The Effect of Nonpersonnel Resources on Educational Outcomes: Evidence from South Africa

MIQUEL PELLICER
University of Duisburg-Essen and Southern Africa Labour and Development Research Unit

PATRIZIO PIRAINO
University of Cape Town

I. Introduction

A substantial body of literature in economics and education investigates the impact of different types of educational inputs on student learning in developing countries (for reviews, see Hanushek 2006; Glewwe et al. 2013). This is a literature of obvious policy relevance, as most countries in the developing world allocate a significant share of total government expenditure to education financing (UNESCO 2014). While much attention in the literature has been devoted to evaluating the role played by personnel inputs in the education production function (e.g., teacher preparedness and absenteeism, pupil-teacher ratio), numerous studies have also attempted to estimate the impact of nonpersonnel resources on educational outcomes.

Nonpersonnel inputs in the education production function typically include pedagogical materials, basic furniture, electronic equipment, as well as basic school infrastructure. Recent reviews of the evidence from developing countries show that while some nonpersonnel inputs appear to have an unambiguous positive effect on student learning (e.g., basic school infrastructure and furniture), the body of evidence on several nonpersonnel resources—including a variety of pedagogical materials—remains too low for any strong conclusions (Glewwe et al. 2013; Snilstveit et al. 2016). This is especially true when limiting the analysis of the evidence to randomized evaluations and studies that use quasi-experimental methods. These findings highlight the dif-

We acknowledge funding from the European Union's Seventh Framework Programme through the NOPOOR project Enhancing Knowledge for Renewed Policies against Poverty. We thank participants at the Southern Africa Labour and Development Research Unit seminar at the University of Cape Town and at the Research Programme 3 German Institute of Global and Area Studies seminar for useful comments. We are grateful to Rob Garlick for detailed and very useful feedback on the manuscript. All errors are ours. Contact the corresponding author, Miquel Pellicer, at pellicer.miquel@gmail.com.

Electronically published May 3, 2019
 0013-0079/2019/6704-0006$10.00

ficulty of gaining precise knowledge of the education production function and of the mechanisms through which nonpersonnel resources can improve student learning in developing countries.

A type of input we know little about is the role of financial resources transferred to schools for nonpersonnel expenses. Yet in the context of increasing decentralization, this input may be of particular relevance. Typically, governments in developing countries do not directly allocate specific nonpersonnel materials to schools. Rather, schools are increasingly offered grants on a per-pupil basis and left to choose how to spend these recurring resources (UNESCO 2005; Weidman and DePietro-Jurand 2009). From a policy point of view, it is thus imperative to understand the effectiveness of this type of funding allocation. Our paper contributes toward filling this gap by examining the effects of resources allocated to schools for recurrent expenses on equipment and consumables used to ensure the proper functioning of the school (e.g., textbooks, stationery, computers). In our setting, these funds are allocated on a per-pupil basis and are separate from infrastructural investments (e.g., the building of schools, classrooms, and the provision of water and sewage).

Empirical evidence on the role of general nonpersonnel per-pupil funding in developing countries is scarce. Besides some cross-country studies that are likely to suffer from confounding factors, only two articles are cited by Glewwe et al. (2013) as dealing with the effect of general per-pupil expenditures on school outcomes: Nannyonjo (2007) and Du and Hu (2008). However, the results from these studies are difficult to interpret, as the estimated effects are based on multivariate regressions that cannot control for potentially relevant unobserved factors.

A credible identification strategy is essential in this context because resources may be endogenously allocated to schools (e.g., via success in attracting funds or via government redistributive motives), and this would confound estimates of their effect on outcomes. A recent study by Das et al. (2013), which is not included in the review by Glewwe et al. (2013), does credibly identify the effect of material nonpersonnel resources on test scores. The authors report the results from a randomized experiment in India where schools were assigned a block grant for nonpersonnel expenditures. They find a positive effect on test scores. While extremely valuable, these results are based on an intervention with some unique characteristics. The block grant was restricted to items to be used directly by students. Moreover, the method of disbursal required schools to make a list of desired items and had the implementing organization (together with the teacher) purchase the items. These requirements may have prevented funds from being spent in grossly inefficient ways.

