The Effects of Nationwide Tuition Fee Elimination on Education Outcomes*

Robert Garlick†

September 17, 2016

Abstract

I study the effects on education outcomes of nationwide primary and secondary school fee eliminations in South Africa. This policy shifted education financing from a mixed user fee and government transfer system to a pure transfer system. This mirrors policy debates in developed and developing countries about the optimal mix of school fees, (subsidized) loans, and transfers to finance primary, secondary, and tertiary education. I find that fee elimination has very small effects on enrollment, grade progression, per-student school resources, and the socio-economic profile of the enrolled students. My results are robust to accounting for school-level selection into fee elimination, differential time trends between fee-eliminating and fee-charging schools, and student transfers between fee-charging and fee-eliminating schools. The enrollment effects imply a price elasticity between -0.03 and 0. This price insensitive demand does not appear to be driven by ceiling effects on enrollment, capacity constraints in schools, or measurement error in administrative data. I argue that the pattern of results may reflect low valuation of additional years of education by youths living near fee-eliminating schools, potentially due to a weak relationship between enrollment and subsequent attainment. My results show that demand-side subsidies may have limited ability to increase education participation when quality is low. They also highlight that fee elimination, and other demand-side subsidies, can be very expensive relative to the effects on education enrollment when demand is this price inelastic.

JEL codes: I25, O15

*I am grateful for helpful suggestions from Manuela Angelucci, John Bound, Jishnu Das, John DiNardo, Erica Field, Deon Filmer, Brian Jacob, David Lam, Jeffrey Smith, Stephen Taylor, Duncan Thomas, and conference and seminar participants at Bristol CPMO, CSAE, Duke, ESSA, IFPRI, Michigan, MIEDC, and NEUDC. Rirhandzu Baloyi, Justice Libago, Christo Lombaard, Erna Lubbe, Hersheela Narsee, Ralph Mehl, Siza Shongwe, and Hylton Visagie from South Africa's Department of Basic Education provided invaluable assistance in obtaining the data used in this project. All errors are my own.

†Duke University, robert.garlick@duke.edu.

1 Introduction

Almost all governments around the world provide some form of public education and must choose how to finance this public education. Different countries and regions employ different mixes of user fees, typically called "school fees" or "tuition fees," and transfers from the national or regional budget. Many countries have shifted this funding mix through time: school fees for tertiary education in Germany, the UK, and US have changed substantially in recent years; tax incentives for education spending in the US have changed over time; and school fees for primary and secondary education have changed substantially in Ghana, Kenya, Nigeria, and South Africa. Different financing mixes may lead to different levels of enrollment in public education, learning, and hence labor force composition. The optimal mix of school fees and transfers is theoretically ambiguous. School fees may promote accountability of public education providers to students and their families, and may screen out students with low returns to enrollment. However, school fees may reduce enrollment for credit-constrained or myopic students with high returns to enrollment. Transfers from government budgets may have distortionary effects on labor supply through the tax system. Given this theoretical ambiguity, empirical evidence on the effects of different education finance mixes is particularly important.

I study the effects of a nationwide shift in public education financing in South Africa. Primary and secondary public schools were historically funded by a mix of mandatory school fees and government transfers.[1] Between 2007 and 2010, school fees were eliminated in approximately 75% of public primary and secondary schools. Fees were first eliminated in neighborhoods with poverty rates above a threshold value, generating both time-series and cross-sectional variation in fee-charging status. I use panel data and regression discontinuity methods with administrative and household survey data to study the effects of fee elimination. I find small increases in enrollment ($\approx$ 1% of baseline enrollment), concentrated in the first two grades of secondary school, and no measureable effects on school resources (class sizes, socio-economic status of enrolled students) or learning outcomes (grade progression). I find no evidence of transfers between fee-eliminating schools and nearby fee-charging schools.

These enrollment effects imply that the price elasticity of demand for enrollment in public schools is at most -0.03. I show that the inelastic demand is not explained by ceiling effects on enrollment, capacity constraints in schools, noncompliance with the fee elimination policy, or measurement error in administrative data on enrollment. Instead, I hypothesize that demand is price inelastic in this setting because students in fee-eliminating schools derive little value from additional years of

[1] School fees for primary and secondary public schools are common in many developing countries, particularly in Africa. Primary school fees have been eliminated in many countries over the past two decades, and several countries are actively considering eliminating secondary school fees at the time of writing.

enrollment. I cannot directly test this hypothesis, but it is consistent with treatment effects being concentrated early in secondary school and with other research showing that grade progression is low and uncertain in South African secondary schools in high-poverty neighborhoods (Lam, Ardington, and Leibbrandt, 2010; van den Berg and Louw, 2007). Low valuation of enrollment for marginal students is perhaps surprising, given that Mincer returns to additional years of secondary education in South Africa exceed 10% (Branson, Ardington, Lam, and Leibbrandt, 2013).[2] This contrasts with findings from multiple countries that returns to enrollment in school and university are relatively high for financially constrained individuals, and not lower than Mincer returns (Card, 1999; Duflo, 2001; Heckman, Lochner, and Todd, 2006). I cannot directly determine why my results break this pattern. Results in other countries may be driven by marginal students attending reasonably high quality schools, while my results are driven by marginal students attending low quality schools.

This paper relates and contributes to four literatures.[3] The first literature examines the effect of demand-side incentives to enroll in (mostly primary) school in developing countries. These incentives include school fee elimination (Deininger, 2003; Fafchamps and Minten, 2007; Lucas and Mbiti, 2012), reductions in non-fee costs of enrollment (Evans, Kremer, and Ngatia, 2009; Hidalgo, Onofa, Oosterbeek, and Ponce, 2013), and conditional cash transfers (Baird, McIntosh, and Ozler, 2011; Filmer and Schady, 2008; Schultz, 2004).[4] The literature on conditional cash transfers is particularly large and includes many well-identified studies.[5] However, conditional cash transfers are a pure demand-side subsidy, while school fee elimination combines a supply-side subsidy with a price ceiling, so the two types of policies may have substantially different effects. My paper is most closely related to studies by Barrera-Osorio, Linden, and Urquiola (2007) in Colombia and Borkum (2012) in South Africa, which study the enrollment effects of school fee elimination using natural experiments. I extend this literature by considering a wider range of outcomes, interpreting treatment effects in terms of an economic model of human capital investment, and exploring the relationship between treatment effects and labor market returns to education.[6] My results suggest that large demand-side subsidies can be ineffective in increasing enrollment when education quality is low, reinforcing recent concerns that about poor learning outcomes in developing country schools (Pritchett, 2013). My results also show that school fee elimination can be very expensive relative

[2]Branson and Leibbrandt (2013) find that these large and convex returns to education persist even in Mincer regressions that control for proxies of school quality in the area where respondents grew up.

[3]There is another related literature that examines the effect of user fees on investments in health, including use of public clinics. This literature considers a range of user fees from insurance co-pays to direct clinic fees. Some papers in this literature explicitly examine the role of user fees in screening low- versus high-valuation considers (Dupas, 2014).

[4]See Kremer and Holla (2009) for a review of this literature. Most studies of fee elimination find relatively large effects on enrollment, though several rely on before-after comparisons that face identification challenges.

[5]See Fizbein and Schady (2009) for a review of this literature.

[6]Relative to Borkum's study of the South African experience, I also use a dataset with wider geographic and longer time coverage, and incorporate multiple sources of household survey data.

to the number of marginal students, consistent with a result for conditional cash transfers in Todd and Wolpin (2006).

The second literature examines the effect of price variation, including demand-side subsidies, on enrollment in and progression through public and private education institutions in developed countries. Some papers consider the effect of price variation on the choice between free public and fee-charging private primary or secondary schools (Dynarski, Gruber, and Li, 2009). More papers consider the effect of price variation on the choice of whether, where, and for how to enroll in either private or public post-secondary education institutions (Dynarski, 2003; Kane, 1994).[7] Unlike these studies, I focus on the margin of public primary/secondary school enrollment versus no enrollment. Like some of these studies, I consider the role of price discrimination in driving decisions about where to enroll in school. Fee elimination in South Africa was geographically targeted and I show that this targeting was relatively effective, in the sense that few students transferred from fee-charging to fee-eliminating schools. This contrasts with results from the US showing students are sensitive to geographic variation in the price of university enrollment (Knight and Schiff, 2016).

The third literature examines the effect of education financing on education outcomes. There is substantial debate about whether the *level* of education financing matters for education outcomes (Hanushek, 2003; Hoxby, 2001; Hyman, 2016; Jackson, Johnson, and Persico, 2016). I instead study the effect of the *source* of education financing on education outcomes, as the loss in school fee revenue was meant to be offset by larger transfers from the national government budget. There are few other papers that directly examine the effects of changing the source of education financing, holding the level of education financing constant. I do not observe school-level transfers and so cannot verify whether these offsetting transfers were received. My analysis instead has an intention-to-treat flavor, examining the effect of fee elimination and intended transfer increases on some measures of school resources and learning outcomes. I find no measureable effects on these outcomes, showing that substantial shifts in the source of education spending need change education outcomes.

The fourth literature examines link between political economy and education financing. School fees may promote local accountability of schools to students and their families, as they face a direct loss of revenue if students disenroll. Andrabi, Das, and Khwaja (2014) find that private primary schools' price and quality are responsive to competition induced by improved information about local education markets. I am aware of no comparable results for fee-charging public schools. My results do not directly speak to this mechanism, but the absence of transfers between fee-charging and fee-eliminating schools suggests that threats of school exit by students and their families may not be very credible to schools. Harding and Stasavage (2014) argue that school fees

[7] This literature is too large to discuss in detail here. See Hoxby (2004) for a review and discussion of policy implications.

are particularly salient and broad-based relative to other sources of education financing, and so are particularly politically unpopular. South Africa's school fee elimination occurred approximately ten years after the country's first democratic election and may partly reflect political, rather than economic, considerations. This possibility is reinforced by the government's rejection of a proposed system of need-based tuition fee exemptions, which are arguably less salient, even though they can effectively increase enrollment by low-income students in some settings (Andrews and Stange, 2016).

The paper is organized in five substantive sections. I develop a simple conceptual framework in section 2 that motivates the empirical analysis. Section 3 provides an overview of the South African education system and fee elimination intervention. I then explain how the design of the intervention motivates the identification strategies. Section 4 reports the treatment effects of school fee elimination on school-level enrollment. I then discuss the magnitude of these effects and a number of robustness checks. I explore a number of explanations for these price-insensitive demand estimates in section 5. Section 6 reports treatment effects on measures of school composition, performance and resources. I also consider how these results relate to the enrollment effects. Section 7 concludes.

2 Conceptual Framework

Consider a simple reduced form model of the enrollment decision. Assume that each youth i decides whether to enroll in the local school s at time t. Panel A of figure 1 depicts this decision in a simple demand and supply framework. The vertical axis V shows the value and cost of enrolling for an additional year of schooling in money metric terms. The horizontal axis shows the proportion $P \in [0, 1]$ of youth in this school's neighborhood who are enrolled. The downward-sloping demand curve D captures the idea that the value of enrollment is heterogeneous across individuals within a school.[8] This heterogeneity may arise from different levels of academic ability, different outside options, and the fact that students have reached different grade levels at time t. Assume for now that the cost of enrollment is identical for all agents, so that the cost curve C is horizontal. This assumption is relaxed below. Equilibrium enrollment occurs at p_1. Note that v_1 is the net value of enrollment for the marginal student (gross value less cost), not a market-clearing price.

Eliminating school fees shifts the cost curve downward from C_1 to C_2. The new equilibrium enrollment rate is $p_2 \geq p_1$. Students can be divided into three categories with respect to the fee elimination intervention. p_1 of the students are inframarginal and enroll whether fees are charged

[8] Throughout this section I assume that the demand curve has no discontinuities in the interior. This follows from the assumption that the density of individual valuations of enrollment within a school is strictly continuous.

or not; $1-p_2$ of the students are inframarginal and do not enroll even if fees are not charged; p_2-p_1 of the students are marginal and enroll if and only if fees are eliminated. This yields two general implications of the framework:

I1 Eliminating school fees increases total enrollment and the enrollment rate. Larger baseline school fees will lead to a larger effect on enrollment. If the baseline enrollment rate is at or near one, eliminating fees will have a smaller effect on enrollment.

I2 The change in the enrollment rate equals the proportion of students whose gross value of enrollment is smaller than the cost of enrollment including fees C_1 but larger than the cost of enrollment excluding fees C_2.

