(0.032)	(0.019)	(0.025)	(0.015)	(0.022)	(0.014)	(0.010)	(0.011)
R^2/observations		0.81	0.94	0.80	0.93	0.78	0.92	718	583
Poverty headcount ratio in 2010	0.31								
$I[X>0]$		–0.023	–0.022	–0.026	–0.010	–0.028	–0.007	–0.005	–0.007
		(0.021)	(0.018)	(0.018)	(0.015)	(0.016)	(0.013)	(0.009)	(0.010)
R^2/observations		0.83	0.91	0.82	0.91	0.81	0.90	950	767
Observations		202	199	297	294	391	387		

Notes: Results for poverty in 1980 and 1991 from Litschig and Morrison (2013, Tables 3 and 8, respectively). Poverty headcount ratios for 1980, 1991, and 2000 were computed by IPEA based on census data. 2010 poverty rate computed by the authors based on the 2010 census. The poverty line in 1980 is about R$95, and in 1991, 2000, and 2010 it is R $75.5, all in August 2000 prices. All specifications pool across the first three cutoffs and include state fixed effects and segment dummies. Columns 1–6 are OLS local linear discontinuity estimates with robust standard errors in parentheses. OLS specifications allow for differential slopes by segment and on each side of the cutoff. Columns 7 and 8 give estimates and standard errors based on Imbens–Kalyanaraman (IK) and Calonico–Cattaneo–Titiunik (CCT) optimal bandwidths. Neighborhood (percent) is percent distance from respective cutoff. Pretreatment covariates (1980 census) include county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas.

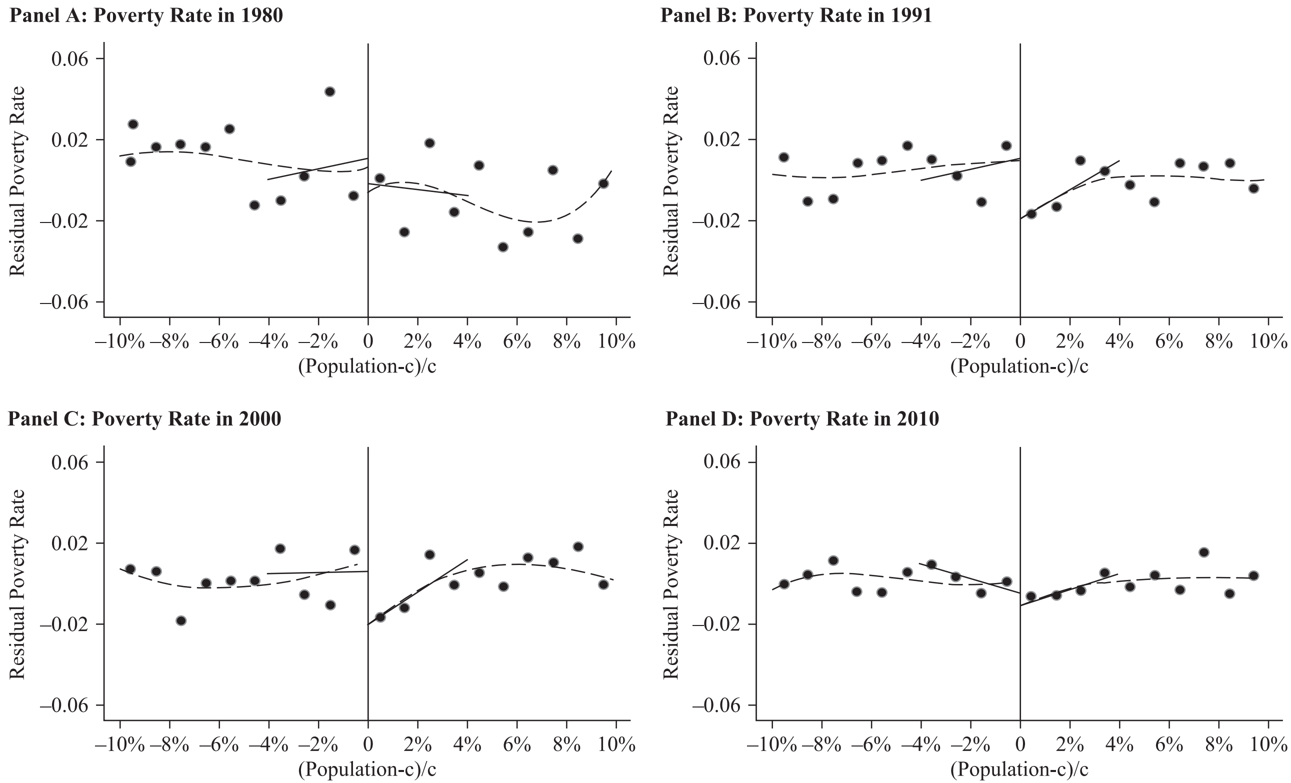

Figure 4

Impacts on the Poverty Rate in 1980, 1991, 2000, and 2010

Source: Results for poverty in 1980 and 1991 from Litschig and Morrison (2013, Table 3 and Figure 5, respectively).
Note: Each dot represents the sample average of residual poverty in a given bin. The residual comes from a regression of the poverty rate on state fixed effects and segment dummies. For 1991, 2000, and 2010 the regression also includes pretreatment covariates from the 1980 census: county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas. The bin width is one percentage point of the respective threshold, $c = 10{,}188$; $13{,}584$; $16{,}980$.

channel alone can thus account for about two percentage points of the estimated total three to five percentage points of poverty reduction in 1991, leaving the remaining one to three percentage points to better local public service provision overall. And since about 10 percent of the overall population in marginal treatment communities had a one-year education advantage by 2000, about two percentage points of the poverty reduction in 2000 might be associated with schooling, leaving another one to three percentage points to unmeasured and persistent public service improvements overall.

VIII. Impacts on Cognitive Development of Next Generation Cohorts

A. Impacts on ENEM Test Scores—2007–2011, High School Graduating Cohorts

Table 5 gives estimates of the jump in the municipality-level mean, 10th, 25th, 50th, 75th, and 90th percentiles of standardized ENEM test scores for high school graduating cohorts, pooled for 2007–2011. OLS estimates in Columns 1–6 of the first row indicate a gain in average test scores of about 0.08 standard deviation. Statistical significance is mostly at 5 or 10 percent. Estimates based on Imbens and Kalyanaraman and Calonico, Cattaneo, and Titiunik optimal bandwidth procedures shown in Columns 7 and 8 are smaller and not significant statistically. OLS results in rows two through five suggest that not just the mean but the entire ENEM test score distribution shifted to the right in municipalities immediately to the right of the population cutoffs. Impact estimates mostly fall in the range of 0.06 to 0.10 standard deviation. Statistical significance is typically at 1 or 5 percent for the tenth percentile and at 5 or 10 percent for the 25th, 50th, and 75th percentiles. Estimates at the top of the distribution (90th percentile) are typically not significant statistically. Estimates for ENEM percentiles based on Imben–Kalyanaraman and Calonico–Cattaneo–Titiunik optimal bandwidth procedures are again smaller and not significant statistically. Figure 5 provides graphical evidence of the discontinuities in the mean, median, 10th, and 90th percentile ENEM statistics at the cutoff. Figures for the 25th and 75th percentiles are similar and available upon request. Although the plots are quite noisy, there is strongly suggestive—even if not fully conclusive—evidence of an upward shift at the cutoff in all four panels. Online Appendix Figure 2 confirms the upward shift by plotting the marginal test score distributions separately for test-takers residing in municipalities with population within 2 percent above and below the first three FPM cutoffs. Overall, there is thus both statistical and graphical evidence of a rightward shift in the entire distribution of ENEM test scores at the cutoff.