This paper adds causal evidence of the impact of nonpersonnel spending on educational outcomes in developing countries on the basis of actual policy for allocating such resources. Evaluations based on existing government programs may offer a more realistic setting for studying the effect of resources on student learning in the context of a developing economy, in this case South Africa. In addition, South Africa represents a particularly interesting case to study with regard to school resources and performance. Inequalities in both variables are massive in the country because of the historical legacy of apartheid.[1] Understanding the effect of resources on school outcomes will thus help assess the impact of progressivity in funding allocations on existing performance gaps.

The peculiar manner in which these resources have been allocated in South Africa since 2007 provides us with what we believe to be a credible identification strategy. Schools are assigned a poverty score that depends on the socioeconomic characteristics of the surrounding community and are divided into quintiles on the basis of this score. Nonpersonnel funding is then determined by the school quintile. Schools in the poorest quintile receive around 800 rands (~80 USD) per pupil, while schools in the remaining quintiles obtain progressively less funding, with schools at the top receiving about 150 rands per pupil. Funding is therefore discontinuous in the poverty score at the thresholds determining the quintiles, and we use a regression discontinuity approach to estimate the effects of resources. Our main outcome variables are attendance and success in the national test that students take at the end of high school (matric).

We use school-level data for five of South Africa's nine provinces. We have information on poverty scores, quintiles, enrollment by grade, as well as matric attendance and results. Our descriptive analysis shows that differences in government funding across the bottom three quintiles of the poverty score distribution are rather small. This leads us to exclude two of the four discontinuities we could have potentially exploited. Using more detailed data from one province, we also show that fees in quintile 5 schools are high enough to render changes in government funding at the 4–5 quintile threshold largely irrelevant for the schools concerned. This leads us to rely exclusively on the cut-off from quintile 3 to quintile 4 for identification purposes.

Importantly, data in all five provinces were collected both before and after the current funding system, which was established in 2006–7. This allows us to use information before the policy change to verify the validity of our identification strategy as well as to control for prepolicy outcomes at the school

[1] Since the democratic transition in the first half of the 1990s, some of these inequalities have started to recede, although they remain very large. For instance, our data show that in formerly white-only schools, 54% of students obtain university access pass rates, which compares with only 15% in formerly African black schools.

level. We show that there is no jump in pretreatment outcomes at the 3–4 threshold. We also show that the number of teachers is balanced at either side of the threshold during the posttreatment period, suggesting that resources are used for nonpersonnel expenditures as intended. This is reassuring for the validity of our key identification assumptions. Moreover, we use additional data on fees in the province of Gauteng to check whether schools may compensate lower transfers from the government with higher fees. We find no evidence of fee compensation at the cut-off from quintile 3 to quintile 4. This provides further confidence on the identification strategy.

Our empirical analysis shows that nonpersonnel school resources have a positive effect on student throughput from grades 10–12.[2] On the other hand, we find no evidence of a positive effect on matric pass rates. If anything, the effects on matric pass rates are somewhat negative. These findings suggest that the additional resources may have the effect of increasing the retention of academically weaker students. We speculate this might be a result of the per-pupil nature of funding, which may create an incentive to keep students past the compulsory level, with no effect on student learning (as measured by pass rates on the final examination). That is, it is possible that schools may respond to additional funding by altering their behavior on the incentivized margin (enrollment) but not on the nonincentivized margin (matric pass).

It is important to note, however, that our results apply to a specific type of schools, notably secondary schools just above the median of the school poverty distribution. It is plausible that resources might have a different (and positive) effect in other contexts. Notably, resources might be more productive in poorer schools (if decreasing returns are important) and/or in primary schools, where students are younger.

The rest of the paper is organized as follows. Section II explains the allocation of nonpersonnel funds across South African schools. Section III describes our data. Section IV presents our empirical approach, while Section V evaluates its strengths and limitations in our context. Section VI presents the results. Section VII provides a discussion and concludes.

II. Allocation of Nonpersonnel Resources in South Africa

Under the apartheid system, school allocations were highly regressive, with white students receiving substantially more funding than black students (Branson, Kekana, and Lam 2013). After the democratic transition, the eradication of these inequalities became a priority. Salaries of black and white teachers were

[2] These are the postcompulsory grades in South Africa. Schooling is mandatory until grade 9 or age 15 (whichever comes first).

equalized. A progressive system for funding nonpersonnel expenditures was set up in 2000. These funds were meant to finance running costs of the schools, such as teaching materials, small equipment, and small repairs.[3]

The system allowed a high level of discretion to the provinces. Most of the funding followed what was denoted as the quintile system: each province constructed a poverty score that ordered schools from poorest to richest on the basis of the surrounding community and the characteristics of the school itself. The quintiles from this score determined the allocation that each school obtained. Funds assigned to each quintile varied significantly by province as well as the amount of nonpersonnel resources channeled through the system (Department of Basic Education, Republic of South Africa 2003).