This simple framework treats enrollment as a static decision, rather than adopting a dynamic discrete choice framework. The framework also abstracts away from uncertainty. In practice, grade progression is far from universal in South Africa and the potentially stochastic relationship between enrollment and grade progression may vary across schools and individuals within a school (Lam, Ardington, and Leibbrandt, 2010). The value of enrollment V can be interpreted as the expected present value of the discounted stream of future benefits from enrollment, which implicitly includes the option value of enrolling in higher grades in future years. Explicitly accounting for dynamics and uncertainty does not appear to generate any empirical predictions that can be implemented with my data.

Panel B of figure 1 considers two potential valuation curves: D^X is relatively elastic in the neighborhood of the equilibrium and inelastic D^Y is relatively inelastic. Eliminating fees increases the equilibrium enrollment rates from p_1^X and p_1^Y to p_2^X and p_2^Y. The former change is clearly larger, generating another implication of the framework:

I3 The effect of fee elimination on enrollment is increasing in the local elasticity of the valuation curve D.

Panel C of figure 1 considers three potential demand curves: convex D^X, linear D^Y, and concave D^Z. The change in enrollment rates induced by eliminating fees is largest for D^Y and smallest for D^Z. The mean valuation amongst the "always-enrollers" is largest for D^Z and smallest for D^X. This generates another implication of the framework:

I4 There is not a monotonic relationship between the mean valuation amongst students enrolled when school fees are charged and the increase in enrollment rate induced by eliminating fees.

I4 is particularly empirically important. It implies that school-level treatment effects of school fee elimination will not necessarily be correlated with proxies for the valuation of enrollment prior

to elimination. Hence, enrollment may rise by a smaller or larger margin in "good" than "bad" schools. In this depiction, the proportion of always-enrollers and their mean valuation of enrollment are positively correlated. However, this does not hold more generally.

The framework can be extended to allow for a number of more realistic elements. In particular:

I5 If baseline costs of enrollment are heterogeneous, the cost curve will be upward-sloping and the effect of fee elimination on enrollment will be attenuated relative to the constant-cost case.

I6 If schools face binding capacity constraints, the equilibrium enrollment rate will be constrained to be below 1 and the effect of fee elimination on enrollment may be attenuated.

I7 If fee elimination reduces the valuation of enrollment, perhaps due to resource constraints or negative peer effects, the demand curve will shift downward and attenuate or potentially reverse the effect of fee elimination on enrollment.

This framework generates several guidelines for the empirical analysis. First, eliminating fees should increase the enrollment rate (I1). Second, if the enrollment rate falls, this must be due to negative effects of fee elimination on the perceived valuation of enrollment (I7). Third, if the change in the enrollment rate is small, this may reflect near universal baseline enrollment (I1), a low initial value of fees (I1), inelastic demand (I2/3), heterogeneous costs (I5), binding capacity constraints (I6), and/or negative effects on the valuation of enrollment (I7). Fourth, the magnitude of the treatment effects need not be explained by the mean valuation of enrollment amongst students enrolled at baseline (I4).

The framework as written assumes that credit constraints never bind on the enrollment decision. The presence of credit constraints will have an ambiguous impact on the relationship between fee elimination and enrollment. In terms of figure 1 panel A, eliminating fees will induce some credit-constrained students with valuations above v_1 to enroll, However, credit-constrained students with valuations between v_2 and v_1 will not be able to enroll. If the former group is larger than the latter, the treatment effect of fee elimination on enrollment will be increased relative to a world with no credit constraints.

3 Background Information and Identification Strategy

3.1 Background on South African Education

South Africa is a middle income country with a history of sharp economic, political, and social inequality. The public education system was racially segregated until the early 1990s, and per

capita government expenditure on white schools was orders of magnitude larger than on black schools. Both enrollment and high school graduation rates differed sharply by race and there is some evidence of large differences by household socio-economic status (Fedderke, Luiz, and de Kadt, 2000; Seekings and Nattrass, 2005). A small number of black students from low income households enrolled in private schools, mostly church-run. Despite recent growth in the number of small, for-profit private schools in low income communities, the best available data suggest that the private sector remains negligible relative to the public sector and considerably smaller than in South Asian countries (Centre for Development Enterprise, 2010).

State expenditure on black education rose substantially in the 1970s, 1980s, and 1990s and this was associated with rapidly rising enrollment rates. However, substantial evidence suggests that the quality of education remained very low in historically black schools. Curricula at these schools had deliberately focused on non-academic subjects until the 1990s, reflecting the *apartheid* government's insistence on preparing black students for manual employment only. Few students completed secondary schooling, pass rates on high school graduation examinations were low, and even fewer students took mathematics or physical science as high school subjects (Fedderke, Luiz, and de Kadt, 2000). The education system was officially desegregated in the early 1990s but the majority of black students continue to attend historically black schools. Sharp racial differences are visible in which schools students attend and whether they attend school at all, with black enrollment up to 10 percentage points lower than white enrollment.

There are two approaches to estimating baseline enrollment rates and they produce slightly different results. The first approach divides the total number of students that schools report enrolling by the population projected from census data. This method yields national enrollment rates of 85% and 78% for for 7-13 and 14-18 year-old youths respectively. The second approach uses nationally representative household survey data and yields enrollment rates of 96% and 85% in these age brackets.[9] Both approaches suggest that enrollment is high at younger ages and tapers into adolescence. These are *net enrollment rates*, which measure the proportion of the age-eligible population enrolled. This may overstate the enrollment rate in a country with substantial rates of late school entry, grade repetition, and temporary dropout (Department of Education, 2009).

These baseline enrollment rates show that there was substantial room for fee elimination to increase enrollment. I am not aware of any nationally representative dataset that would allow calculation of the baseline enrollment rate in neighborhoods where fees were eliminated. There is a strong negative correlation between youth enrollment and household socio-economic status, suggesting that baseline enrollment was substantially lower in treated than in untreated neighborhoods.

[9] These data are reported in Department of Education (2009) and Department of Education (2011). The survey data are drawn from the General Household Survey, a nationally representative annual project conducted by Statistics South Africa. The survey is relatively comparable to its namesake in the United Kingdom.

The difference between enrollment behavior in low- and high-income households led some policy-makers to advocate interventions to reduce the pecuniary cost of schooling (Pampallis, 2008). Such policies aligned closely with the post-*apartheid* government's long-standing stated commitment to free education. They may also have been motivated by widespread primary school fee eliminations in other African countries during the 1990s and 2000s. Edmonds (2006) shows that rural households that received a fully anticipated income increase (a state old age pension) were more likely to enroll their children than those that did not. This is consistent with credit constrained enrollment decisions.

Low income and credit constraints are not the only reasons offered for low enrollment rates in some neighborhoods. Dropout in secondary school has also been ascribed to the low quality of schools in low income neighborhoods and may be a rational response to low returns to education in these schools. There are no direct measures of returns to education at different types of schools, so these arguments typically infer low returns to education from evidence of low school quality. While measuring "school quality" is a difficult process, it is true that South African students attending schools in low-income neighborhoods perform considerably worse on international literacy and numeracy assessments than poorer students from other African countries (van den Berg and Louw, 2007). South Africa has a system of nominal school choice, under which individuals from low income neighborhoods can in principle enroll in schools in high income neighborhoods. The limited data available on this phenomenon suggests that it is uncommon. This may reflect a combination of high commuting costs in cities that are still highly segregated by income and race, and social and cultural barriers that limit low income students' ability to integrate into schools in high income neighborhoods.

Two interventions were introduced in order to reduce the pecuniary cost of education and promote higher enrollment: means-tested *individual-level school fee waivers* and *school-level school fee eliminations*. The former intervention was introduced in 1996 and required that schools grant partial fee waivers to any household that either earned less than 10 times the per-student school fee or was eligible to receive a means-tested government child grant. The latter requirement meant that a large proportion of the country's students were eligible for the waiver. However, household survey data from 2005 show that approximately 2% of students benefitted from fee waivers and some media reports suggest that parents were discouraged by schools from requesting waivers (Borkum, 2012; Hanes, 2006). This intervention reduced the actual price of education below its nominal level and so risks attenuating any effect of the second intervention on enrollment behavior. I argue below that any bias is likely to be small, due to the rarity of fee waivers and the discontinuous implementation of the fee elimination intervention.

3.2 School Fee Elimination

The school fee elimination intervention was announced in 2006 and implemented in 2007.[10] Schools treated by this intervention were required to eliminate all tuition and enrollment fees, though the status of additional fees for extra-curricular activities was not regulated. These "no fee" schools were chosen by a complex three-stage interaction between provincial and national governments, laid out in guidelines published by the national Department of Education.

In the first stage, provincial governments assigned each school in their province a "poverty score" based on characteristics of the electoral ward in which it was located.[11] These scores ranked all schools within the province from least to most poor, with ties permitted. The national Department of Education provided each province with ward-level data on income, employment, education, health, and "living environment" from the 2001 census as a starting point for the assignment of poverty scores. Provinces were permitted choose their own weighting of these five data series and to make *ad hoc* adjustments to the resultant score based on within-ward heterogeneity. They were not permitted to use any data collected directly from schools, such as administrative data on school's physical facilities or student-teacher ratios. Wildeman (2008) conducted anonymous interviews with provincial officials responsible for creating the poverty scores and reported that most of the *ad hoc* adjustments were made for schools near the boundaries of electoral wards, as the socio-economic characteristics of their students may have differed from those of the electoral ward. Wildeman's interviewees reported no incidents of schools lobbying provincial officials to change their scores, although lobbying may have occurred in the third stage described below. The formulae used to determine the poverty scores were left to the discretion of the provinces and no province has made its formula publicly available.[12]

In the second stage, the national government divided all schools in the country into five "quintiles" based on these poverty scores. Each quintile was intended to contain approximately 20% of the students in the country (based on 2006 enrollment data) and each quintile would contain approximately "equally poor" school neighborhoods in each province. So in the relatively poor Eastern Cape province, 35% and 6% of all schools were assigned to the first and fifth quintiles respectively; in the relatively low poverty Western Cape province, 7% and 23% of all schools were assigned to the first and fifth quintiles respectively. The choice of how many schools were to be treated in each province was based on province-level data from the 2001 census but the exact algorithm used

[10] South Africa's academic year runs from January to December.

[11] The electoral ward is not an administrative unit in South Africa. Assignment took place at this level because it is the smallest geographic unit at which census data is available.

[12] Each province used a different scale for the poverty scores. I standardize these by recentering them at the cutoff between quintiles 2 and 3 and rescaling them to have standard deviation one within each province. The results are reasonably robust to alternative standardizations: rescaling the range to one within each province or rescaling the variance to minimize the sum of the differences between the quintile 1/2 cutoff and the quintile 3/4 cutoff.

for this decision is unclear. All schools in quintiles 1 and 2 were intended to be no fee schools. By determining the number of schools to be treated in each province, the national government implicitly specified a cutoff value of the poverty score above which all schools were the "intention to treat" and below which all schools were the "intention to control" group. National government officials report that the number of schools to be treated was chosen after the poverty scores had already been assigned, so it was not possible for poverty scores to be precisely manipulated in the neighborhood of the cutoff.

In the third and final stage, provincial governments decided which schools were to abolish fees, which created an "actual treatment" group of schools. The actual treatment assignments followed the intended treatment assignments relatively closely: 98% of quintile 1/2 schools above the cutoff eliminate fees while only 6% of quintile 3/4/5 schools below the cutoff do so. The discrepancies may reflect lobbying by schools above the cutoffs who wished to continue charging fees or by schools below the cutoff who wished to eliminate them. The frequency of these discrepancies varies across provinces: 30% of schools have different intended and actual treatment statuses in the least compliant province (Northern Cape), while intended and actual treatment statuses are identical for all schools in one other province (Gauteng).

I use data from nationally representative Income and Expenditure Surveys in 2005/6 and 2010/11 to verify that this school fee elimination actually occurred. I calculate the total expenditure on primary and secondary school fees (distinct from other education costs) in each household and plot the density of fees in figure 2. The proportion of households paying fees fell by 32 percentage points. I cannot match households to schools and hence determine their eligibility for treatment, a comparison of the densities strongly suggests that the reduction in fee-paying is driven entirely by households that were previously paying low fees. This pattern is consistent with the intended design of the fee elimination policy.

3.3 Identification Strategy

The design of the fee elimination intervention makes it a natural candidate for analysis using both difference-in-differences and regression discontinuity methods. I compare schools below the cutoffs, the intended control group, with schools above the cutoffs, the intended treatment group. If the poverty scores are "as good as randomly assigned" (Lee and Lemieux, 2010) in the neighbourhood of the cutoffs, these two groups differ only in their treatment status and so any differences in enrollment between the two groups may be interpreted as a causal effect of the fee elimination intervention.