B. Impacts on Prova Brasil Test Scores—2007, 2009, and 2011; Eighth- or Ninth-Graders

Table 6 presents estimates of discontinuities in the municipality-level mean, 10th, 25th, 50th, 75th, and 90th percentiles of standardized Prova Brasil test scores for students in eighth or ninth grade pooled across 2007, 2009, and 2011. Estimated gains for mean test scores fall mostly in the 0.05–0.10 standard deviation range and are typically significant at 5 or 10 percent. Results in Rows 2–5 suggest that the mean increase in Prova Brasil

Table 5
Impacts on the Distribution of ENEM Test Scores, 2007–2011, High School Graduating Cohorts

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
Average test score	–0.47								
$I[X>0]$		0.082	0.082	0.108	0.078	0.075	0.033	0.026	0.025
		(0.057)	(0.045)	(0.045)	(0.036)	(0.041)	(0.033)	(0.022)	(0.026)
R^2/observations		0.76	0.84	0.75	0.83	0.73	0.83	918	902
10th percentile	–1.44								
$I[X>0]$		0.094	0.102	0.089	0.073	0.062	0.038	0.017	0.018
		(0.042)	(0.034)	(0.033)	(0.027)	(0.029)	(0.025)	(0.018)	(0.017)
R^2/observations		0.73	0.81	0.73	0.80	0.69	0.78	803	861
25th percentile	–1.05								
$I[X>0]$		0.080	0.081	0.092	0.070	0.057	0.027	0.013	0.009
		(0.052)	(0.041)	(0.040)	(0.033)	(0.036)	(0.029)	(0.020)	(0.025)
R^2/observations		0.74	0.82	0.74	0.81	0.72	0.81	800	766
Median test score	–0.55								
$I[X>0]$		0.080	0.080	0.109	0.080	0.066	0.026	0.015	0.015
		(0.061)	(0.050)	(0.049)	(0.040)	(0.043)	(0.036)	(0.025)	(0.029)
R^2/observations		0.75	0.83	0.75	0.82	0.72	0.81	817	849

(*continued*)

Table 5 (*continued*)

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
75th percentile	0.04								
$I[X>0]$		0.077	0.079	0.121	0.085	0.089	0.036	0.032	0.032
		(0.068)	(0.056)	(0.055)	(0.045)	(0.050)	(0.042)	(0.025)	(0.028)
R^2/observations		0.76	0.83	0.74	0.82	0.72	0.82	1114	875
90th percentile	0.62								
$I[X>0]$		0.098	0.088	0.128	0.080	0.090	0.026	0.029	0.021
		(0.080)	(0.066)	(0.064)	(0.052)	(0.058)	(0.048)	(0.030)	(0.039)
R^2/observations		0.73	0.81	0.72	0.80	0.70	0.81	1083	834
Observations		202	199	297	294	391	387		

Notes: ENEM 2007–2011 test-taker samples, high school (12th grade) graduating cohorts. All specifications pool across the first three cutoffs and include state fixed effects and segment dummies. Columns 1–6 are OLS local linear discontinuity estimates with robust standard errors in parentheses. OLS specifications allow for differential slopes by segment and on each side of the cutoff. Columns 7 and 8 give estimates and standard errors based on Imbens–Kalyanaraman (IK) and Calonico–Cattaneo–Titiunik (CCT) optimal bandwidths. Neighborhood (percent) is percent distance from respective cutoff. Pretreatment covariates (1980 census) include county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas.

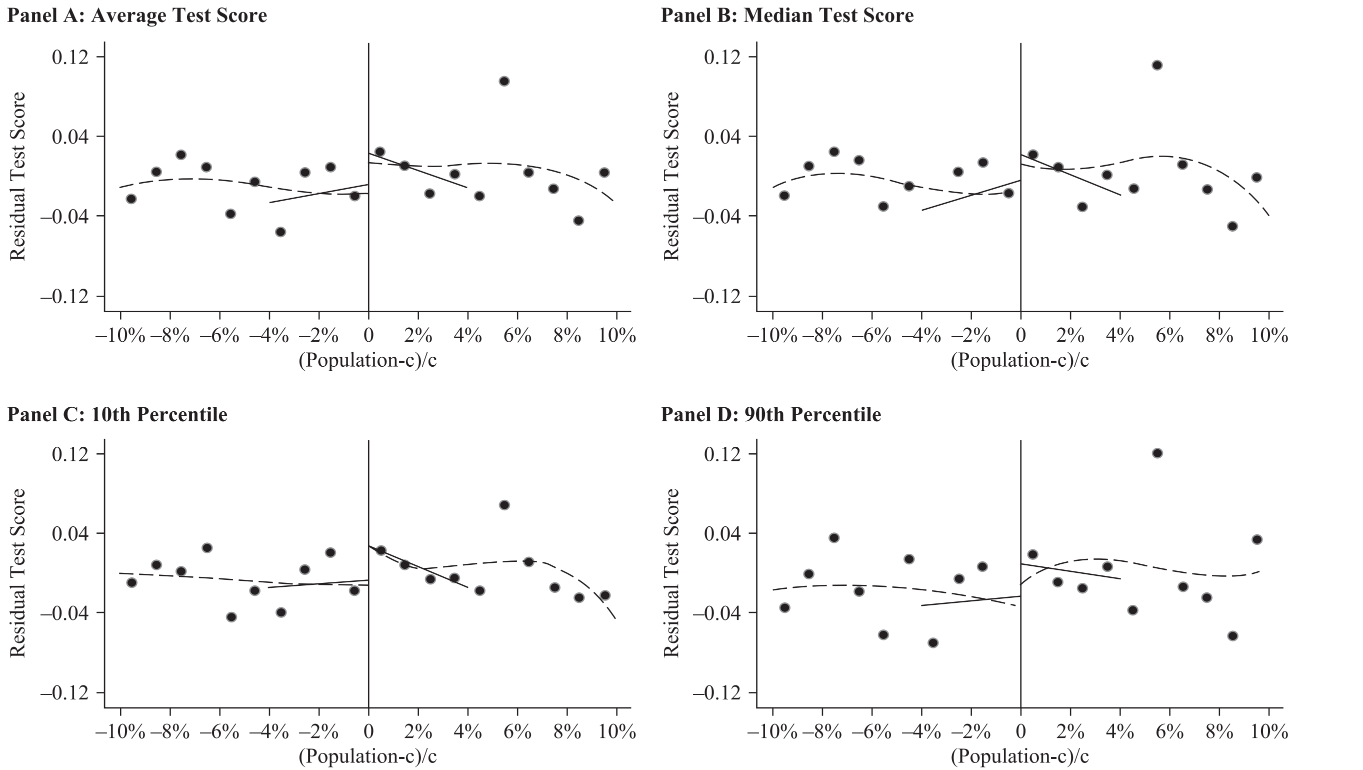

Figure 5

Impacts on the Distribution of ENEM Test Scores—2007–2011, High School Graduating Cohorts

Source: ENEM 2007–2011 test-taker samples, high school (12th grade) graduating cohorts.

Notes: Each dot represents the sample average of the residual mean or given percentile test score in a given bin. The residual comes from a regression of the mean or given percentile test score on state fixed effects, segment dummies, and pretreatment covariates from the 1980 census: county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas. The bin width is one percentage point of the respective threshold, $c = 10{,}188$; 13,584; 16,980.

Table 6
Impacts on the Distribution of Prova Brasil Test Scores—2007, 2009, and 2011; Eighth- or Ninth-Graders

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
Average test score	−0.14								
$I[X>0]$		0.032	0.048	0.105	0.089	0.113	0.078	0.052	0.054
		(0.068)	(0.063)	(0.056)	(0.050)	(0.048)	(0.043)	(0.026)	(0.043)
R^2/observations		0.72	0.78	0.72	0.77	0.70	0.76	1267	696
10th percentile	−1.28								
$I[X>0]$		0.039	0.065	0.103	0.090	0.112	0.084	0.060	0.070
		(0.074)	(0.068)	(0.059)	(0.054)	(0.050)	(0.045)	(0.028)	(0.044)
R^2/observations		0.61	0.68	0.61	0.67	0.60	0.67	1173	683
25th percentile	−0.78								
$I[X>0]$		0.060	0.086	0.136	0.125	0.140	0.110	0.058	0.076
		(0.074)	(0.069)	(0.060)	(0.055)	(0.052)	(0.047)	(0.027)	(0.046)
R^2/observations		0.67	0.73	0.68	0.74	0.66	0.72	1300	695
Median test score	−0.17								
$I[X>0]$		0.023	0.039	0.111	0.093	0.125	0.086	0.051	0.057
		(0.076)	(0.070)	(0.061)	(0.055)	(0.053)	(0.047)	(0.028)	(0.047)
R^2/observations		0.70	0.76	0.70	0.76	0.68	0.75	1286	719

(*continued*)

Table 6 (*continued*)

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
75th percentile	0.47								
$I[X>0]$		0.015	0.025	0.087	0.069	0.100	0.060	0.048	0.038
		(0.074)	(0.067)	(0.059)	(0.053)	(0.052)	(0.047)	(0.030)	(0.046)
R^2/observations		0.73	0.79	0.73	0.79	0.72	0.78	1166	696
90th percentile	1.06								
$I[X>0]$		0.011	0.013	0.092	0.070	0.090	0.049	0.045	0.034
		(0.070)	(0.067)	(0.057)	(0.053)	(0.050)	(0.047)	(0.029)	(0.046)
R^2/observations		0.76	0.80	0.75	0.79	0.74	0.79	1200	697
Observations		202	199	297	294	391	387		

Notes: PB 2007 eighth-grade, 2009 ninth-grade, and 2011 ninth-grade test-taker samples. All specifications pool across the first three cutoffs and include state fixed effects and segment dummies. Columns 1–6 are OLS local linear discontinuity estimates with robust standard errors in parentheses. OLS specifications allow for differential slopes by segment and on each side of the cutoff. Columns 7 and 8 give estimates and standard errors based on Imbens–Kalyanaraman (IK) and Calonico–Cattaneo–Titiunik (CCT) optimal bandwidths. Neighborhood (percent) is percent distance from respective cutoff. Pretreatment covariates (1980 census) include county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas.