In 2006 the scheme changed. It was observed that poorer provinces were allocating fewer resources to education than richer ones, and the resulting amounts were considered not sufficient for certain schools. Therefore, the allocation rule became more centralized. A quintile system based on poverty scores still remained the basis for funding, but the quintiles were to be at the national and not the provincial level. As a consequence, the poverty scores were recomputed in a more homogeneous way across provinces. Importantly, the new scores were not allowed to account for school characteristics to avoid perverse incentives. The variables that formed the base for the poverty score were to be the same across all provinces; these included income, unemployment, and the level of education/literacy in the area. The geographic unit of analysis was to be the electoral ward, although some provinces appear to have used smaller units.[4] Note that the provincial departments were given discretion on the weighting to apply to the various socioeconomic indicators. This—in addition to the fact that each province used a different scale—makes poverty scores not directly comparable across provinces. As detailed in Section IV, poverty scores are interacted with provincial dummies when fitting the smooth function of the running variable in the regression discontinuity model to account for these different scales.

The reform of school funding also incorporated a change of the fee policy. Policy makers wanted to allow free access to school for poorer learners, and the quintile system used for school funding was also used as a basis for fee waiving. Schools in quintiles 1 and 2 were declared no-fee schools and therefore were

[3] As noted in the introduction, there exist other funds in South Africa for infrastructure upgrading (e.g., new classrooms or facilities). The type of funds on which our paper focuses should therefore be thought of as both nonpersonnel and noninfrastructural.

[4] Wards are not administrative units in South Africa; the level of aggregation was dictated by the information available in the census (Garlick 2013). The national Department of Education provided each province with socioeconomic data from the 2001 census.

not allowed to charge fees to students. The higher amount of resources these quintiles received by virtue of the funding system was considered enough to compensate for the lost fees. Starting in 2008, schools in quintile 3 also began to shift to a no-fee regime on a rollout basis that varied by province.

For any given year, the government releases guidelines for funding to each quintile and the provinces decide on the actual disbursements. Wildeman (2008) indicates that some provinces deviated from the guidelines, although the differences do not appear to be significant. Figure 1 shows the evolution of funding for each quintile over time (in constant 2008 rands). Quintiles 1 and 2 received around 700 rands per learner, while the figures for quintile 4 and 5 are around 400 and 150 rands. Values for all quintiles are relatively constant except for quintile 3, which increases from 2007 to 2010 from around 500 to 700 rands. The change for quintile 3 reflects the change in the fee policy, with quintile 3 schools shifting to a no-fee regime over time. The figure also makes clear that the difference in funding between quintiles 1 and 2 and (from 2009 onward)