Baseline difference-in-differences specification: My primary estimation sample consists of schools in quintiles 2 and 3. I use quintile 2 schools as the "intention to treat" group and quintile

3 schools as the control group. By eliminating quintile 1, 4, and 5 schools, I restrict the sample to schools with relatively similar baseline characteristics in line with the spirit of the regression discontinuity design. I begin by estimating

$$\text{Fee elimination}_i = \alpha_0 + \alpha_1 \text{High poverty score}_i + \epsilon_i \tag{1}$$

by ordinary least squares. The coefficient α_1 captures the difference in the probability of fee elimination between schools with high poverty scores (i.e. in quintile 2) and schools with low poverty scores (i.e. in quintile 3). This tests whether assignment to treatment is broadly consistent with the process described in the previous subsection. I use the cross-section of schools in 2007 to estimate this model.

I then estimate

$$\begin{aligned} \text{Enrollment}_{it} =& \beta_0 + \beta_1 \text{High poverty score}_{it} + \beta_2 \mathbf{1}\{\text{Year} \geq 2007\} \\ &+ \beta_3 \text{High poverty score}_{it} \times \mathbf{1}\{\text{Year} \geq 2007\} + \nu_{it} \end{aligned} \tag{2}$$

The coefficient β_3 captures the difference in the change in enrollment from pre-2007 to post-2007 between schools with high poverty scores (i.e. in quintile 2) and schools with low poverty scores (i.e. in quintile 3). This tests whether schools in the intention to treat group experience a larger increase in enrollment from 2005/6 to 2007/8 than other schools. I use a panel of school-level enrollment between 2005 and 2008 to estimate this model and I use a cluster-robust variance estimator that allows unrestricted intertemporal correlation in ν_{it} for each school i.[13]

I finally estimate

$$\begin{aligned} \text{Enrollment}_{it} =& \gamma_0 + \gamma_1 \text{Fee elimination}_{it} + \gamma_2 \mathbf{1}\{\text{Year} \geq 2007\} \\ &+ \gamma_3 \text{Fee elimination}_{it} \times \mathbf{1}\{\text{Year} \geq 2007\} + \eta_{it} \end{aligned} \tag{3}$$

using instrumental variables, with the indicator for fee elimination instrumented by high poverty score indicator and the interaction term treated analogously. The coefficient γ_3 captures the difference in the change in enrollment from pre-2007 to post-2007 between no fee schools and fee charging schools who comply with their intended school fee policy. This captures the effect of fee elimination on enrollment for schools that comply with their intended treatment status. I again use a panel of school-level enrollment between 2005 and 2008, with a cluster-robust variance estimator. The results in section 4 verify that the instruments used in estimating equation (3) easily pass the appropriate tests for instrument strength.

[13] I pool 2005 and 2006 together and 2007 and 2008 together to smooth out potential measurement error in the reported enrollment data. Formal tests do not reject equality of enrollment in 2005 and 2006 and in 2007 and 2008.

Identification of $(\alpha_1, \beta_3, \gamma_3)$ relies on the assumption that the counterfactual trend in enrollment for quintile 2 and quintile 3 schools from 2005/6 and 2007/8 would have been identical if fees had not been eliminated. This assumption may be problematic if the two groups of schools differ on observed or unobserved characteristics that are associated with enrollment trends. If such "confounding" occurs, the identification assumption will fail. I use two strategies to address this potential concern.

Reweighted difference-in-differences specification: I first consider the possibility that quintile 2 (intention to treat) and quintile 3 (control) schools may differ on observed characteristics. The panel structure of the difference-in-differences design is equivalent to including school-level fixed effects in equations (1) – (3). This accounts for any differences in the level of enrollment between intention to treat and control schools. However, it does not account for the possibility that enrollment trends may be systematically correlated with baseline observed characteristics that differ systematically between intention to treat and control schools. I observe data on a vector of baseline characteristics from 2005 and 2006: enrollment, number of grades offered, phase (primary, intermediate, or secondary), location (urban or rural), historical racial classification (Asian, black, white, mixed race, or founded after desegregation), designation as a mathematics and science specialization school, partial self-governance status, class size, student-teacher ratio, proportions of part-time and temporary teachers, dropout rate, grade promotion rate, proportion of orphans in the school, and proportion of students in the school whose families receive government social grants. The means of almost all of these characteristics differ significantly between quintile 2 and 3 schools and the χ^2 test statistic for equality of all characteristics is 2323. I therefore construct a sample of control schools weighted to have the same distribution of observed characteristics as the intention to treat schools. I follow Abadie (2005) and DiNardo, Fortin, and Lemiuex (1996) in using the reweighting function

$$\omega(X_{it}) = \frac{Pr(\text{High poverty score}_i = 1|X_{it})}{1 - Pr(\text{High poverty score}_i = 1|X_{it})} \tag{4}$$

for all schools in the control group. This term assigns high weight to control schools whose observed characteristics X_{it} in the baseline period (2005 and 2006) "look like" those of the intention to treat schools.[14] I estimate the predicted probability using a logistic regression of an indicator for high poverty scores (i.e. quintile 2 schools) on the full vector of observed characteristics and quadratic terms in the continuous variables.[15] I then estimate equations (1) – (3) using weighted least

[14] I include missing data indicators where values of any of the observed characteristics are not reported in the dataset.

[15] Although the two groups differ on average characteristics, there is reasonable overlap in the predicted probabilities $\hat{Pr}(\text{High poverty score}_i = 1|\text{High poverty score}_i = 1, X_{it})$ and $\hat{Pr}(\text{High poverty score}_i = 1|\text{High poverty score}_i = 0, X_{it})$. The maximum predicted probabilities in the two samples are almost identical and less than 1% of the control observations have predicted probabilities below the minimum in the intention to treat group.

squares.[16]

After constructing the weights, I test whether the means of the observed characteristics differ between the intention to treat and the reweighted control group. I fail to reject the null hypothesis of equal means for any of the individual observed characteristics and fail to reject the null of joint equality (χ^2 test statistic of 28.3 with p-value 0.45) from a weighted seemingly unrelated (SUR) regression. This confirms that the reweighted sample of control schools are statistically indistinguishable from the intention to treat schools on observed baseline characteristics. Hence, the weighted least squares estimates from equations (1) – (3) will be purged of any confounding due to differences in enrollment trends correlated with observed school-level characteristics.

Regression discontinuity differences specification: The strategy above will account for any confounding due to differences in observed baseline characteristics between intention to treat and control schools. However, it will not account for confounding due to differences in *unobserved* characteristics that are correlated with enrollment trends between the two groups. It will also be problematic if enrollment trends are correlated with socio-economic status as measured by poverty scores. These determine schools' status in the intervention and so are disjoint between intention to treat and control schools and cannot be used in the reweighting algorithm. I address this concern by taking advantage of the process used to assign schools to treatment status. Specifically, I estimate

$$\begin{aligned}\text{No fee school}_i =& \mathbf{1}\{\text{Poverty score} \geq 0\} \times f^+(\text{Poverty score}_i) \\ &- \mathbf{1}\{\text{Poverty score} < 0\} \times f^-(\text{Poverty score}_i) + \epsilon_i\end{aligned} \tag{5}$$

and

$$\begin{aligned}\Delta\text{Enrollment}_i =& \mathbf{1}\{\text{Poverty score} \geq 0\} \times f^+(\text{Poverty score}_i) \\ &- \mathbf{1}\{\text{Poverty score} < 0\} \times f^-(\text{Poverty score}_i) + \epsilon_i\end{aligned} \tag{6}$$

where f^+, f^-, g^+, and g^- are polynomial or local linear functions of poverty scores. The idea behind these models is to specify flexibly the relationship between the outcome of interest (respectively, no fee status and the change in enrollment from 2005/6 to 2007/8). I can then evaluate

$$\lim_{\text{Poverty score}\downarrow 0} \hat{f}^+(\text{Poverty score}_i) - \lim_{\text{Poverty score}\uparrow 0} \hat{f}^-(\text{Poverty score}_i)$$

to estimate the magnitude of the change in the outcome of interest that occurs in the neighborhood of the cutoff poverty score that separates intention to treat and control schools. Provided schools are

[16] I approximate the standard errors of the resultant estimators using 100 replications of a boostrap algorithm that iterates over both stages of the estimation: construction of the weights and estimation of the linear difference-in-differences models. The bootstrap algorithm resamples school clusters with replacement.

not able to manipulate precisely the poverty score they are assigned by province, untreated schools on one side of the cutoff should be a valid counterfactual for treated schools on the other side of the cutoff. Note that I use a time-differenced outcome variable in equation (6) so the specification already removes any time-invariant observed or unobserved characteristics. This identification strategy may be considered "doubly robust" relative to standard regression discontinuity designs that use only cross-section data.

This identification strategy generates two testable predictions. First, I verify that none of the observed characteristics listed above have statistically significant or economically meaningful "jumps" at the cutoff. Second, I verify that there is no evidence that the density of the poverty score variable jumps at the cutoff (see figure 3). McCrary (2008) notes that such a jump would be consistent with manipulation of the poverty scores in order to control treatment assignment.

4 Effect of fee elimination on enrollment

Figure 4 shows the time trend in enrollment rates calculated from the General Household Survey. Comparing 2006 and earlier years to 2007 and subsequent years suggests that enrollment amongst primary school-aged youths was largely unaffected by the intervention, while enrollment amongst secondary school-aged youths rose very slightly. Table 1 presents the more formal difference-in-differences analysis comparing quintile 2 (intention to treat schools) to quintile 3 (control schools). Column 1 shows that quintile 2 schools are 94 percentage points more likely to eliminate school fees than quintile 3 schools. This high rate of compliance with the intervention means that poverty score is a strong instrument (first stage F-statistic over 50000) for fee elimination. Column 2 shows that enrollment in schools that eliminated fees rose from 2005/6 to 2007/8 by 3.6 students more on average than in other schools. The 95% confidence interval for this intention to treat estimate is 0 to 7.1 students and the point estimate is marginally significantly different to zero. The corresponding instrumental variables estimate in column 3 is 3.8 students per school (95% confidence interval 0 to 7.5).

This effect size can be expressed in several different metrics. First, it directly measures the number of additional students induced to enroll by the intervention in each school. Second, the intention to treat and instrumental variables estimates respectively imply 0.9% and 1% increases in total baseline enrollment (confidence intervals respectively 0 to 1.8% and 0 to 1.9%). Third, the effect can be converted into a change in the enrollment rate under some additional assumptions. The baseline enrollment rate for youth aged 7-18 in 2005/6 in the General Household Survey was 92%. If this enrollment rate applied to the neighborhoods around fee-eliminating schools, the intervention

would increase the enrollment rate by 0.83 (ITT) to 0.91 (IV) percentage points.[17] The actual enrollment rate is strongly negatively correlated with household socio-economic status and is likely to be considerably lower near treated schools. Hence, the change in the enrollment rate for any given change in the level of enrollment will be smaller. The upper bound of the 95% confidence intervals allow me to rule out effect sizes larger than 1.7 percentage points for the intention to treat estimate or 1.8 percentage points for the instrumental variables estimates. This implies that a minimum of 6% of youths remained unenrolled despite the elimination of fees, with the rate in affected neighborhoods probably substantially higher.

Converting this effect size into an elasticity is complicated by the lack of available data on baseline school fees. The South African Department of Education has collected these data for all schools around the country but I have not yet been able to obtain access to them. Without these data, it is not possible to assign a monetary value to the eliminated fees or to calculate how this compares to other pecuniary costs of enrollment (transport, uniforms, books, etc.). I use two sources of survey data to obtain a rough value of the elasticity. First, the proportion of youths aged 7-18 in the General Household Survey who were enrolled and paying school fees fell from 98% in 2005/6 to 73% in 2007/8. Although these respondents cannot be matched to schools and hence to treatment status, I assume that the proportion of students in each quintile is equal and that all students not paying fees at baseline were in the treated quintiles 1 and 2. This implies a baseline fee payment rate of 95%. If all students who paid fees after the intervention were in the control quintiles 3 - 5, then the fee payment rate in treated schools fell to 32%.[18] This implies that the arc price elasticity of enrollment with respect to paying any fees is

$$\begin{aligned}\epsilon_1 &= \frac{\Delta\text{Enrollment rate}}{\Delta\text{Fee payment rate}} \times \frac{\text{Fee payment rate}}{\text{Enrollment rate}} \\ &= -\frac{0.0083}{0.63} \times \frac{0.6 \times 0.95 + 0.4 \times 0.32}{0.92} \\ &= -0.010\end{aligned}$$

for the intention to treat estimate and -0.011 for the instrumental variables estimate. Even using the upper bound of the 95% confidence intervals for the change in the enrollment rate yields elasticities of -0.019 and -0.020.

The second approach to calculating an elasticity uses calculations by Branson, Lam, and Zuze (2012) based on data from the National Income Dynamics Survey.[19] They report that the average

[17] This follows from multiplying the baseline enrollment rate by the percentage increase in the baseline enrollment level.