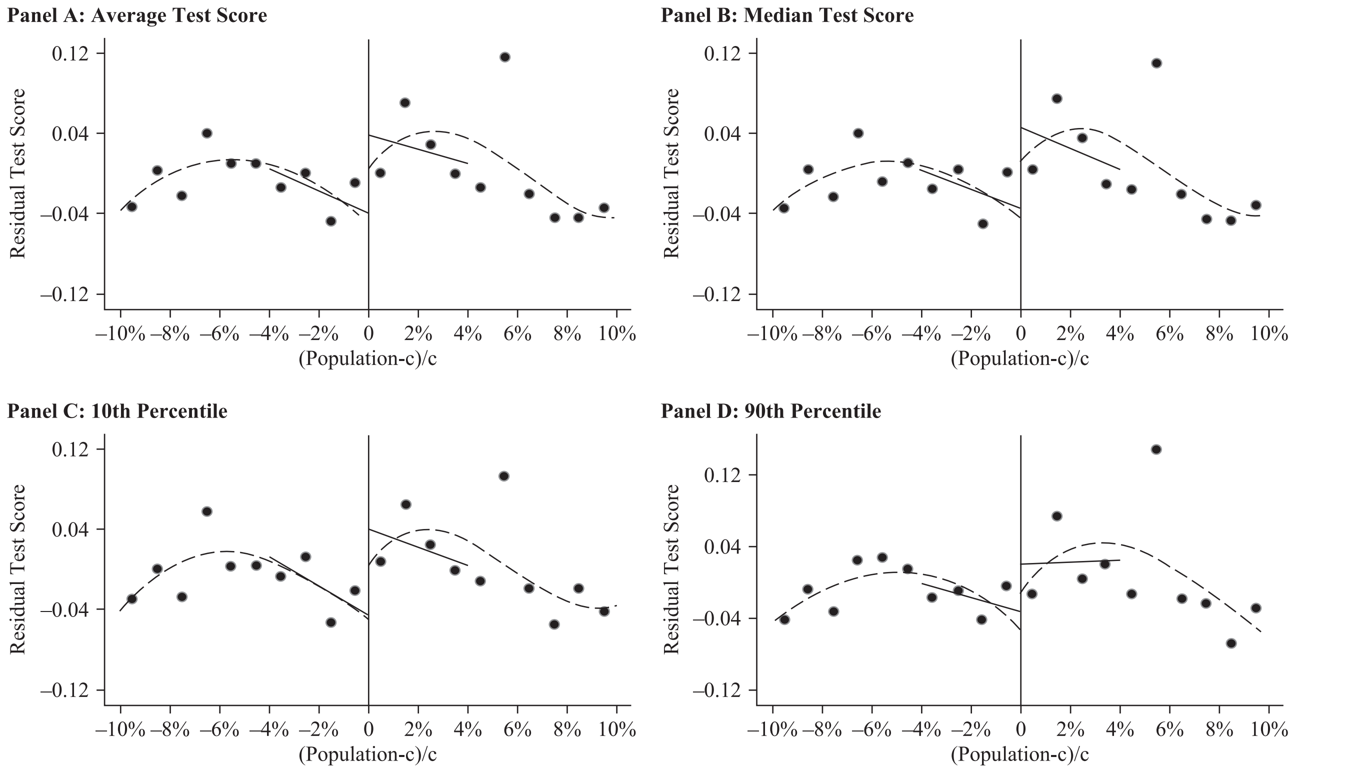

Figure 6

Impacts on the Distribution of Prova Brasil Test Scores—2007, 2009, and 2011; Eighth- or Ninth-Graders

Source: PB 2007 eighth-grade, 2009 ninth-grade, and 2011 ninth-grade test-taker samples.

Notes: Each dot represents the sample average of the residual mean or given percentile test score in a given bin. The residual comes from a regression of the mean or given percentile test score on state fixed effects, segment dummies, and pretreatment covariates from the 1980 census: county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas. The bin width is one percentage point of the respective threshold, $c = 10{,}188$; 13,584; 16,980.

test scores is more strongly driven by the bottom of the distribution than by the top. For the 10th and 50th percentiles, estimates range mostly from 0.06 to 0.10, and for the 25th percentile most estimates fall within 0.08 to 0.13 standard deviation. For the 10th, 25th, and 50th percentiles, statistical significance is typically at 5 or 10 percent. For the 75th and 90th percentiles, estimates mostly range from 0.04 to 0.09 and are typically not statistically different from zero. Figure 6 presents graphical evidence of the discontinuities in Prova Brasil mean, median, 10th, and 90th percentile test score statistics for eighth- or ninth-graders. It is clear in Panels A, B, and C that there is an upward shift of similar magnitude at the cutoff point for mean, median, and tenth percentile test scores, respectively, and that there is a somewhat smaller upward shift for 90th percentile test scores in Panel D. Figures for the 25th and 75th percentiles are similar and available upon request. Online Appendix Figure 3 again confirms the upward shift by plotting the marginal distributions of Prova Brasil test scores for students in eighth or ninth grade separately for test-takers residing in municipalities with population within 2 percent above and below the first three FPM cutoffs. Overall, there is thus clear statistical and graphical evidence of a rightward shift in the distribution of Prova Brasil test scores at the cutoff.

C. Impacts on Prova Brasil Test Scores—2007, 2009, and 2011; Fourth- or Fifth-Graders

Online Appendix Table 14 presents estimates of the jump in the municipality-level mean, 10th, 25th, 50th, 75th, and 90th percentiles of standardized Prova Brasil test scores for students in fourth or fifth grade, pooled across 2007, 2009, and 2011. Results are similar across estimation approaches, showing no evidence of gains anywhere in the test score distribution. The statistical evidence is also in line with the graphical evidence (available upon request). Online Appendix Figure 4 shows that the marginal distributions of Prova Brasil test scores for students in fourth or fifth grade in municipalities with population within 2 percent above and below the first three FPM cutoffs completely overlap. As noted earlier, the zero effect for fourth- or fifth-graders might be due to sample selection bias since in treatment communities a larger proportion of children enrolled in private primary schools and therefore did not take the Prova Brasil exam.

IX. Mechanisms

A. Impacts on Student-to-Teacher Ratio in Public Primary Schools, 1991–2011

Table 7 presents discontinuity estimates for the student-to-teacher ratio in public primary schools over the period 1991–2011. Results vary little across OLS and optimal bandwidth approaches and are significant at least at 10 percent in virtually all specifications. In 1991, the student-to-teacher ratio was about two to three students lower in marginal treatment communities, down from a mean of 20.8 just to the left of the cutoff. In 1995–1996, the reduction amounted to about three to four students per teacher, and during the 1997–2003 period, the reduction of the average student-to-teacher ratio was by about two students. An average class size reduction of two to three students per teacher persisted through 2007–2011. Figure 7 provides

Table 7
Impacts on Student-to-Teacher Ratio—Primary Public Schools, 1991–2011

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:	Comparison Mean:	N	Y	N	Y	N	Y	Y	Y
Primary school student-to-teacher ratio in 1991:	20.8								
$I[X>0]$		−3.605	−3.257	−2.728	−2.410	−2.405	−1.879	−1.669	−2.536
		(1.510)	(1.473)	(1.180)	(1.143)	(1.061)	(0.990)	(0.791)	(1.091)
R^2		0.45	0.51	0.46	0.53	0.46	0.53		
Observations		175	172	263	260	345	341	591	427
Primary school student-to-teacher ratio in 1995–1996:	21.3								
$I[X>0]$		−5.021	−4.666	−3.336	−3.318	−3.695	−3.526	−3.115	−3.830
		(1.576)	(1.499)	(1.241)	(1.208)	(1.119)	(1.080)	(0.906)	(1.244)
R^2		0.53	0.59	0.55	0.60	0.54	0.59		
Observations		177	174	263	260	347	343	578	428
Primary school student-to-teacher ratio in 1997–2003:	23.7								
$I[X>0]$		−2.889	−2.468	−2.150	−1.818	−2.187	−1.702	−1.516	−1.761
		(1.198)	(1.173)	(0.941)	(0.916)	(0.811)	(0.801)	(0.601)	(0.712)
R^2		0.60	0.64	0.62	0.65	0.62	0.66		
Observations		193	190	286	283	375	371	744	709