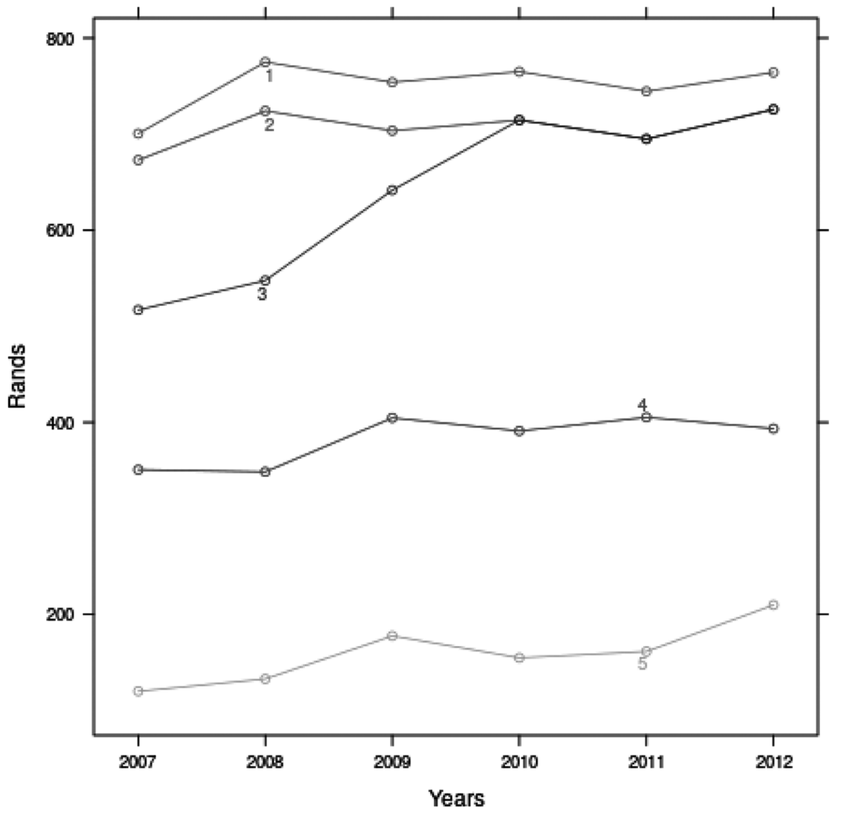

Figure 1. Evolution of per pupil funding by quintile. A color version of this figure is available online.

between quintiles 2 and 3 is relatively small, a point that will be relevant for our analysis.

III. Data

We use school-level administrative data on five of South Africa's nine provinces: Western Cape, Gauteng, Northern Cape, Eastern Cape, and KwaZulu-Natal. These were the provinces from which we were able to obtain school poverty scores. The data come from different sources and have been merged by the authors. Poverty scores come from either the provincial Department of Education or from provincial Gazettes. They refer to the years 2005 (Western Cape), 2006 (Eastern Cape), 2007 (Gauteng), and 2008 (KwaZulu-Natal and Northern Cape). These data have been merged with two administrative and freely available data sets on basic school characteristics. The first data set is the Education Management Information System of the national Department of Basic Education, which provides data on school quintiles for different years and on the former apartheid era school categorization (e.g., whether the school belonged to a white area or an African black township; http://www.education.gov.za/EMIS/tabid/57/Default.aspx). The second data set is the SNAP survey, a school-level survey with information on enrollment levels by grade for different years (https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/482). In addition, these data sets have been merged with information on matric results obtained from the Department of Basic Education. These include the number of students writing the exam and those passing for each school. All variables are available from 2008–12. We also have enrollment data from 2003–5 and matric information for 2005. Our analytical sample includes only schools for which we have information on all years, which are a large majority of them (83%).

For one province (Gauteng), we obtained additional data from the provincial Department of Education. The additional information for Gauteng includes the government nonpersonnel funding actually received by the school and fee amounts at the school level. Moreover, the data from Gauteng contain information on the number of computers in the school as well as on the race composition of the student body. We use these variables as imperfect proxies of material resources owned by the school and students' socioeconomic status, respectively. Unfortunately, the available information is not constant across all years, so we use different subsamples in alternative estimations. In particular, fee information is available from 2008–12, with the exception of 2010. Information on the number of computers is available only for 2010–12.

For the general data set with all provinces, our funding variable is the amount of funds to be assigned to each school—on the basis of its quintile and province—as stated in the provincial official documents. It is important to note that this is based on the allocations that schools are supposed to obtain on the basis of the quintile reported in our data, not on the money they actually obtain. There are two possible sources of measurement error on this variable. First, the provincial documents and/or our quintile data could contain mistakes. In order to check for this possibility, we use the Gauteng data on the allocations to individual schools, which we have for 2008. The correlation between the official Gauteng record and our funding variable is 0.97, which suggests that measurement error is not likely to be problematic in this regard. A second potential source of error may result from leakages in government disbursements, which would imply that schools receive less funding than they are entitled to. Reinikka and Svensson (2004) provide an example of this problem in Uganda. However, South Africa has a significantly better record of accountability in government expenses compared with other countries in the region. Also, since the information on the amount schools should receive is publicly available, this would not appear to be a major concern in our setting.

We take as treatment for time t the average funding from the 3 years before t. This is because funding for learning is cumulative, and funding at $t - 1$ ought to contribute to learning as much as funding at t.[5] The funding variable is expressed in hundreds of rands. This is the order of magnitude of the observed jumps in funding across quintiles.