[18] This calculation abstracts away from the fact that the number of enrolled students is changed by the intervention. The small magnitude of this change means that the results are robust to taking this into account.

[19] The National Income Dynamics Survey is a nationally representative panel dataset that began in 2008. The baseline questionnaire included some questions on retrospective schooling history. Branson, Lam, and Zuze (2012)

school fees paid by enrolled youths of any age in 2007 were 64 rands[20] in quintile 1/2 schools and R301 rands in quintile 3/4 schools. Under the extreme assumption that these values were identical prior to the fee elimination, the intervention reduced fees by 79% of their baseline value. If fees were the only pecuniary cost of education, this implies that the arc price elasticity of enrollment with respect to the value of fees

$$\begin{aligned}\epsilon_2 &= \frac{\Delta\text{Enrollment rate}}{\Delta\text{Fee amount}} \times \frac{\text{Fee amount}}{\text{Enrollment rate}} \\ &= -\frac{0.0083}{237} \times \frac{0.6 \times 301 + 0.4 \times 64}{0.92} \\ &= -0.008\end{aligned}$$

for the intention to treat estimate and - 0.009 for the instrumental variables estimate. The price elasticity of enrollment with respect to the total pecuniary cost of enrollment $\tilde{\epsilon}_2$ will be weakly larger. If, for example, fees are only half of the total cost of enrollment, then $\tilde{\epsilon}_2$ = -0.023 (ITT) or -0.026 (IV).[21]

The discussion above implies that fee elimination induced less than four additional students per school to enroll, increasing the baseline enrollment level by 1% and the baseline enrollment rate by less than one percentage point. Informal calculations suggest elasticities with respect to any fee or to the value of the fee no larger than -0.01. By any measure, this implies that demand for enrollment in the treated schools was highly price insensitive. Demand for broader measures of school participation that also take into account attendance by marginal students may be even lower. The remainder of this section shows that the result is robust to a variety of alternative estimation strategies. Section 5 explores potential explanations for this result.

Robustness checks: The results reported above are valid estimates of the causal effect of the fee elimination intervention on enrollment only if fee-charging and fee-eliminating schools would have experienced the same change in enrollment from 2005/6 to 2007/8 in the absence of the intervention. This assumption may fail if the two groups of schools have different observed or unobserved characteristics. To address this possibility, I estimate equations (1) – (3) using the weights in equation (4) to equalize the distribution of observed baseline school characteristics. Column 7 of table 1 shows that the difference in fee elimination rates between high and low poverty schools is unaffected by reweighting. However, the enrollment effect is halved from 3.6 students per school to 1.7 students per school (95% confidence interval -1.5 to 4.9). This implies even more

are able to match approximately 90% of enrolled youths to schools in the Department of Education's public database.

[20] One South African rand was equal to approximately $0.145 in 2007 without adjusting for purchasing power parity.

[21] To derive this result, define E, F, and O as respectively the enrollment rate, the cost of fees and the cost of other education inputs (transport, uniforms, etc.). Then $\tilde{\epsilon}_2 = \frac{\Delta E}{\Delta(F+O)} \times \frac{F+O}{E}$. If O is unaffected by fee elimination, then $\tilde{\epsilon}_2 = \frac{\Delta E}{\Delta F} \times \frac{F}{E} + \frac{\Delta E}{\Delta F} \times \frac{O}{E} = \epsilon_2 \times \left(1 + \frac{O}{F}\right)$.

price insensitive demand than in the unweighted results. I also transform equations (3) – (3) from level to first difference specifications and include the vector of baseline school characteristics in the regression. The results, shown in columns 4 – 6 of table 1, are essentially identical to those from the unweighted regression. The "doubly-robust" analysis that uses regression and reweighting generates results very similar to those using just reweighting. All of these results suggest a very high rate of compliance with the intervention by schools and small effects on enrollment.[22]

An alternative approach to robustness is to estimate treatment effects for schools in the neighborhood of the poverty score cutoff between intention to treat and control schools. Figures 5 and 6 show estimates of the regression discontinuity models in equations (5) and (6) respectively. The estimates are obtain using local linear regression with bandwidth chosen to minimize the mean squared error of the treatment effect, following Imbens and Kalyanaraman (2012). The first figure shows that the probabilty of fee elimination rises from approximately 8% for schools just less poor than the cutoff to approximately 86% for schools just more poor than the cutoff. The difference of 76 percentage points is smaller than the difference between all intention to treat and control schools but is still substantial. The second figure shows that enrollment increases from 2005/6 to 2007/8 by 2.8 students more in intention to treat than control schools (confidence interval -0.7 to 6.3). This is slightly smaller than the difference for all schools, though the estimates are not statistically distinguishable at conventional significance levels. I also implement the regression discontinuity using global linear, quadratic and cubic models, and comparing means in the 50%, 25%, and 10% of the sample closest to the cutoff. The point estimates for the change in enrollment are somewhat sensitive to specification but are all between 0.8 and 4 students. This reinforces the earlier result that fee elimination had a relatively small effect on student enrollment.

As a final robustness check, I consider a longer time series of enrollment that includes four years of pre-treatment data. Figure 7 shows the raw enrollment data for intention to treat and control schools. The two groups have very different pre-treatment levels of enrollment but there is little evidence of differential trends. The enrollment time series for the reweighted control schools looks almost identical to the treated schools. The graph supports the identifying assumption of parallel pre-treatment trends. I use a more formal test of this idea from Heckman and Hotz (1989) by constructing a "difference-in-second-differences" estimator. This equals the conventional difference-in-differences estimator (2007/8 versus 2005/6 enrollment in intention to treat and control schools) minus the pre-treatment difference-in-differences estimator (2005/6 versus 2003/4 enrollment in intention to treat and control schools). This is a consistent estimator of the treatment effect of

[22]The difference between the weighted and unweighted results reflects heterogeneous enrollment trends over some observed characteristics, discussed in section 5 below. The unweighted estimators, with or without regression adjustment, are not consistent estimators of the relevant treatment effect in the presence of this form of heterogeneity. The identifying assumptions upon which the unadjusted and regression-adjusted models are based are subtly different, as discusssed in Imbens and Wooldridge (2009).

interest under the weaker assumption that the two groups have linear but potentially non-parallel trends and are subject to common shocks. The estimated values of this parameter with and without reweighting are respectively 5.9 and 3.1 students (standard errors 4.7 and 4.2), compared to the standard estimate of 3.6 and 1.7 students (s.e. 1.8 and 1.6). This implies that pre-treatment enrollment in intention to treat schools was declining slightly relative to control schools. However, it still implies a small overall effect on enrollment.

5 Explaining the price-insensitive demand

This section builds off the conceptual framework developed in section 2 to explore possible reasons for the small enrollment effects that I find. The framework suggests that this may be due to (1) high baseline enrollment or "ceiling effects," (2) low baseline fees, (3) inelastic demand, (4) binding capacity constraints, (5) credit constraints, or (6) negative effects on the valuation of enrollment. I do not directly observe baseline fees and so cannot test the second explanation. I explore the first, fourth and fifth explanations in this section and find that they cannot fully account for the small effects. The sixth explanation is left to the next section. The analysis as a whole suggests an important role for the residual explanation, inelastic demand.

Treatment effects by poverty level: The analysis to date has concentrated on quintile 2 schools and omitted the quintile 1 schools in poorer neighborhoods that also eliminated fees. Given the positive correlation between household income and enrollment found in survey data, baseline enrollment rates are likely to be lower in quintile 1 than quintile 2 schools. There is thus less scope for ceiling effects on enrollment. I estimate quintile-specific versions of the difference-in-differences model in equations (1) – (3), with and without reweighting.[23] Table 2 column 1 reports that the treatment compliance rates are very high: 99 and 95% for quintile 1 and 2 schools respectively. These are unaffected by reweighting. Enrollment increased from 2005/6 to 2007/8 in quintile 1 schools by 7.7 students (95% confidence interval 4.3 to 11 students), which represents a 2% increase in baseline enrollment. The equivalent change in quintile 2 schools is 3.4 students (confidence interval -0.1 to 7 students) or 0.9% of baseline enrollment. Combining these estimates using instrumental variables shows that fee elimination increased enrollment by an average of 6.1 students across all fee eliminating schools. As in the previous section, reweighting substantially reduces the magnitude of the treatment effects.

[23] Reweighting quintile 3 schools to "look like" quintile 2 schools was successful in the sense that the mean values of the treatment and reweighted control schools' baseline characteristics are neither substantively nor significantly different. Applying the same procedure to compare quintile 1 and 3 schools was less successful. The χ^2 test statistic for the null hypothesis that the means of all 28 baseline characteristics are equal across quintile 1 and reweighted quintile 3 schools is 54.3 (p-value 0.002). It is only possible to reduce this difference by using substantially higher order polynomial expressions in the reweighting functions in equation (4). The estimated treatment effects in quintile 1 should therefore be interpreted with a degree of caution.

The significantly larger change in enrollment in quintile 1 schools may be due to ceiling effects on enrollment in quintile 2 schools. It may also reflect more elastic demand in quintile 1 schools.[24] This result is unlikely to be generated by differences in baseline school fees. Fees were higher in quintile 2 than quintile 1 schools (Borkum, 2012), implying larger treatment effects in the former unless students are credit constrained. I discuss the possibility of credit constraints in more detail below.

Grade-specific treatment effects: Baseline enrollment rates calculated from the General Household Survey are substantially higher for primary school-aged youths (96%) than secondary school-aged youths (85%). Even if these overstate the relevant enrollment rates due to grade repetition, they suggest that ceiling effects on enrollment are substantially more likely in primary than secondary school. I therefore estimate treatment effects on enrollment in each grade from 0 to 12 using basic and reweighted difference-in-differences models. The results are shown in figure 8. Panels A and C show that there are near-zero enrollment changes in quintile 1 and quintile 2 primary schools respectively. The treatment effects are largest in the early grades of secondary school (8-10) and largely die out by grades 11 and 12. The reweighted estimates in panels B and D have a similar pattern though the point estimates are considerably smaller.

What do these effect sizes imply for enrollment rates? Averaging across quintile 1 and 2 schools, enrollment increases by 8.9 students in each of grades 8 and 9 and by 0.4 students in each grades 10, 11 and 12. These imply respective increases of 14.8% and 0.4% of the baseline enrollment level.[25] The 2005 and 2006 General Household Surveys report enrollment rates of 87% for ages 14-15 and 83% for ages 16-18. Matching these age and grade brackets implies that eliminating fees increased age 14-15 enrollment from 87% to almost 100% and left age 16-18 enrollment almost unchanged at 83%.[26]

Capacity constraints: Some schools may have a binding upper limit on the number of students they can accommodate, due to limited personnel, classroom space or other physical facilities. This would lead them to deny enrollment to students who would otherwise be induced to enroll by the policy change. The legal status of such denials is subject to an ongoing court challenge but anecdotal reports suggest that it is uncommon. I formally test for the existence of capacity constraints in two steps. I first calculate the maximum enrollment in each school between 2003 and 2006. This

[24] These explanations are not entirely independent. A highly concave demand curve will result in near universal baseline enrollment and will be inelastic in the range affected by fee elimination.

[25] The difference between the level and percentage measures arise because the average school contains 53 students in each of grades 8 and 9, compared to 112 students in each of grades 10, 11 and 12.

[26] Matching age and grade brackets assumes that no youths start school after the mandatory age, repeat grades or temporarily drop out of schooling. This assumption is clearly incorrect but unavoidable. In practice, many of the youths aged 16-18 will be enrolled in grades 8 and 9. This will bias the estimated treatment effect on the enrollment rate. The direction of the bias depends on the relative frequency of overage students in the inframarginal and marginal populations.

is one measure of each school's maximum capacity. I then compare the frequency with which 2007 or 2008 enrollment exceeds the previous maximum by quintile. This occurs in 26% of quintile 3 (control), 26% of quintile 2 schools, and 33% of quintile 1 schools. A substantial fraction of schools are thus able to accommodate additional students and this proportion is higher in fee-eliminating than fee-charging schools. Furthermore, the change in enrollment from 2006 to 2007 is smaller than at least one previous annual change in enrollment (2003 to 2004, 2004 to 2005 or 2005 to 2006) in over half of the sample. Hence, many schools are able to accommodate larger increases in enrollment than they experience when fees are eliminated. The results do not conclusively rule out capacity constraints but they do not appear to be of central importance.