(*continued*)

Table 7 (*continued*)

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
Primary school student-to-teacher ratio in 2004–2006:	22.3								
$I[X>0]$		−3.987	−3.920	−2.922	−2.938	−2.642	−2.523	−1.587	−1.961
		(1.439)	(1.343)	(1.117)	(1.100)	(0.977)	(0.951)	(0.629)	(0.763)
R^2		0.55	0.63	0.53	0.58	0.53	0.57		
Observations		140	139	206	205	273	271	594	526
Primary school student-to-teacher ratio in 2007–2011:	17.1								
$I[X>0]$		−3.330	−3.202	−3.166	−2.928	−2.537	−2.086	−1.575	−2.461
		(1.285)	(1.255)	(1.075)	(1.064)	(0.927)	(0.913)	(0.682)	(0.919)
R^2		0.70	0.73	0.68	0.70	0.68	0.70		
Observations		199	196	294	291	385	381	839	638

Notes: Primary schools include only municipal public schools, not state-run or private primary schools. Authors' calculations of student-to-teacher ratios are based on the school census for each year. All specifications pool across the first three cutoffs and include state fixed effects and segment dummies. Columns 1–6 are OLS local linear discontinuity estimates with robust standard errors in parentheses. OLS specifications allow for differential slopes by segment and on each side of the cutoff. Columns 7 and 8 give estimates and standard errors based on Imbens–Kalyanaraman (IK) and Calonico–Cattaneo–Titiunik (CCT) optimal bandwidths. Neighborhood (percent) is percent distance from respective cutoff. Pretreatment covariates (1980 census) include county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas.

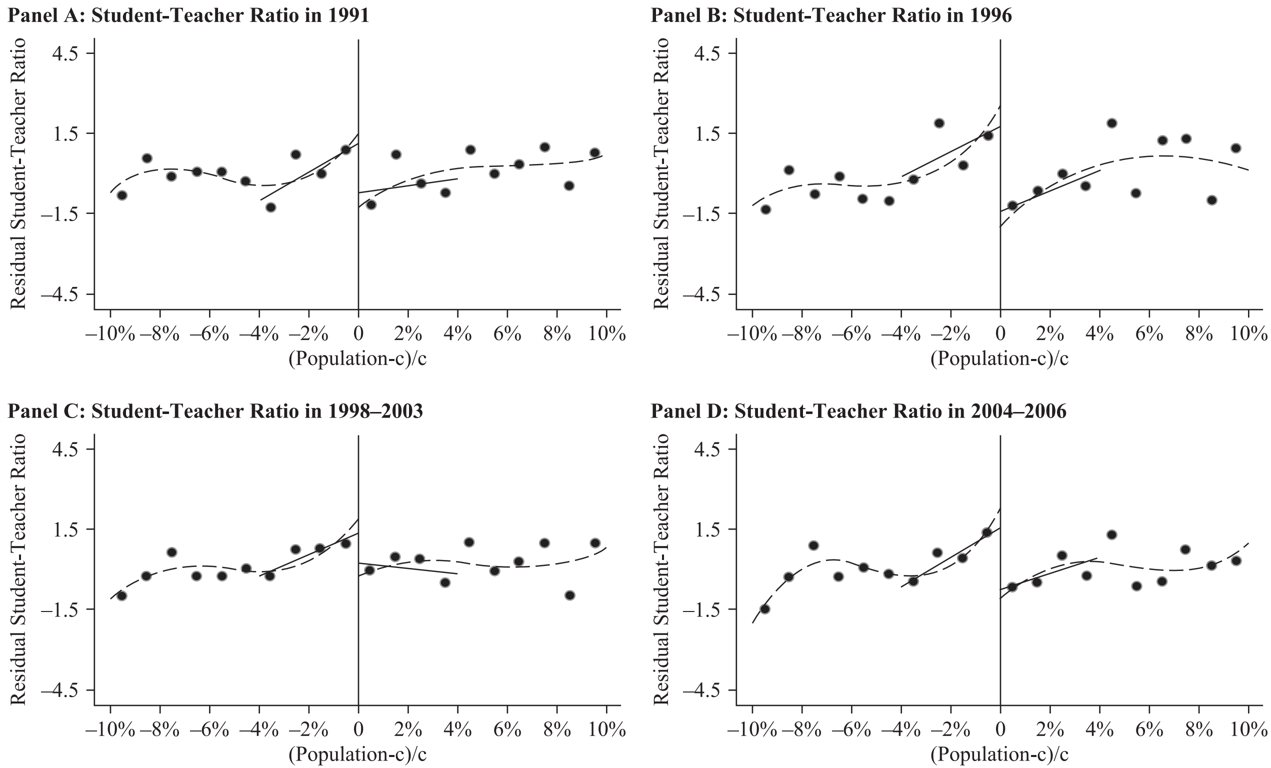

Figure 7

Impacts on Student-to-Teacher Ratio—Primary Public Schools, 1991–2006

Source: Authors' calculations of student-to-teacher ratios in municipal primary public schools are based on the school census for each year.

Notes: Each dot represents the sample average of the residual student-to-teacher ratio in a given bin. The residual comes from a regression of the student-to-teacher ratio in a given year or range of years on state fixed effects, segment dummies, and pretreatment covariates from the 1980 census: county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas. The bin width is one percentage point of the respective threshold, $c = 10{,}188$; 13,584; 16,980.

clear graphical evidence of these discontinuities. Exact temporal patterns should be interpreted with caution because sample sizes in 1991 and 2004–2006 in particular are 10–20 percent lower due to missing school census information for those years. Nonetheless, the available evidence clearly points to a reduction in local public primary school average class size of about two to three students per teacher that persisted throughout the 1990s and 2000s.

B. Impacts on Parental Education, ENEM and Prova Brasil, Late 2000s

Table 8 gives estimates of the jump in education levels of the (most-educated) parents of ENEM high school graduating cohorts, pooled from 2007 to 2011. OLS estimates suggest that the proportion of parents with no more than a primary school education decreased by about three to four percentage points, and that corresponding increases in parental education levels are about equally distributed among the proportions with some middle school (up to Grade 8) and some high school (up to Grade 12). The proportion of parents with college education is no different between treatment and comparison communities. Statistical significance reaches 5 percent in the 3 percent and 4 percent discontinuity samples and is weaker in the 2 percent sample. Calonico–Cattaneo–Titiunik results are in line with OLS results. Imbens–Kalyanaraman discontinuity estimates are small and indistinguishable from zero. Figure 8 shows clear evidence of a reduction in the proportion of parents with no more than primary education and corresponding increases in the proportions of parents with some middle or high school education.

In Table 9 we present estimates of discontinuities in education levels of the parents of Prova Brasil test-takers in eighth or ninth grade, pooled across 2007, 2009, and 2011. Results vary little across OLS and optimal bandwidth approaches and indicate that the proportion of parents with no more than a primary school education decreased by about three percentage points. The corresponding increase in parents' education is observed mostly in the proportion of parents who completed at least high school. Statistical significance is typically at 1 or 5 percent in the 3 and 4 percent and optimal bandwidth samples and is weaker in the 2 percent OLS sample. Figure 9 shows clear evidence of a reduction in the proportion of parents with no more than primary education and a corresponding increase in the proportion of parents who completed at least high school.

Online Appendix Table 15 shows estimates of discontinuities in education levels of the parents of Prova Brasil test-takers in fourth or fifth grade, pooled across 2007, 2009, and 2011. Results vary little across OLS and optimal bandwidth approaches and suggest that the proportion of parents with no more than a fourth grade education decreased by about two to three percentage points. The corresponding increase in parents' education is observed mostly in the proportion of parents who completed at least high school. Statistical significance for the discontinuity estimate of the lowest education category is mostly at 1 or 5 percent across OLS and optimal bandwidth samples and is weaker for other parental education categories. Graphical evidence (available upon request) shows clear evidence of a reduction in the proportion of parents with no more than primary education and a less striking increase in the proportion of parents who completed at least high school.