Regarding our outcome variables, we focus on matric pass rates as well as on student throughput. The simple pass rate—that is, the number of student passing the exam over those writing it—may not be the best measure of school performance. It is in fact possible that schools discourage low-achieving students from writing the test as a result of the emphasis placed on matric pass rates in a variety of accountability processes.[6] For this reason, we also use (and tend to prefer as a measure of student learning) the number of students passing the examination relative to the number of pupils in grade 10 from 2 years earlier. The other outcome variables we use are the student throughput from grades 10–12 and the ratio of students writing the matric exam relative to both grades 12 and 10 from 2 years earlier.

Table 1 shows the descriptive statistics of our sample by quintile. We have 500–700 high schools per quintile.[7] As expected, the poorest schools are in

[5] The results are not affected by using funding for time t only.

[6] Borkum (2012) also notes the possibility of "conscious gate keeping whereby schools block the progression of weak students at grade 10 in order to avoid having them eventually take the matric exams in grade 12" (376).

[7] The uneven number of schools across quintiles results from the under- and overrepresentation of the five provinces in our sample in different quintiles.

TABLE 1
DESCRIPTIVE STATISTICS BY QUINTILE

Variables	Q1	Q2	Q3	Q4	Q5
Size:					
Number of schools	550.00	536.00	754.00	520.00	446.00
Apartheid department:					
Townships	.10	.16	.32	.30	.04
Coloured	.02	.02	.06	.25	.14
Homeland	.68	.64	.48	.18	.02
Indian	.00	.00	.01	.09	.18
Natal	.03	.04	.03	.04	.03
New	.14	.09	.07	.02	.01
White	.01	.01	.02	.12	.57
Province:					
Eastern Cape	.31	.30	.39	.17	.16
Gauteng	.04	.08	.13	.26	.39
KwaZulu-Natal	.61	.56	.39	.38	.32
Northern Cape	.01	.02	.02	.03	.01
Western Cape	.03	.04	.06	.17	.11
Pretreatment:					
log grade 10	4.80	4.97	5.09	5.16	5.14
Throughput grades 10–12	.59	.58	.58	.57	.65
Matric attendance rate	.91	.92	.93	.94	.97
Matric pass rate	.61	.62	.61	.66	.83
Matric attendance rate over grade 10	.50	.49	.50	.52	.59
Matric pass rate over grade 10	.29	.28	.29	.33	.51
Treatment:					
Government funding (hundreds of rands)	7.47	6.97	6.57	3.84	1.60
Posttreatment:					
log grade 10	4.73	4.75	5.01	5.31	5.20
Throughput grades 10–12	.58	.57	.56	.53	.72
Matric attendance rate	.85	.87	.89	.91	.95
Matric pass rate	.59	.62	.62	.74	.89
Matric attendance rate over grade 10	.49	.49	.49	.49	.69
Matric pass rate over grade 10	.28	.29	.29	.37	.63

Note. Apartheid department refers to the government department managing different types of schools under apartheid. Data in these rows refer to the distribution of schools in different departments for a given quintile. Failure of the rows to sum up to 1 comes from missing values. Data for treatment-related variables correspond to the average in each quintile. Pretreatment variables refer to either 2005 (for matric variables) or to the average between 2003 and 2005 (for enrollment-related variables). Posttreatment variables refer to the average between 2008 and 2012.

the provinces of the Eastern Cape and KwaZulu-Natal, many of which belonged to the homelands under the apartheid regime. At the opposite extreme, in quintile 5, schools are predominantly white (and, to a lesser extent, Indian and Coloured), according to the old classification system. Quintiles 3 and 4 are more mixed, containing a significant number of township schools and with a more even distribution across all provinces (except Northern Cape, which is sparsely populated and has fewer schools in any quintile).

Table 1 also shows the treatment values as well as pre- and posttreatment outcomes. The treatment (government funding averaged between 2008 and

2012) indeed decreases with the quintile, showing small differences between quintiles 1 and 3 and larger jumps between quantiles 3 and 5. Almost all pre- and posttreatment variables (notably pass rates) increase by quintile, suggesting that poverty scores do reflect factors harmful for school progression and learning. Most patterns are clearly convex in quintile, with a particularly sharp difference between quintile 5 and the rest. This is consistent with expectations, reflecting the historical legacy of discrimination in South Africa, and is much in line with a variety of other socioeconomic outcomes (Leibbrandt et al. 2010). The throughput and pass rates for all quintiles (except the fifth) imply that during the period under analysis, less than one-third of students enrolled in grade 10 successfully completed matric 2 years later.