Credit constraints: How important might credit constraints be in explaining the pattern of results discussed above? Section 2 noted that credit constraints might attenuate or augment the treatment effects of fee elimination. Edmonds (2006) reports some evidence of credit constraints to school enrollment in rural but not urban areas. Lam, Ardington, Branson, Goostrey, and Leibbrandt (2010) study a largely urban population and find little evidence of credit constraints to enrollment in tertiary education, which is typically much more expensive than primary or secondary enrollment. These results motivate estimating treatment effects separately for urban and rural schools.[27] Table 3 reports the results of this exercise. The compliance rates are equally high in both groups of schools. The enrollment effects are substantially larger in urban than rural areas, although the urban sample is smaller and the estimates less precise.[28]

If credit constraints are indeed present in rural but not urban areas, these results suggest that there are a substantial number of rural students who would be induced to enroll by fee elimination if not for credit constraints. This is consistent with Edmonds' results, as the income shock he uses to test for credit constraints is substantially larger than the value of school fees. However, the results could also be due to larger baseline fees in urban areas or more elastic demand in urban areas. Even the larger effects in urban areas imply increases of less than 2% of baseline enrollment in quintile 1 and quintile 2 schools. Eliminating fees therefore leaves a large fraction of the population unenrolled in both rural and urban areas.

Taken together, these results demonstrate that the small enrollment effects are not explained by ceiling effects, capacity constraints or credit constraints. The next section explores whether they may be explained by a declining valuation of enrollment.

[27] This includes all non-rural schools, urban and suburban.

[28] The substantial difference between urban and rural schools' enrollment results explains the discrepancy between the reweighted and regression-adjusted results in table 1. The regression-adjusted estimates fail to take this heterogeneity into account.

6 Effects of fee elimination on school composition and outcomes

Fee elimination may change the resources available to schools and the composition of their student body. Compensating government transfers may be larger or smaller than the foregone fee revenue. Even if per-student revenues are unaffected, education inputs such as classrooms and teachers may adjust to changes in student numbers with a lag. Student expectations about changes in education resources may in turn affect their enrollment decisions and so attenuate or augment the enrollment response to fee elimination.

This section explores the equilibrium effects of fee elimination on education outcomes (measured by dropout rates, grade repetition rates, and class sizes) and the socio-economic profile of the enrolled students (measured by the proportion of students eligible for means-tested government social grants and the proportion who have had at least one parent die). These effects are equilibrium in that they are conditional on student enrollment decisions. They provide no direct information about the function mapping enrollment to education outcomes or *vice versa*. The section also explores the enrollment trends in control schools close to and far from treatment schools in order to address the possibility of spillovers.

Education outcomes: Table 4 column 2 shows fee elimination causes the dropout rate in fee-eliminating schools to fall from 2.8% to 2.1%. The rate of grade repetition also falls from 9.6% to 9.2% (column 4), though this effect is quite imprecisely estimated. Average class size increases from 39.3 to 39.8 students (column 6). In sum, fee elimination appears to have marginally increased enrollment and decreased per student education resources but reduced grade repetition and substantially reduced dropout. The dropout effect is most striking relative to its baseline level. This pattern would arise if marginal students are substantially less likely to drop out than their inframarginal peers or if fee elimination reduces credit constraint-induced dropout by inframarginal students. None of the other effects are large enough to suggest that education outcomes in schools are substantially affected by fee elimination. Future research will also examine whether high school completion, measured by passing a national graduation examination, is affected by fee elimination.

Socio-economic profile: The effect of fee elimination on the socio-economic profile of the enrolled student population is shown in columns 8 and 10 of table 4. The proportion of students who have lost at least one parent to death falls from 20.7% to 19.5%[29] while the proportion of students eligible for means-tested government social grants stays constant. The socio-economic profile of the fee-eliminating schools appears relatively stable, suggesting that the marginal students

[29]The effect on the proportion of students who have lost both parents to death falls by 0.2 percentage points from a baseline of 3.8%. The patterns for both definition of orphan are therefore consistent.

are not substantially different on these dimensions to the inframarginal students. However, these data are collected as teachers' reports of students' self-reports and so are likely to be measured with substantial error. The effects should therefore be interpreted with caution.

Transfers between fee-eliminating and fee-charging schools: My estimation strategy assumes that the school fee elimination policy has no effect on enrollment levels at the control schools that continue to charge fees. This assumption may be violated if students who attended control schools before the policy change transfered to treatment schools after the policy change to take advantage of their lower cost. Such behavior would result in an upward bias in the estimated treatment effect of the fee elimination intervention.

I do not observe student-level data on transfers that would permit a direct test of this hypothesis. I therefore implement an indirect test that examines whether control schools that are geographically closer to treatment schools experience falls in enrollment from 2006 to 2007 relative to farther away control schools and relative to their own enrollment change in previous years. Figure 9 shows a local linear regression of the change in grade-level enrollment from 2006 to 2007 at control schools against the distance from the nearest treatment school offering the same grade.[30] Control schools nearer to treatment schools actually experience small gains in enrollment relative to control schools farther away. I cannot reject that this pattern of changes is identical to that observed between 2005 and 2006, before the fee elimination policy was implemented. The result is robust to restricting the sample to control schools within one half standard deviation of the cutoff. I interpret this as strong evidence against the spillover hypothesis.

I also estimate a linear regression of change in enrollment by grade at fee-charging schools from 2006 to 2007 on the same measure at the nearest fee-eliminating school. If the treatment effect is driven entirely by transfers, the slope coefficient should be approximately equal to -1. Instead, it equals 0.045 (standard error 0.016). This is not significantly different to the coefficient in the equivalent regression using changes from 2005 to 2006 (0.037, with standard error 0.08). This result is robust to weighting the regression by the inverse distance between treatment and control schools, to restricting the sample to control schools within one half standard deviation of the cutoff, and to excluding control schools that are more than 10 miles from the nearest treatment school. These results strongly suggest that the treatment effects are not driven by transfers between schools. However, I observe only net transfers and not gross inflow and outflow of students into each school. I cannot rule out the possibility that approximately equal numbers of students transfer in each direction between fee-charging and fee-eliminating schools.

One interpretation of these results is that geographically targeted variation in the prices of public services may be an effective alternative to individual-level means testing. The absence of spillovers

[30] I construct this distance measure using GIS codes for every school in the sample.

in this context suggests that setting lower prices for public services in poor neighborhoods does not induce people from wealthier neighborhoods to adjust their behavior to take advantage of the lower cost services. Geographic targeting may be a desirable alternative to individual means-testing when the latter is expensive. However, caution should be exercised in generalizing this result to settings where use of public services is more price-sensitive or transport costs are lower.

7 Conclusion

Increasing participation in formal schooling is an important thrust of public policy throughout the developing world. A wide range of countries have employed a wide range of demand- and supply-side interventions to increase student enrollment and attendance. Empirical work in microeconomics has found evidence of positive relationships between formal schooling and individual earnings, health and future children's human capital accumulation, at least some of which seem to capture causal effects of education. The macroeconomic literature points to a potential role of formal schooling in explaining cross-country differences in per capita income and growth.

This paper contributes to this literature and policy by studying the effect of geographically targeted school fee eliminations on primary and secondary school enrollment in South Africa. Fee-reducing and -eliminating interventions have become popular in developing countries over the past two decades. There is, however, relatively little empirical evidence on their effects. I find that eliminating fees to enroll in South African schools in high-poverty neighborhoods increased enrollment by a relatively small margin. My preferred estimates suggest that 1 to 3 additional students were induced to enroll in the average school. This increased the baseline enrollment level by 1% or less and the baseline enrollment rate by less than 1 percentage point. Back-of-the-envelope calculations suggest a price elasticity of enrollment in the neighborhood of -0.01. School-level composition, education outcomes and resources are also largely unaffected by fee elimination.

I explore a variety of explanations for this pattern. The results are not explained by pre-treatment time trends, pre-treatment differences between high- and low-poverty schools' characteristics, overstated enrollment levels, or transfers between fee-eliminating schools and fee charging schools. They are not fully explained by ceiling effects on enrollment, capacity constraints in schools or credit constraints in households.

I therefore conclude that demand for schooling is relatively price insensitive in the neighborhoods treated by the fee elimination intervention. This is more likely due to low returns to education for marginal students in these schools than to high labor market opportunity costs. A substantial body of prior research has documented low "quality" of education in high-poverty South African schools, measured by enrolled students' test scores and graduation rates. Particularly low returns

to enrollment for marginal students would be unsurprising.

My results imply that reducing enrollment costs may have limited impact on school participation levels in settings where enrollment does not translate into substantial learning or grade progression. This emphasizes a potential complementarity between cost-reduction and quality-upgrading interventions and points to an important avenue for future research.

References

Abadie, A. (2005): "Semiparametric Difference-in-Differences Estimators," *Review of Economics and Statistics*, 72, 1–19.

Andrabi, T., J. Das, and A. Khwaja (2014): "Report Cards: The Impact of Providing School and Child Test Scores on Educational Markets," Working paper, Pomona College.

Andrews, R., and K. Stange (2016): "Price Regulation and Equality of Opportunity in Higher Education: Evidence from Tuition Deregulation in Texas," Working paper, University of Texas at Dallas.

Baird, S., C. McIntosh, and B. Ozler (2011): "Cash or Condition? Evidence from a Cash Transfer Experiment," *Quarterly Journal of Economics*, 126(4), 1709–1753.

Barrera-Osorio, F., L. Linden, and M. Urquiola (2007): "The Effects of User Fee Reductions on Enrollment: Evidence from a Quasi-experiment," Working paper, Columbia University.

Borkum, E. (2012): "Can Eliminating School Fees in Poor Districts Boost Enrollment? Evidence from South Africa," *Economic Development and Cultural Change*, 60(2), 359–398, Forthcoming.

Branson, N., C. Ardington, D. Lam, and M. Leibbrandt (2013): "Changes in education, employment and earnings in South Africa – A cohort analysis," Discussion Paper 105, Southern Africa Labour and Development Research Unit.

Branson, N., D. Lam, and L. Zuze (2012): "Education: Analysis of the NIDS Wave 1 and 2 Datasets," Discussion Paper 81, Southern Africa Labour and Development Research Unit.

Branson, N., and M. Leibbrandt (2013): "Education quality and labour market outcomes in South Africa," Discussion Paper (2013)13, Organisation for Economic Cooperation and Development.

Card, D. (1999): "The Causal Effect of Education on Earnings," in *Handbook of Labor Economics Volume 3A*, ed. by D. Card, and O. Ashenfelter, pp. 1801–1863. Elsevier.

Centre for Development Enterprise (2010): "Hidden Assets: South Africa's Low-Fee Private Schools," .

Deininger, K. (2003): "Does the Cost of Schooling Affect Enrollment by the Poor? Universal Primary Education in Uganda," *Economics of Education Review*, 22, 291–305.

Department of Education (2009): *Trends in Education Macro Indicators Report*. Government Printers, Pretoria, ZA.

——— (2011): *Trends in Education Macro Indicators Report*. Government Printers, Pretoria, ZA.

DiNardo, J., N. Fortin, and T. Lemiuex (1996): "Labor Market Institutions and the Distribution of Wages, 1973 - 1992: A Semiparametric Approach," *Econometrica*, 64(5), 1001–1044.

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

Dupas, P. (2014): "Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence From a Field Experiment," *Econometrica*, 82(1), 197–228.

Dynarski, S. (2003): "Does Aid Matter? Measuring the Effect of Student Aid on College Attendance and Completion," *American Economic Review*, 93(1), 279–288.

Dynarski, S., J. Gruber, and D. Li (2009): "Cheaper by the Dozen: Using Sibling Discounts at Catholic Schools to Estimate the Price Elasticity of Private School Attendance," Discussion Paper 15461, National Bureau of Economic Research.

Edmonds, E. (2006): "Child Labor and Schooling Responses to Anticipated Income in South Africa," *Journal of Development Economics*, 81(2), 386–414.

Evans, D., M. Kremer, and M. Ngatia (2009): "The Impact of Distributing School Uniforms on Children's Education in Kenya," Mimeo.

Fafchamps, M., and B. Minten (2007): "Public service provision, user fees and political turmoil," *Journal of African Economies*, 16(3), 485–518.

Fedderke, J., J. Luiz, and R. de Kadt (2000): "Uneducating South Africa: The Failure to Address the 1910–1993 Legacy," *International Review of Education*, 46(3/4), 257–281.

Filmer, D., and N. Schady (2008): "Getting Girls into School: Evidence from a Scholarship Program in Cambodia," *Economic Development and Cultural Change*, 56(3), 581–617.

Fizbein, A., and N. Schady (2009): *Conditional Cash Transfers: Reducing Current and Future Poverty*. World Bank.

Hanes, S. (2006): "Easing the Burden of School Fees in Africa," Christian Science Monitor.

Hanushek, E. (2003): "The Failure of Input-Based Schooling Policies," *Economic Journal*, 113, 64–98.