Table 8
Impacts on Parents' Education Levels—ENEM 2007–2011, High School Graduating Cohorts

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
No more than primary school	0.40								
$I[X>0]$		–0.025	–0.027	–0.045	–0.031	–0.041	–0.024	–0.002	–0.022
		(0.026)	(0.021)	(0.022)	(0.017)	(0.019)	(0.015)	(0.010)	(0.013)
R^2/observations		0.61	0.77	0.55	0.75	0.52	0.75	928	517
Some middle school	0.19								
$I[X>0]$		0.011	0.015	0.018	0.019	0.015	0.012	0.006	0.010
		(0.012)	(0.012)	(0.009)	(0.009)	(0.008)	(0.007)	(0.005)	(0.007)
R^2/observations		0.51	0.59	0.51	0.58	0.45	0.53	803	623
Some high school	0.26								
$I[X>0]$		0.013	0.015	0.023	0.014	0.024	0.015	0.003	0.012
		(0.016)	(0.014)	(0.013)	(0.011)	(0.013)	(0.011)	(0.006)	(0.009)
R^2/observations		0.42	0.60	0.37	0.59	0.33	0.56	1378	607
Some college	0.15								
$I[X>0]$		0.001	–0.002	0.004	–0.002	0.003	–0.003	–0.007	–0.007
		(0.011)	(0.011)	(0.010)	(0.009)	(0.008)	(0.008)	(0.005)	(0.007)
R^2/observations		0.61	0.70	0.54	0.62	0.53	0.62	1042	840
Observations		202	199	297	294	391	387		

Notes: ENEM 2007–2011 test-taker samples, high school (12th grade) graduating cohorts. The four categories correspond to the highest education level of the most educated parent. All specifications pool across the first three cutoffs and include state fixed effects and segment dummies. Columns 1–6 are OLS local linear discontinuity estimates with robust standard errors in parentheses. OLS specifications allow for differential slopes by segment and on each side of the cutoff. Columns 7 and 8 give estimates and standard errors based on Imbens–Kalyanaraman (IK) and Calonico–Cattaneo–Titiunik (CCT) optimal bandwidths. Neighborhood (percent) is percent distance from respective cutoff. Pretreatment covariates (1980 census) include county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas.

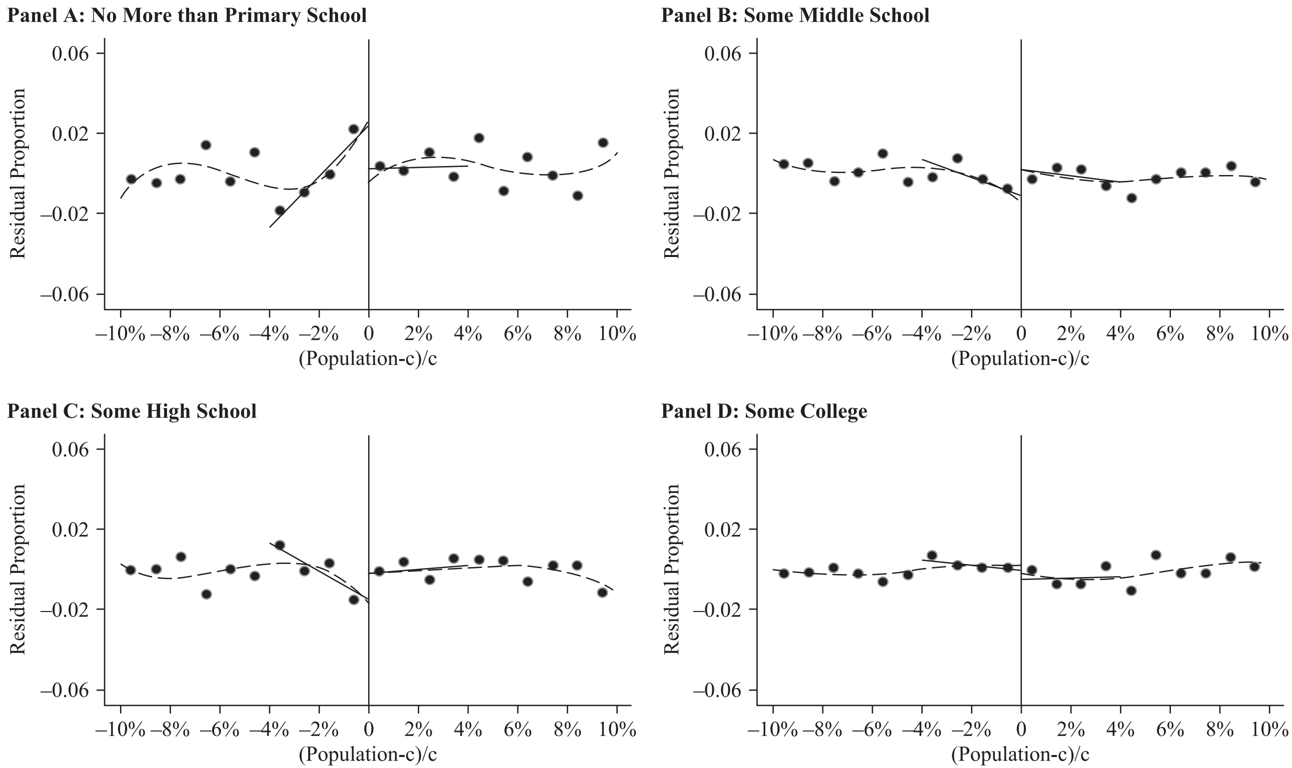

Figure 8

Impacts on Parents' Education Levels—ENEM 2007–2011, High School Graduating Cohorts

Source: ENEM 2007-2011 test-taker samples, high school (12th grade) graduating cohorts.

Notes: The four categories correspond to the highest education level of the most educated parent. Each dot represents the sample average of the residual proportion of test-takers in a given education category in a given bin. The residual comes from a regression of the proportion of test-takers in a given education category on state fixed effects, segment dummies, and pretreatment covariates from the 1980 census: county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas. The bin width is one percentage point of the respective threshold, $c = 10{,}188$; 13,584; 16,980.

Table 9

Impacts on Parents' Education Levels—Prova Brasil 2007, 2009, and 2011; Eighth- or Ninth-Graders

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
No more than primary school	0.30								
$I[X>0]$		–0.018	–0.022	–0.030	–0.024	–0.040	–0.030	–0.014	–0.023
		(0.016)	(0.012)	(0.012)	(0.009)	(0.011)	(0.008)	(0.007)	(0.009)
R^2/observations		0.72	0.85	0.73	0.84	0.70	0.83	674	473
Some middle school	0.28								
$I[X>0]$		–0.002	–0.003	–0.013	–0.008	–0.015	–0.011	–0.013	–0.018
		(0.015)	(0.015)	(0.011)	(0.011)	(0.010)	(0.010)	(0.008)	(0.010)
R^2/observations		0.45	0.51	0.39	0.49	0.33	0.46	626	458
Some high school	0.15								
$I[X>0]$		0.005	0.010	0.009	0.008	0.011	0.007	0.006	0.008
		(0.010)	(0.009)	(0.008)	(0.007)	(0.007)	(0.006)	(0.004)	(0.006)
R^2/observations		0.53	0.69	0.49	0.65	0.46	0.62	701	561
Completed at least high school	0.27								
$I[X>0]$		0.016	0.016	0.034	0.024	0.044	0.033	0.018	0.033
		(0.017)	(0.018)	(0.014)	(0.014)	(0.013)	(0.012)	(0.009)	(0.013)
R^2/observations		0.57	0.66	0.55	0.65	0.51	0.63	696	431
Observations		202	199	297	294	391	387		

Notes: PB 2007 eighth-grade, 2009 ninth-grade, and 2011 ninth-grade test-taker samples. The four categories correspond to the highest education level of the most educated parent. All specifications pool across the first three cutoffs and include state fixed effects and segment dummies. Columns 1–6 are OLS local linear discontinuity estimates with robust standard errors in parentheses. OLS specifications allow for differential slopes by segment and on each side of the cutoff. Columns 7 and 8 give estimates and standard errors based on Imbens-Kalyanaraman (IK) and Calonico, Cattaneo, and Titiunik (CCT) optimal bandwidths. Neighborhood (percent) is percent distance from respective cutoff. Pretreatment covariates (1980 census) include county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, illiterate percentage of people over 14 years old, infant mortality, enrollment of 7–14-year-olds and percent of population living in urban areas.