Harding, R., and D. Stasavage (2014): "What Democracy Does (and Doesn't do) for Basic Services: School Fees, School Inputs, and African Elections," *Journal of Politics*, 76(1), 229–245.

Heckman, J., and J. Hotz (1989): "Choosing Among Alternative Nonexperimental Methods for Estimating the Impact of Social Programs: The Case of Manpower Training," *Journal of the American Statistical Association*, 84(408), 862–880.

Heckman, J., L. Lochner, and P. Todd (2006): "Earnings functions, rates of return and treatment effects: The Mincer equation and beyond," in *Handbook of the Economics of Education Volume 2*, ed. by E. Hanushek, and F. Welch. North-Holland.

Hidalgo, D., M. Onofa, H. Oosterbeek, and J. Ponce (2013): "Can Provision of Free School Uniforms Harm Attendance? Evidence from Ecuador," *Journal of Development Economics*, 103(3), 43–51.

Hoxby, C. (2001): "All School Finance Equalizations Are Not Created Equal," *Quarterly Journal of Economics*, 116(4), 1189–1231.

——— (ed.) (2004): *College Choices: The Economics of Where to Go, When to Go, and How to Pay For It*. National Bureau of Economic Research.

Hyman, J. (2016): "Does Money Matter in the Long Run? Effects of School Spending on Educational Attainment," Working paper, University of Connecticut.

Imbens, G., and K. Kalyanaraman (2012): "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *Review of Economic Studies*, 79(3), 933–959.

Imbens, G., and J. Wooldridge (2009): "Recent Developments in the Econometrics of Program Evaluation," *Journal of Economic Literature*, 47(1), 5–86.

Jackson, K., R. Johnson, and C. Persico (2016): "The Effects of School Spending on Educational and Economic Outcomes: Evidence from School Finance Reforms," *Quarterly Journal of Economics*, 131(1), 157–218.

Kane, T. (1994): "College Entry by Blacks since 1970: The Role of College Costs, Family Background, and the Returns to Education," *Journal of Political Economy*, 102(5), 878–911.

Knight, B., and N. Schiff (2016): "The Out of State Tuition Distortion," Working paper, Brown University.

Kremer, M., and A. Holla (2009): "Improving Education in the Developing World: What Have We Learned from Randomized Evaluations?," *Annual Review of Economics*, 1, 513–542.

Lam, D., C. Ardington, N. Branson, K. Goostrey, and M. Leibbrandt (2010): "Credit Constraints and the Racial Gap in Post-secondary Education in South Africa," University of Michigan.

Lam, D., C. Ardington, and M. Leibbrandt (2010): "Schooling as a Lottery: Racial Differences in School Advancement in Urban South Africa," *Journal of Development Economics*, 95(2), 121–136.

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

Lucas, A., and I. Mbiti (2012): "Access, Sorting and Achievement: The Short-Run Effects of Free Primary Education in Kenya," *American Economic Journal: Applied Economics*, 4(4), 226–253.

McCrary, J. (2008): "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 142(2), 698–714.

Pampallis, J. (2008): *School Fees*. Centre for Education Policy Development, Johannesburg, ZA.

Pritchett, L. (2013): *The Rebirth of Education: Schooling Ain't Learning*. Center for Global Development, Washington, DC.

Schultz, P. (2004): "School Subsidies for the Poor: Evaluating the Mexican PROGRESA Poverty Program," *Journal of Development Economics*, 74(1), 199–250.

Seekings, J., and N. Nattrass (2005): *Class, Race and Inequality in South Africa*. Yale University Press, New Haven, CT.

Todd, P., and K. Wolpin (2006): "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility," *American Economic Review*, 96(5), 1384–1417.

van den Berg, S., and M. Louw (2007): "Lessons Learnt From SACMEQII: South African Student Performance in Regional Context," Discussion Paper 16, Stellebosch University.

Wildeman, A. (2008): "Reviewing Eight Years of the Implementation of School Funding Norms," Idasa Economic Government Programme Research Paper.

Table 1: Treatment effects of fee elimination on school-level enrollment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Outcome	Fee Elimination	Enroll	Enroll	Fee Elimination	Δ Enroll	Δ Enroll	Fee Elimination	Enroll	Enroll	Fee Elimination	Δ Enroll	Δ Enroll
Estimator	OLS	OLS	IV	OLS	OLS	IV	OLS	OLS	IV	OLS	OLS	IV
High poverty	0.944***			0.941***	3.24*		0.943***		0.944***	1.51		
	(0.004)			(0.004)	(1.84)		(0.004)		(0.004)	(1.59)		
High poverty ×		3.57**						1.72				
post treatment		(1.81)						(1.63)				
Eliminated fees						3.44*						1.60
						(1.95)						(1.68)
Eliminated fees ×			3.78**						1.82			
post treatment			(1.92)						(1.72)			
Regression				×	×	×				×	×	×
Reweighting							×	×	×	×	×	×
Adjusted R2	0.901	0.049	0.047	0.905	0.052	0.052	0.904	0.000	-	0.908	0.047	0.047
# clusters	6181	6181	6181	6181	6181	6181	6181	6181	6181	6181	6181	6181
# observations	12362	24724	24724	12362	12362	12362	12362	24724	24724	12362	12362	12362

Notes: Columns 1, 4, 7 and 10 show the rates of fee elimination in high-poverty schools (quintile 2) relative to low-poverty schools (quintile 3). Columns 2, 5, 8 and 11 show the change in enrollment from 2005/6 to 2007/8 in high-poverty schools relative to low-poverty schools. Columns 3, 6, 9, and 12 show the change in enrollment from 2005/6 to 2007/8 in fee-eliminating schools relative to fee-charging schools with fee status instrumented by poverty level. Columns 1, 2, 3, 7, 8, and 9 are estimated in levels using data from both periods. Columns 4, 5, 6, 10, 11, and 12 are estimated in differences and control for baseline school characteristics using regression. Columns 7, 8, 9, 10, 11, and 12 are estimated using weighted least squares with weights assigned to low poverty schools to equate their distribution of baseline characteristics with those of the high poverty schools. Standard errors are shown in parentheses. In columns 1, 2, 3, 4, 5, and 6 the standard errors are estimated using a cluster-robust variance estimator allowing unrestricted error correlation within each school unit. In columns 7, 8, 9, 10, 11, and 12 the standard errors are constructed from 100 replications of a bootstrap algorithm that resamples school units and is stratifed by treatment group. *, **, and ***denote significance at the 10, 5, and 1% levels respectively.

Table 2: Treatment effects of fee elimination on enrollment in high- and very-high poverty schools (quintiles 2 and 1)

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome	Fee Elimination	Enroll	Enroll	Fee Elimination	Enroll	Enroll
Estimator	OLS	OLS	IV	OLS	OLS	IV
High poverty	0.949***			0.945***		
	(0.004)			(0.004)		
Very high poverty	0.995***			0.995***		
	(0.001)			(0.002)		
High poverty ×		3.41*			0.67	
post-treatment		(1.81)			(1.74)	
Very high poverty ×		7.65***			4.91***	
post-treatment		(1.73)			(1.66)	
Eliminated fees ×			6.13***			3.26**
post-treatment			(1.66)			(1.58)
Reweighting				×	×	×
Adjusted R2	0.930	0.059	0.058	0.957	0.001	0.001
# clusters	10235	10235	10235	10235	10235	10235
# observations	20470	40940	40940	20470	40940	40940

Notes: Columns 1 and 4 show the rates of fee elimination in high- and very high-poverty schools (quintiles 2 and 1) relative to low-poverty schools (quintile 3). Columns 2 and 5 show the change in enrollment from 2005/6 to 2007/8 in high- and very high-poverty schools relative to low-poverty schools. Columns 3 and 6 show the change in enrollment from 2005/6 to 2007/8 in fee-eliminating schools relative to fee-charging schools with fee status instrumented by indicators for high- and very high-poverty schools. Columns 4, 5, and 6 are estimated using weighted least squares with weights assigned to low poverty schools to equate their distribution of baseline characteristics with those of the high- and very high-poverty schools. Standard errors are shown in parentheses. In columns 1, 2 and 3 these are estimated using a cluster-robust variance estimator allowing unrestricted error correlation within each school unit. In columns 4, 5 and 6 these are constructed from 100 replications of a bootstrap algorithm that resamples school units and is stratifed by treatment group. *, **, and ***denote significance at the 10, 5, and 1% levels respectively.

Table 3: Treatment effects of fee elimination on enrollment in rural and urban schools

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome	Fee Elimination	Enroll	Enroll	Fee Elimination	Enroll	Enroll
Estimator	OLS	OLS	IV	OLS	OLS	IV
Sample	Urban schools			Rural schools		
High poverty	0.979***			0.944***		
	(0.006)			(0.005)		
Very high poverty	1.000***			0.997***		
	(0.000)			(0.001)		
High poverty ×		9.18**			-0.04	
post-treatment		(3.93)			(1.93)	
Very high poverty ×		7.50**			4.75**	
post-treatment		(3.83)			(1.86)	
Eliminated fees ×			8.40***			3.12*
post-treatment			(3.27)			(1.79)
Reweighting				×	×	×
Adjusted R2	0.982	0.050	0.049	0.902	0.016	0.017
# clusters	2791	2791	2791	7293	7293	7293
# observations	5582	11164	11164	20470	40940	40940

Notes: Columns 1 and 4 show the rates of fee elimination in high- and very high-poverty schools (quintiles 2 and 1) relative to low-poverty schools (quintile 3). Columns 2 and 5 show the change in enrollment from 2005/6 to 2007/8 in high- and very high-poverty schools relative to low-poverty schools. Columns 3 and 6 show the change in enrollment from 2005/6 to 2007/8 in fee-eliminating schools relative to fee-charging schools with fee status instrumented instrumented by indicators for high- and very high-poverty schools. Standard errors are shown in parentheses and are estimated using a cluster-robust variance estimator allowing unrestricted error correlation within each school unit. *, **, and ***denote significance at the 10, 5, and 1% levels respectively.

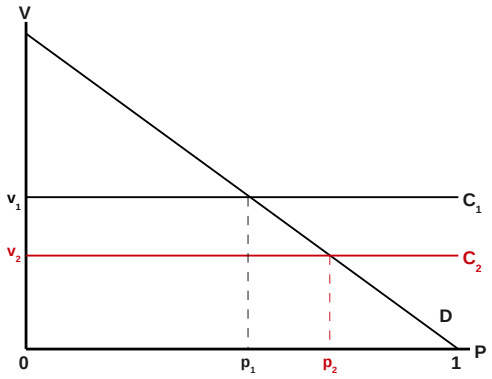
Table 4: Treatment effects of fee elimination on measures of school resources, grade progress, and composition

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Outcome	Dropout rate		Grade repetition		Class size		Proportion of student grant recipients		Proportion of student orphans	
Estimator	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV
High poverty ×	-0.28**		-0.44**		0.89***		1.66**		-0.98***	
post-treatment	(0.11)		(0.22)		(0.30)		(0.74)		(0.32)	
Very high poverty ×	-0.96***		-0.14		0.28		-1.10		-1.18***	
post-treatment	(0.10)		(0.20)		(0.25)		(0.72)		(0.31)	
Eliminated fees ×		-0.72***		-0.26		0.50**		-0.03		-1.13***
post-treatment		(0.09)		(0.19)		(0.23)		(0.65)		(0.28)
Pre-treatment mean	2.78		9.63		39.31		37.31		20.66	
Adjusted R2	0.003	0.002	0.003	0.002	0.006	0.006	0.003	0.003	0.002	0.002
# clusters	9857	9857	9856	9856	10147	10147	10235	40929	10131	36282
# observations	33399	33399	33396	33396	36986	36986	10235	40929	10131	36282

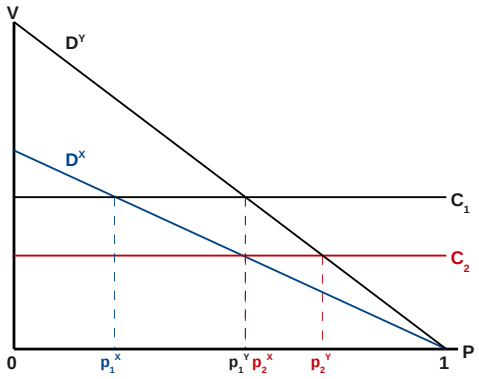
Notes: Columns 1, 3, 5, 7 and 9 and 4 show the change in the relevant outcome from 2005/6 to 2007/8 in high-and very high-poverty schools relative to low-poverty schools. Columns 2, 4, 6, 8 and 10 show the change in the relevant outcome from 2005/6 to 2007/8 in fee-eliminating schools relative to fee-charging schools, with fee status instrumented by indicators for high- and very high-poverty schools. Standard errors are shown in parentheses and are estimated using a cluster-robust variance estimator allowing unrestricted error correlation within each school unit. *, **, and ***denote significance at the 10, 5, and 1% levels respectively.