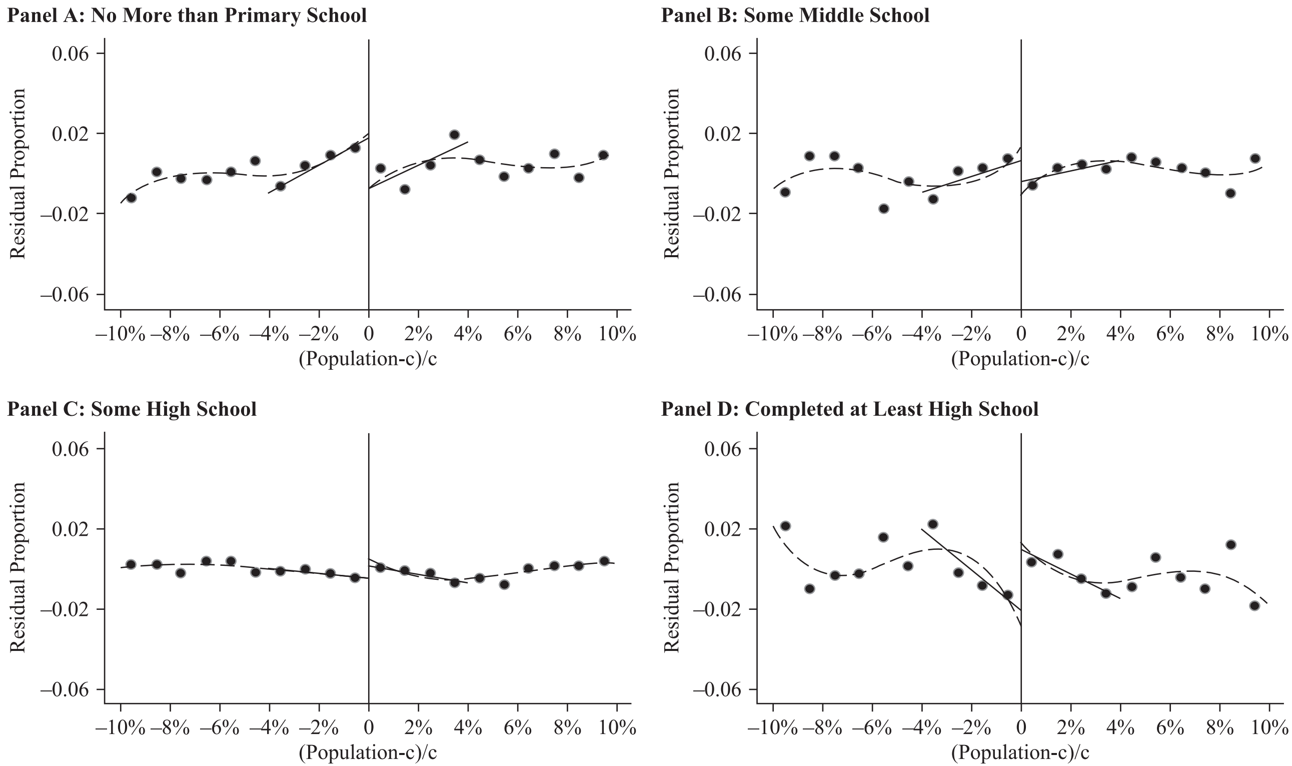

Figure 9

Impacts on Parents' Education Levels—Prova Brasil 2007, 2009, and 2011; Eighth- and Ninth-Graders

Source: PB 2007 eighth-grade, 2009 ninth-grade, and 2011 ninth-grade test-taker samples.

Notes: The four categories correspond to the highest education level of the most educated parent. Each dot represents the sample average of the residual proportion of test-takers in a given education category in a given bin. The residual comes from a regression of the proportion of test-takers in a given education category on state fixed effects, segment dummies, and pretreatment covariates from the 1980 census: county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds, and percent of population living in urban areas. The bin width is one percentage point of the respective threshold, $c = 10{,}188$; 13,584; 16,980.

Table 10
Impacts on Household Income—ENEM 2007–2009, High School Graduating Cohorts

Neighborhood (Percent):		2	2	3	3	4	4	IK	CCT
Pretreatment Covariates:		N	Y	N	Y	N	Y	Y	Y
	Comparison Mean:								
Income up to one minimum wage	0.37								
$I[X>0]$		–0.038	–0.046	–0.055	–0.040	–0.035	–0.018	–0.015	–0.022
		(0.032)	(0.021)	(0.025)	(0.017)	(0.021)	(0.015)	(0.012)	(0.015)
R^2/observations		0.80	0.90	0.78	0.88	0.77	0.87	793	694
Income between one and two minimum wages	0.36								
$I[X>0]$		0.005	–0.003	0.006	–0.001	–0.002	–0.006	0.006	0.011
		(0.020)	(0.020)	(0.016)	(0.016)	(0.014)	(0.014)	(0.011)	(0.012)
R^2/observations		0.29	0.36	0.28	0.32	0.27	0.30	668	771
Income between two and five minimum wages	0.21								
$I[X>0]$		0.027	0.041	0.032	0.029	0.024	0.017	0.004	0.012
		(0.024)	(0.018)	(0.018)	(0.014)	(0.015)	(0.012)	(0.008)	(0.012)
R^2/observations		0.75	0.86	0.73	0.84	0.70	0.82	1100	597
Income above five minimum wages	0.06								
$\mathrm{I}[\mathrm{X}>0]$		0.006	0.008	0.017	0.012	0.013	0.007	0.000	0.001
		(0.013)	(0.011)	(0.010)	(0.009)	(0.009)	(0.008)	(0.005)	(0.007)
R^2/observations		0.67	0.78	0.64	0.77	0.61	0.75	959	799
Observations		202	199	297	294	391	387		

Notes: ENEM 2007-2009 test-taker samples, high school graduating cohorts. Income refers to the test-taker's household. 1 minimum wage was R$380, R$415, and R$465 in 2007, 2008, 2009, respectively. All specifications pool across the first three cutoffs and include state fixed effects and segment dummies. Columns 1–6 are OLS local linear discontinuity estimates with robust standard errors in parentheses. OLS specifications allow for differential slopes by segment and on each side of the cutoff. Columns 7 and 8 give estimates and standard errors based on Imbens-Kalyanaraman (IK) and Calonico, Cattaneo, and Titiunik (CCT) optimal bandwidths. Neighborhood (percent) is percent distance from respective cutoff. Pretreatment covariates (1980 census) include county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, infant mortality, enrollment of 7–14-year-olds and percent of population living in urban areas.

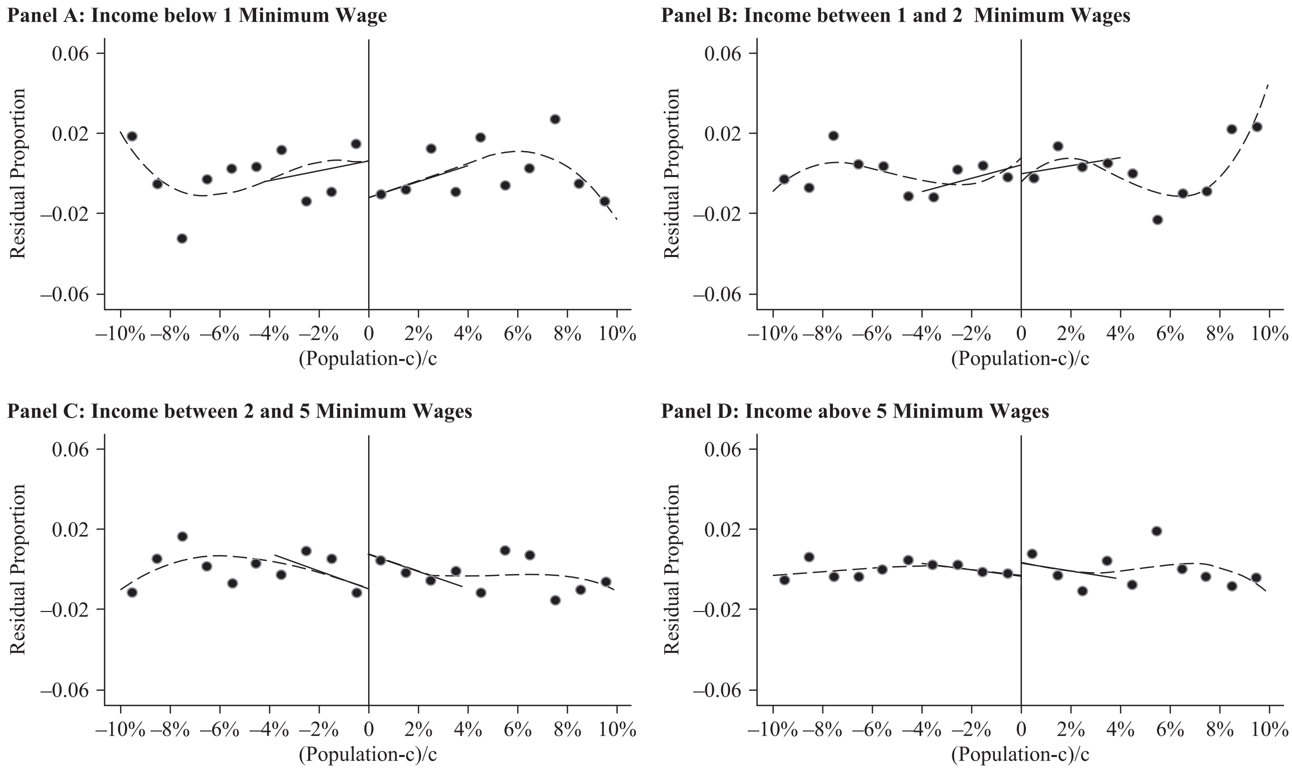

Figure 10

Impacts on Household Income—ENEM 2007–2009, High School Graduating Cohorts

Source: ENEM 2007–2009 test-taker samples, high school graduating cohorts.

Notes: Income refers to the test-taker's household. One minimum wage was R$380, R$415, and R$465 in 2007, 2008, 2009, respectively. Each dot represents the sample average of the residual proportion of test-takers in a given income category in a given bin. The residual comes from a regression of the proportion of test-takers in a given income category on state fixed effects, segment dummies, and pretreatment covariates from the 1980 census: county income per capita, average years of schooling for individuals 25 years and older, poverty headcount ratio, percentage of people over 14 years old who are illiterate, enrollment of 7–14-year-olds, and percent of population living in urban areas. The bin width is one percentage point of the respective threshold, $c = 10{,}188$; 13,584; 16,980.

C. Impacts on Household Income

In Table 10 we show impact estimates on the distribution of household income for high school graduating cohorts, pooled from 2007 to 2009. OLS estimates in the first row indicate that the proportion of households with monthly income up to one minimum wage is about four percentage points lower in marginal treatment communities, down from an average proportion of 0.37 in marginal comparison communities. OLS estimates in Rows 3 and 4 indicate that the proportions of ENEM test-taker households with income between two and five minimum wages and with more than five minimum wages are, respectively, three and one percentage points higher at the cutoff. Estimates of this shift are typically significant at 5 or 10 percent. Optimal bandwidth results are smaller in magnitude and indistinguishable from zero. Figure 10 provides evidence of a drop in the proportion of households with income below one minimum wage in communities just to the right of the cutoff and corresponding increases in the proportions of households with income between two and five and more than five minimum wages, respectively. Overall, there is thus both graphical and statistical evidence of a rightward shift in household income of ENEM graduating cohorts.

X. Discussion

As outlined in Section III, test score gains of the next generation $\mathrm{d}A_3/\mathrm{d}F_0$ can be decomposed as follows:

$$\frac{\mathrm{d}A_3}{\mathrm{d}F_0} = \frac{\partial A_3}{\partial C}\frac{\partial C}{\partial F_0} + \frac{\partial A_3}{\partial S^P}\frac{\partial S^P}{\partial F_0} + \frac{\partial A_3}{\partial I^P}\frac{\partial I^P}{\partial F_0} + \frac{\partial A_3}{\partial O}\frac{\partial O}{\partial F_0}.$$

Our estimates put the class size reduction $\partial C/\partial F_0$ at about −3, the impact of extra funding on the share of parents with no more than primary education $\partial S^P/\partial F_0$ at about −0.03, and the effect on the share of parents with income up to one minimum wage $\partial I^P/\partial F_0$ at about −0.04. We also estimate the partial correlations between mean ENEM 12th grade test scores at the end of the 2000s on the left-hand side and average class size from 1997 to 2003, the share of parents with no more than primary education, and the share of parents with income up to one minimum wage on the right-hand side using a cross-section of all Brazilian municipalities. The corresponding estimates and 95 percent confidence intervals are $\partial A_3/\partial C = -0.009[-0.008, -0.010]$, $\partial A_3/\partial S^P = -0.128[-0.198, -0.059]$, and $\partial A_3/\partial I^P = -1.219[-1.266, -1.172]$. Using the decomposition above, we have:

$$\frac{\mathrm{d}A_3}{\mathrm{d}F_0} = (-0.009)\times(-3) + (-0.128)\times(-0.03) + (-1.219)\times(-0.04) = 0.079$$

This exercise suggests that the above mechanisms together plausibly account for the approximately 0.06–0.10 impact on cognitive skills we estimate.

Since the parameter estimates linking class size and parental education and income to children's test scores based on Brazilian data are purely correlational, we also perform the same accounting exercise using estimates of $\partial A_3/\partial C$ and $\partial A_3/\partial S^P$ from prior studies. The drawback of these better identified estimates is that they are only available from developed country settings. In their followup study on the STAR experiment in Tennessee, Krueger and Whitmore (2001) show that smaller classes in Grades K–3 are associated with small but persistent test score gains through Grade 8, as well as a 0.13 standard deviation increase in college entrance exam scores, once they account for sample selection. Since the STAR experiment reduced class size by about seven students per teacher, we might expect an effect size of about $3/7 \times 0.13 \approx 0.056$ with the class size reduction of about three that we find in our data. Another related study by Fredriksson, Öckert, and Oosterbeek (2013) documents beneficial long-run impacts of smaller class size in late primary school (ages 10–13) on cognitive skills, completed education, and wages and earnings using Swedish data. Their estimates for academic achievement at age 16 suggest that a class size reduction of three students per teacher would increase test scores by 0.069 standard deviation three years after exposure. Although gains in cognitive skills from reduced class size in early grades may be muted in Brazil due to more frequent teacher absenteeism and less adequate teacher qualifications, the available evidence from developed countries suggests that the test score gains reported here are quantitatively plausible.

To what extent parental education might account for their children's test score gains is again difficult to tell because little is known about the causal effect of parental schooling on children's test score performance. Causal studies of intergenerational effects that exploit compulsory schooling reforms, such as Oreopoulos, Page, and Stevens (2006), focus on grade repetition in the United States, while Black, Devereux, and Salvanes (2005) and Holmlund, Lindahl, and Plug (2011) look at completed schooling in Norway and Sweden, respectively. The one study we are aware of (Carneiro, Costas, and Parey 2013) finds that an additional year of maternal schooling increases math and reading test scores of U.S. children ages 7–8 and 12–14 years by about 0.1 standard deviation. Assuming that about one-fifth of the children in our sample had a parent with an additional year of education (consistent with our results for school-age cohorts in the early 1980s), we would expect a $0.2 \times 0.1 = 0.02$ standard deviation increase when we look at all test-takers. Together with the class size effect of about 0.056–0.069 standard deviation discussed above, the class size and intergenerational channels add up to about 0.076–0.089 standard deviation, again close to the approximately 0.06–0.10 impact on cognitive skills we estimate in this study. We conduct a final test of the relevance of the class size and parental income and education channels by including these intermediary outcomes as controls in the outcome regression for test scores. Holding these intermediary outcomes constant should substantively attenuate the effect estimate on test scores if these are indeed the key drivers of the test score gains. Online Appendix Table 16 shows the results for ENEM test scores. Comparing columns without and with intermediary outcomes in a given neighborhood, impact estimates for average test scores, 25th percentiles, and median test scores fall by about half and lose statistical significance. Impact estimates at the tenth percentile are reduced by less than half, while impact estimates at the 75th and 90th percentiles are reduced by more than half. For Prova Brasil the attenuation is even more dramatic and happens across the entire test score distribution as shown in Online Appendix Table 17. Overall, the available evidence on mechanisms indicates that the cognitive gains of next-generation students are plausibly accounted for by reduced class size in primary school, intergenerational knowledge spillovers, and household income gains.

XI. Alternative Mechanisms

A. Age at First Birth and Fertility

Increased schooling as a result of the transfer windfall might have delayed age at first birth and reduced the number and age of children taking the tests. If so, the higher test score performance in beneficiary communities might be driven at least in part by a differential sibling and age composition of test-takers. The first row of Online Appendix Table 18 shows impacts on age at first birth for the younger of the directly affected cohorts (ages 0–9 in 1982, 28–37 in 2010). Point estimates are occasionally significant and positive but small, about 0.2 to 0.3, compared to a mean age at first birth of 22. The second row shows impacts on the average number of children of test-taking age (10–18-year-olds in 2010) by this cohort of parents. Point estimates are invariably small and never significant statistically. In the third row we show impacts on the average age of these children. Point estimates tend to be negative, but they are again small and never significant. Rows 4–6 of Online Appendix Table 18 show similar results for the older directly exposed cohort (10–19 in 1982, 38–47 in 2010) and their offspring in 2010. Overall, these results suggest that age at first birth and fertility were unaffected by the temporary funding boost.