Figure 1: Conceptual framework showing the effect of fee elimination $C_1 \rightarrow C_2$ on the proportion of students enrolled P

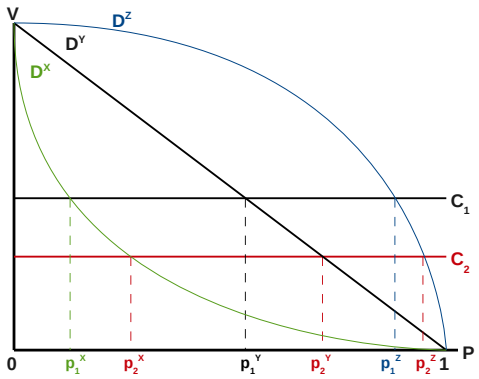
Panel A: Basic framework



Panel B: Different elasticities of demand

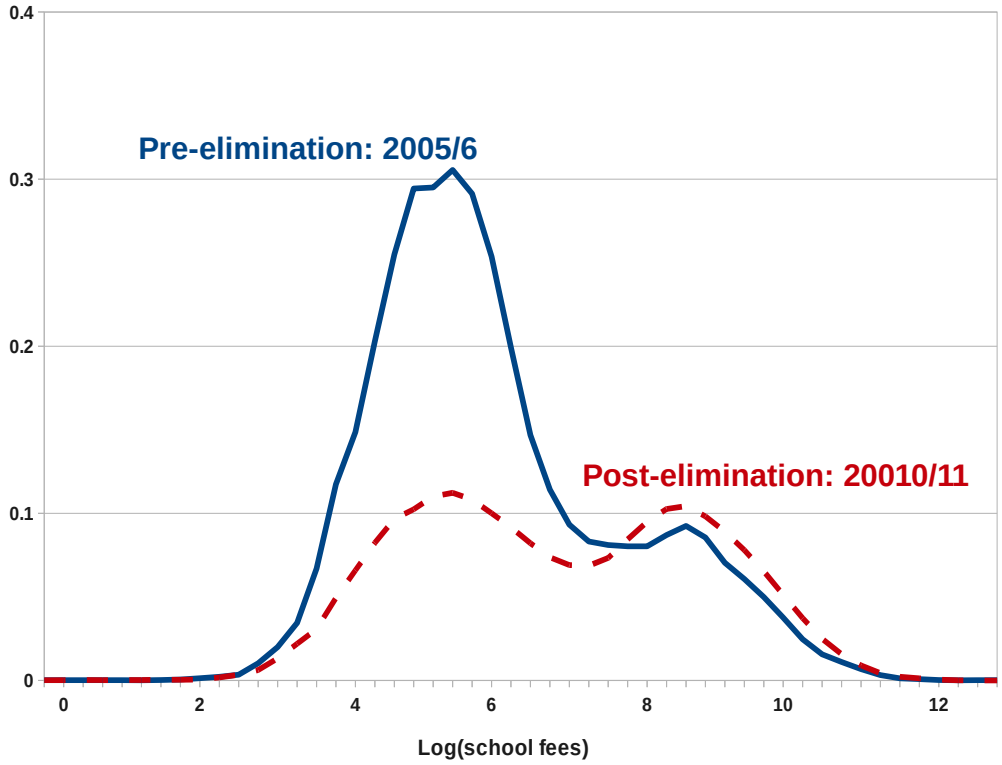


Panel C: Relating baseline valuation & enrollment changes



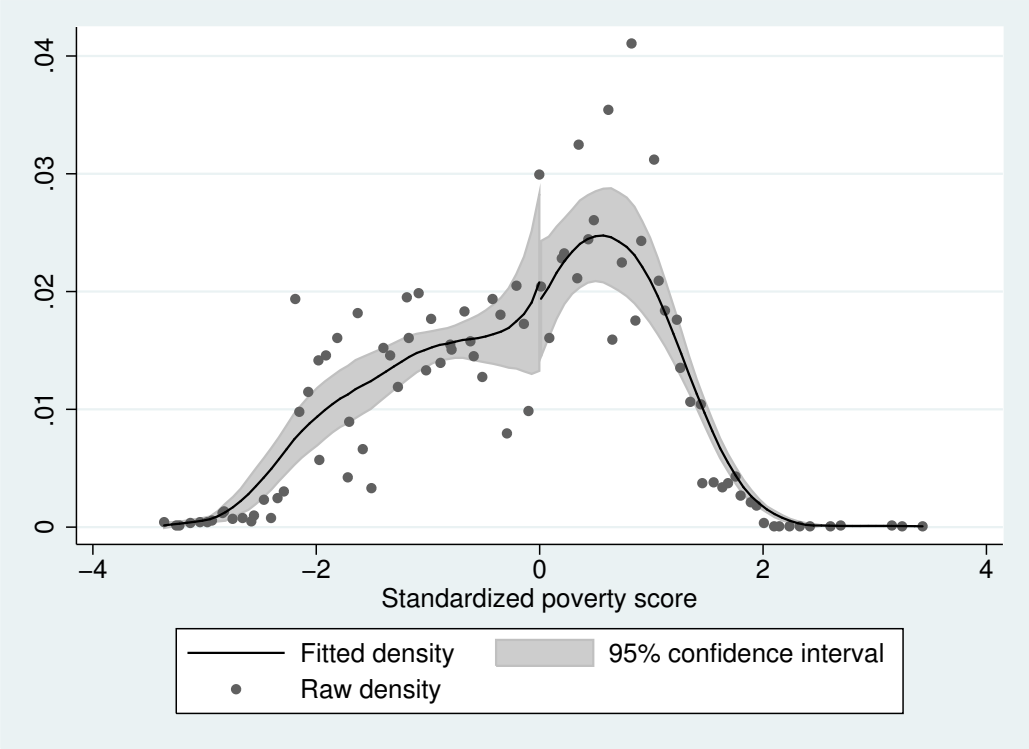
Notes: This figure illustrates the conceptual framework used in the paper. The proportion of students enrolled P is increased by eliminating school fees (panel A), the magnitude of the effect is larger when demand is more elastic (panel B), and the magnitude of the effect need not be correlated with mean baseline valuation (panel C).

Figure 2: Density of school fee payments by household before and after fee elimination



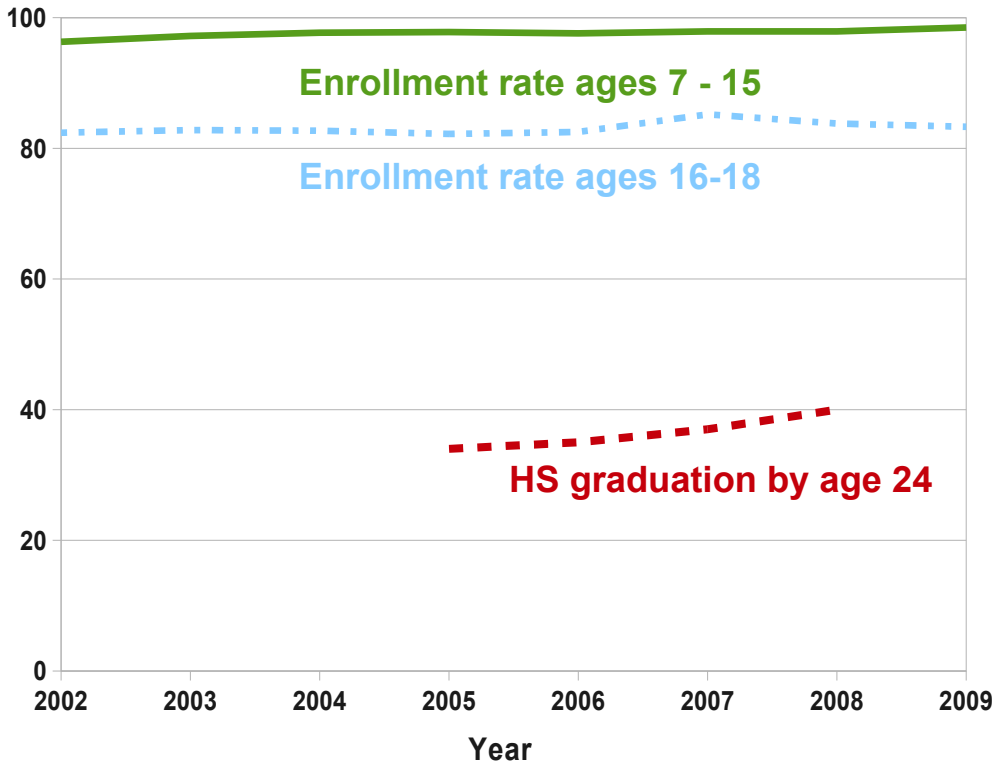
Notes: This figure shows the density of log school fees paid by households in the nationally representative Income and Expenditure Surveys from 2005/6 and 2010/1. Densities are calculated using survey weights and the vertical scales for the two densities are adjusted to reflect the higher fraction of households that report paying no fees in 2010/1 (52%) than 2005/6 (20%). School enrollment is not observed in these datasets, so households reporting zero expenditure on school fees may have no members enrolled in school.

Figure 3: Falsification test for a discontinuity in the density of standardized poverty scores



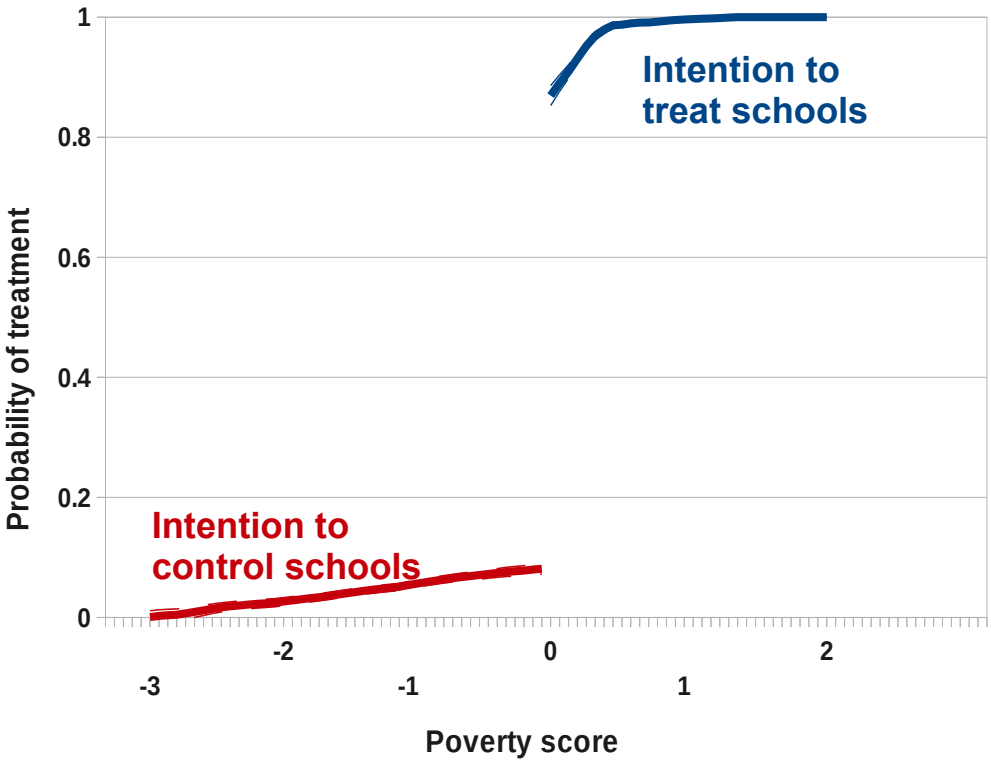
Notes: This figure estimates the density of poverty scores for schools in low- and high-poverty neighborhoods (left- and right-hand sections of the figure respectively). The difference between the densities evaluated at the cutoff between low- and high-poverty schools is -0.002 and is not statistically significant. This provides reassuring evidence that there was not systematic manipulation of poverty scores in order to control schools' assignment to (intended) treatment status. The densities are estimated separately on either side of the cutoff using local linear regression with a plug-in bandwidth selection. The results are robust to alternative bandwidth choices.