B. Primary and Middle School Enrollment in 2000

As noted above, we find that in treatment communities a two to three percentage point larger proportion of 7–10-year-old children enrolled in private primary schools in 2010 and therefore did not take the Prova Brasil exam. This raises the possibility that the test score gains for eighth- and ninth-graders and high school graduating cohorts could be due to them attending private schools when they were younger. Online Appendix Table 19 presents impacts on primary and middle school enrollment in 2000 to investigate this potential channel. The first two rows show that average overall net enrollment rates in primary and middle school are about 94 percent and 92 percent, respectively, and that there is no differential enrollment at the cutoff. Rows three and four show that average net enrollment rates in public primary and middle school are about 90 percent and 89 percent, respectively, and that there is again no differential enrollment at the cutoff for public schools. Together, these results suggest that the test score gains we find for eighth- and ninth-graders and high school graduating cohorts are unlikely to be driven by private primary or middle school attendance when these cohorts were younger.

C. Political Competition

Litschig and Morrison (2013) provide suggestive evidence that the reelection probability of local incumbent parties in the 1988 elections improved by about ten percentage points at the cutoff. If the corresponding decrease in political competition continued in subsequent elections, it is possible that the quality of public spending suffered as a result and that the test score gains were in fact lower than they could have been with constant electoral competition. To investigate this possibility, we test whether the mayor election win margin (winner vote share – runner-up vote share) in subsequent elections was affected by the initial funding boost. The first row of Online Appendix Table 20 shows that in the 1996 mayoral elections, the average win margin was 17 percentage points to the left of the cutoffs. Most discontinuity estimates show a reduced win margin of a few percentage points, but only one out of eight is statistically different from zero. Similar results appear in the 2000 mayoral elections, shown in Row 2. In the 2004 and 2008 elections, the discontinuity estimates are closer to zero, and again only one is statistically different from zero as shown in Rows 3 and 4. Overall, this evidence suggests that the funding windfall had no persistent effects on local political competition.

XII. Conclusion

This work shows that a temporary increase in transfers to local governments in Brazil led to long-lasting schooling and literacy gains of school-age cohorts, as well as persistent poverty reduction in the community overall. Extra transfers also led to gains in cognitive skills of children and youth born after the extra transfers had expired. Available evidence on mechanisms indicates that these cognitive gains are plausibly accounted for by reduced class size in primary school and intergenerational spillovers. An important advantage of our study in terms of external validity is that the additional funds were distributed through and used by the regular Brazilian bureaucracy under routine conditions. We also need not worry about experimenter effects since our study population was not surveyed for a particular purpose, nor did stakeholders have any incentive to make the intervention look effective, for example. As with any regression discontinuity analysis, the impacts presented here apply only to relatively small municipalities with population levels at the respective cutoffs. Keeping track of government finances may be harder in larger cities and for other sources of revenues, such as oil royalties, which might weaken accountability and attenuate the schooling, income, and learning gains documented here. Given how much of decentralized public service delivery is financed by intergovernmental transfers, it is important that additional work in other settings assess the external validity of the findings reported here.

References

Behrman, Jere Richard, and Nancy Birdsall. 1983. "The Quality of Schooling: Quantity Alone is Misleading." *American Economic Review* 73(5):928–46.

Behrman, Jere Richard, Nancy Birdsall, and Robert Kaplan. 1996. "The Quality of Schooling and Labor Market Outcomes." In *Opportunity Foregone: Education in Brazil*, ed. Nancy Birdsall and Richard Sabot, 245–67. Baltimore, MD: Johns Hopkins University Press.

Birdsall, Nancy. 1985. "Public Inputs and Child Schooling in Brazil." *Journal of Development Economics* 18(1):67–86.

Black, Sandra Eilene, Paul Devereux, and Kjell Gunnar Salvanes. 2005. "Why the Apple Doesn't Fall Far: Understanding Intergenerational Transmission of Human Capital." *American Economic Review* 95(1):437–49.

Blimpo, Moussa, and David Evans. 2011. "School-Based Management and Educational Outcomes: Lessons from a Randomized Field Experiment." Unpublished.

Boadway, Robin, and Anwar Shah. 2009. *Fiscal Federalism: Principles and Practices of Multi-order Governance*. Cambridge, UK: Cambridge University Press.

Brollo, Fernanda, Tommaso Nannicini, Roberto Perotti, and Guido Tabellini. 2013. "The Political Resource Curse." *American Economic Review* 103(5):1759–96.

Carneiro, Pedro, Meghir Costas, and Mathias Parey. 2013. "Maternal Education, Home Environments, and the Development of Children and Adolescents." *Journal of the European Economic Association* 11(S1):123–60.

Case, Anne, and Angus Deaton. 1999. "School Inputs and Educational Outcomes in South Africa." *Quarterly Journal of Economics* 114(3):1047–84.

Caselli, Francesco, and Guy Michaels. 2013. "Do Oil Windfalls Improve Living Standards? Evidence from Brazil." *American Economic Journal: Applied Economics* 5(1):208–38.

Calonico, Sebastian, Matias Damian Cattaneo, and Rocio Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 86(2):2295–326.

De Carvalho, Magno. 1997. "Demographic Dynamics in Brazil: Recent Trends and Perspectives." *Brazilian Journal of Population Studies* 1:5–23.

Duflo, Esther. 2001. "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment." *American Economic Review* 91(4): 795–813.

Duflo, Esther, Pascaline Dupas, and Michael Kremer. 2015. "School Governance, Teacher Incentives, and Pupil–Teacher Ratios: Experimental Evidence from Kenyan Primary Schools." *Journal of Public Economics* 123:92–110.

Fredriksson, Peter, Björn Öckert, and Hessel Oosterbeek. 2013. "Long-Term Effects of Class Size." *Quarterly Journal of Economics* 128(1):249–85.

Gadenne, Lucie. 2017. "Tax Me, but Spend Wisely? Sources of Public Finance and Government Accountability." *American Economic Journal: Applied Economics* 9(1):274–314.

Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw. 2001. "Identification and Estimation of Treatment Effects with a Regression Discontinuity Design." *Econometrica* 69:201–9.

Holmlund, Helena, Mikael Lindahl, and Erik Plug. 2011. "The Causal Effect of Parent's Schooling on Children's Schooling: A Comparison of Estimation Methods." *Journal of Economic Literature* 49(3):614–50.

Imbens, Guido, and Karthik Kalyanaraman. 2012. "Optimal Bandwidth Choice for the Regression Discontinuity Estimator." *Review of Economic Studies* 79(3):933–59.

Instituto Brasileiro de Geografia e Estatística. 2002. "Estimativas Populacionais do Brasil, Grandes Regiões, Unidades da Federação e Municípios." IBGE background paper. Rio de Janeiro: IBGE.

Krueger, Alan, and Diane Whitmore. 2001. "The Effect of Attending a Small Class in the Early Grades on College-Test Taking and Middle School Test Results: Evidence from Project Star." *Economic Journal* 468:1–28.

Lee, David, and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature* 48(2):281–355.

Litschig, Stephan, and Kevin Morrison. 2013. "The Impact of Intergovernmental Transfers on Education Outcomes and Poverty Reduction." *American Economic Journal: Applied Economics* 5(4):206–40.

Monteiro, Joana, and Claudio Ferraz. 2010. "Does Oil Make Leaders Unaccountable? Evidence from Brazil's Offshore Oil Boom." Unpublished.

Muralidharan, Karthik, and Venkatesh Sundararaman. 2011. "Teacher Performance Pay: Experimental Evidence from India." *Journal of Political Economy* 119(1):39–77.

Olken, Benjamin, Junko Onishi, and Susan Wong. 2014. "Should Aid Reward Performance? Evidence from a Field Experiment on Health and Education in Indonesia." *American Economic Journal: Applied Economics* 6(4):1–34.

Olsson, Ola, and Michele Valsecchi. 2015. "Resource Windfalls and Local Government Behavior: Evidence from a Policy Reform in Indonesia." Unpublished.

Oreopoulos, Philip, Marianne Page, and Ann Huff Stevens. 2006. "The Intergenerational Effects of Compulsory Schooling." *Journal of Labor Economics* 24(4):358–68.
Pradhan, Menno, Daniel Suryadarma, Amanda Beatty, Maisy Wong, Arya Gaduh, Armida Alisjahbana, and Rima Prama Artha. 2014. "Improving Educational Quality through Enhancing Community Participation: Results from a Randomized Field Experiment in Indonesia." *American Economic Journal: Applied Economics* 6(2):105–26.
Rosenzweig, Mark, and Kenneth Wolpin. 1988. "Migration Selectivity and the Effects of Public Programs." *Journal of Public Economics* 37(3):265–89.
Shah, Anwar. 1991. "The New Fiscal Federalism in Brazil." World Bank Discussion Paper 124. Washington, DC: World Bank.
World Bank. 1985. *Brazil: Finance of Primary Education*. Washington, DC: World Bank.