Figure 4: Time trend in net enrollment rate



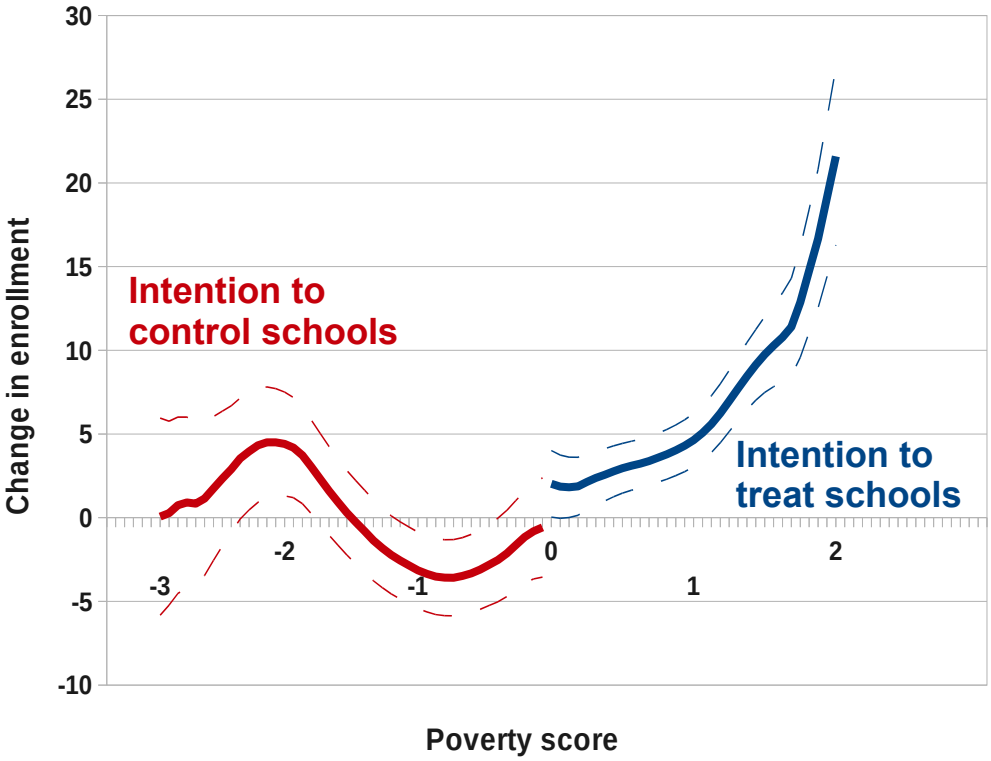
Notes: This figure shows the time series of the net enrollment rate. Data are calculated from the General Household Survey using appropriate sampling weights.

Figure 5: Probability that schools eliminate fees, by poverty score



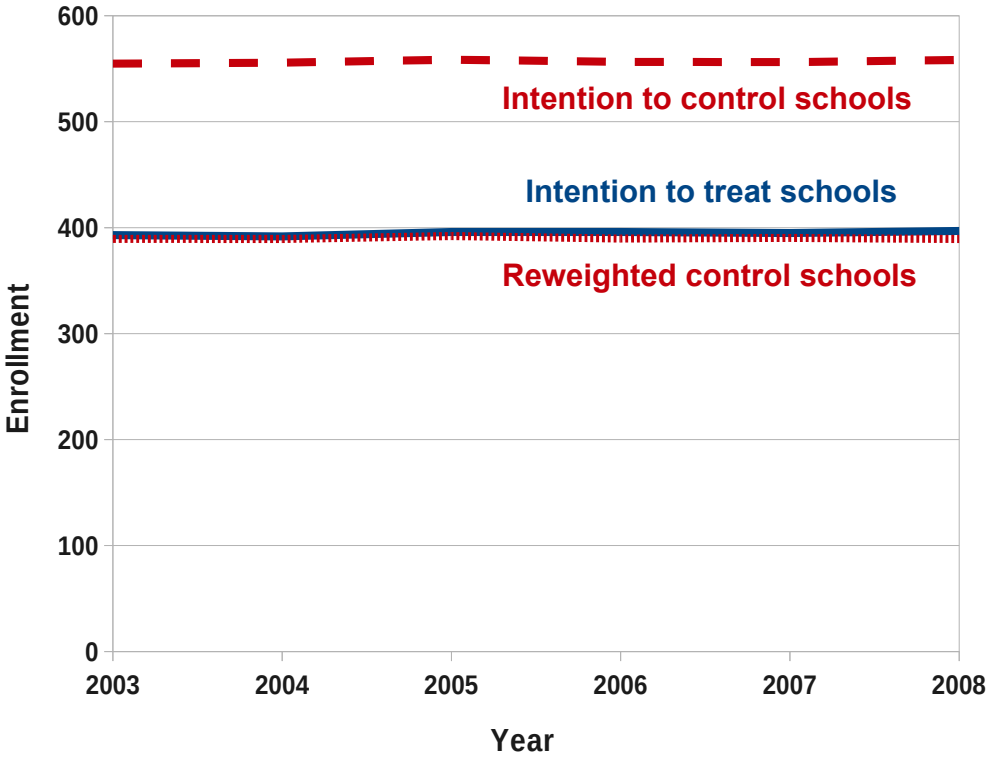
Notes: This figure shows the fitted probability that schools eliminate fees by their assigned poverty score. The fitted curves and 95% confidence intervals are from local linear regressions estimated separately on either side of the cutoff, with bandwidth choices following Imbens and Kalyanaraman (2012). The estimated difference between the curves at the threshold value is 76 percentage points with a standard error of 1 percentage point.

Figure 6: Change in enrollment from 2005/6 to 2007/8, by poverty score



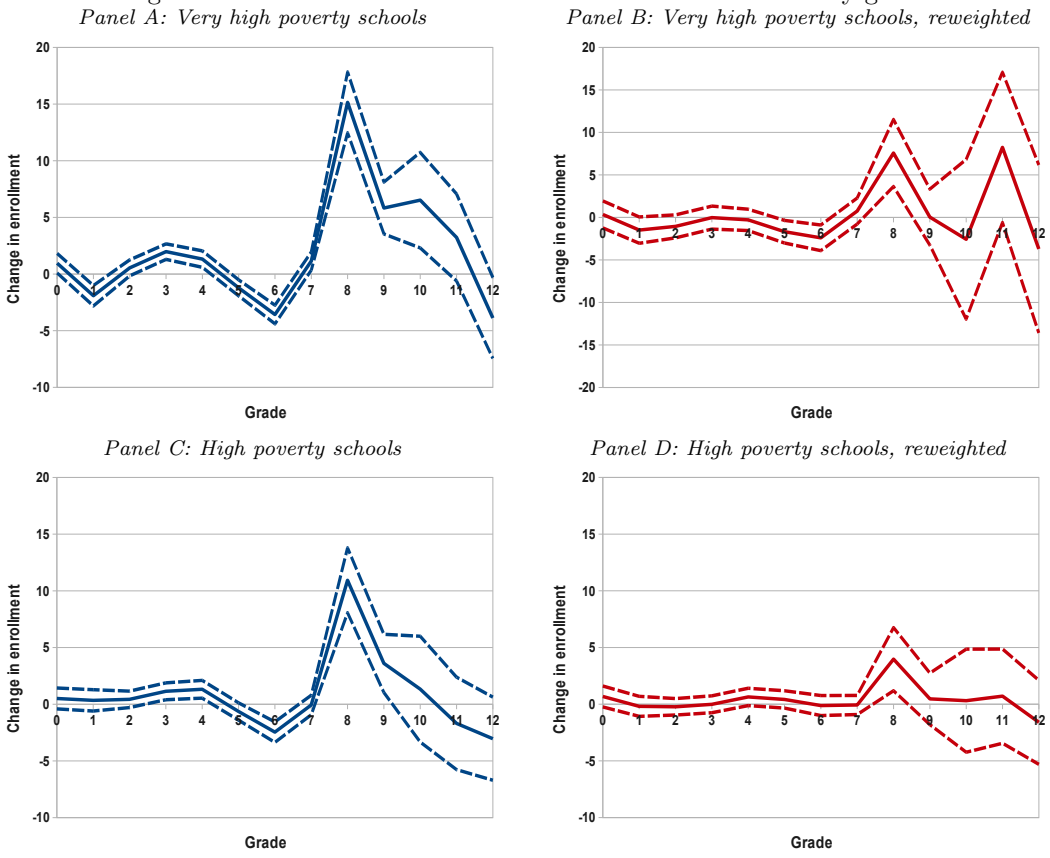
Notes: This figure shows the fitted change in school enrollment from 2005/6 to 2007/8 by schools' assigned poverty scores. The fitted curves and 95% confidence intervals are from local linear regressions estimated separately on either side of the cutoff, with bandwidth choices following Imbens and Kalyanaraman (2012). The estimated difference between the curves at the threshold value is 2.8 students with a standard error of 1.8 students.

Figure 7: Level of enrollment from 2003 to 2008 for intention to treat schools, control schools, and reweighted control schools



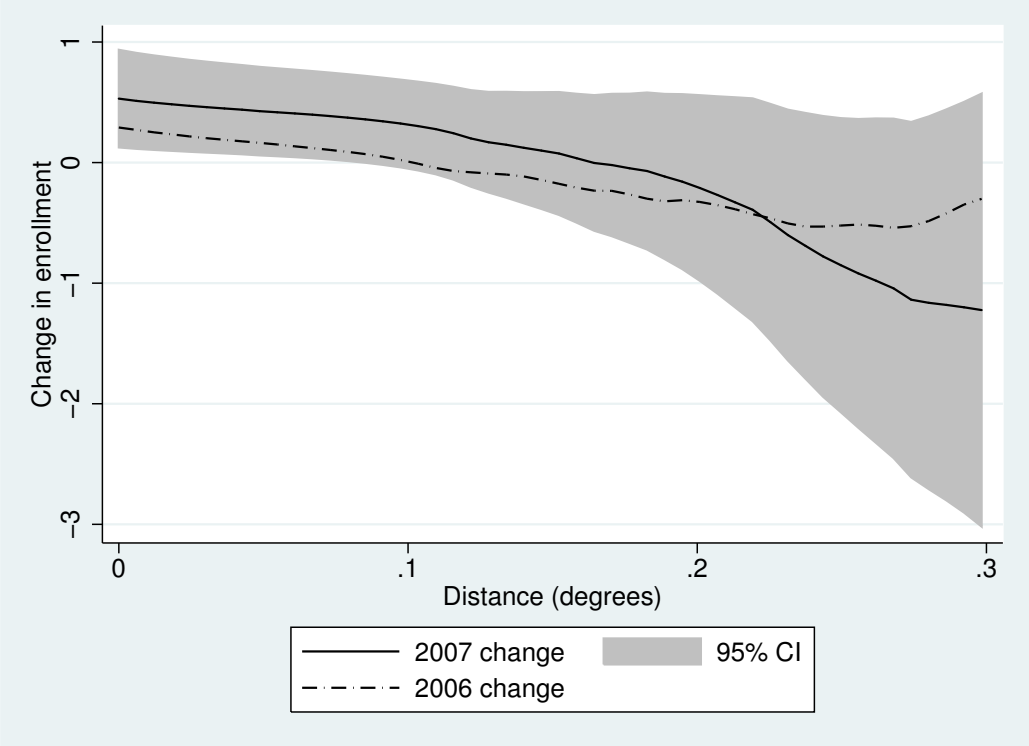
Notes: This figure shows the level of enrollment for different groups of schools in each year from 2003 to 2008. The solid blue line shows intention to treat (quintile 2) schools, the dashed red line shows intention to control (quintile 3) schools, and the dotted red line shows the latter group of schools after reweighting them to have the same distribution of baseline observed characteristics as the intention to treat schools.

Figure 8: Treatment effects of fee elimination on enrollment by grade



Notes: This figure shows grade-specific treatment effects of fee elimination on enrollment. Panels A and B show results for schools in very high poverty neighborhoods (quintile 1) while panels C and D show results for schools in less high poverty neighborhoods (quintile 2). Panels B and D show results after reweighting the control schools (quintile 3) to have the same distribution of baseline observed characteristics. These are intention-to-treat effects. The compliance rate does not vary substantially by grade so the overall pattern of instrumental variables effects is similar. The 95% confidence intervals in panels A and C are obtained from a school-level cluster robust variance estimator. The intervals in panels B and D are obtained from 100 replications of a bootstrap that resamples schools and stratifies by quintile.

Figure 9: Changes in enrollment from 2005 to 2006 and 2006 to 2007 for fee-charging schools located at different distances from fee-eliminating schools



Notes: This figure uses local linear regression to estimate the relationship between the change in enrollment from 2006 to 2007 and proximity to a fee-eliminating school (solid line) for the sample of fee-charging schools. It estimates the same relationship for the change from 2005 to 2006 (dashed line). A strong negative slope for the former relationship relative to the latter relationship would suggest that fee-charging schools near fee-eliminating schools are losing students to those treated schools. The results strongly reject this hypothesis and are robust to alternative bandwidth choices.