IV. Empirical Approach: Regression Discontinuity Model

We use a regression discontinuity approach to estimate the effect of school resources on educational outcomes. Since funding jumps at the poverty score thresholds that determine the quintiles, indicator functions of such jumps should deliver the local causal effect of funding once we control for a smooth function of the poverty score. In our data, the quintiles that determine funding do not perfectly follow the poverty score; as time passes, a number of schools are assigned to different (typically lower) quintiles. The design is therefore fuzzy and leads us to adopt a standard instrumental variables approach using the original quintile assignment as instrument for resources. In our analysis, outcome and endogenous variables are averaged over the years 2008–12.

More formally, consider an outcome Y_{ij} of school i in province j. Denote the funding received by each school as F_{ij} and the poverty score associated to each school as s_{ij}. Because poverty scores differ by province, we will fit the smooth function of the poverty score separately by province, interacting the score with provincial dummies, denoted by P_j for province j. The equation for the second stage is thus

$$Y_{ij} = P_j + P_j g_j\left(s_{ij}\right) + \rho F_{ij} + \beta' z_{ij} + u_{ij}, \qquad (1)$$

where g_j is a flexible function of the poverty score that is allowed to differ by province and z_{ij} is a vector of control variables. Notice that ρ, which is our coefficient of interest, is assumed not to differ by province, so that we obtain one coefficient that aggregates the effects across the different provinces.

The term F_{ij} might be correlated with u_{ij} because schools may succeed in attracting funds depending on political connections or general characteristics themselves linked with the performance of the school. For this reason, we instrument F_{ij} with an indicator function that captures the original quintile assignment on the basis of the poverty score. If poverty scores were constructed using technical considerations, as argued above, the initial assignment (and therefore government funding) should be random around the threshold that determines quintile jumps. For reasons to be detailed below, our analysis is based on just one discontinuity from quintile 3 to quintile 4, and we focus on that discontinuity alone here. The indicator function taking a value of 1 for schools assigned to quintile 3 or lower (as opposed to quintile 4 or higher) is denoted by $F34_{ij}$. Our first stage is thus

$$F_{ij} = P_j + P_j h_j(s_{ij}) + \rho F34_{ij} + \gamma' z_{ij} + v_{ij}. \tag{2}$$

We estimate these equations by two-stage least squares. If no variable relevant for the outcome Y jumps at the threshold defining the original quintile assignment $F34_{ij}$, the error term u_{ij} will be uncorrelated with the predictions of the first stage and ρ will be consistently estimated. Because the poverty score is computed on the basis of geographical features, it is spatially correlated and so is likely to be the outcome we consider. We thus cluster our standard errors at the magisterial district level, a unit larger than the ward that typically served as the basis for the computation of poverty scores. There are 239 magisterial districts in our sample containing an average of 13 schools.

We estimate different specifications of these equations using alternative windows around the threshold for the jump from quintile 3 to quintile 4. The poverty score is centered around this threshold and standardized by dividing it by its standard deviation. The specifications we use are a polynomial of degree 3 for $g_j(s_{ij})$ with a window of ± 1.5 standard deviations and a linear specification with a smaller window (± 0.5) that allows for different slopes at either side of the threshold. Our benchmark specifications control for the apartheid period school classification and indicators for quintile jumps other than the one of interest from quintile 3 to quintile 4.[8] We also perform analyses with and without pre-2007 outcomes as controls.

An important challenge of the analysis is that we are interested in the effect of nonpersonnel school resources but we observe only funding from the government. Specifically, we do not observe the fees charged by the school. This may be problematic for a few reasons. First, if government funding is only a

[8] Some schools have poverty scores that assign them to quintiles other than 3 and 4 in the specifications with a large window. For instance, schools in Eastern Cape with a poverty score above 0.975 (and below 1.71) are assigned to quintile 2. This implies that resources might jump at these other thresholds as well. We thus include indicators for such quintile assignment in all our regressions. This appears to matter only for the first stage, where ignoring the jump in resources at other thresholds makes a difference for estimates of the 3–4 jump. The instrumental variables results are virtually the same, regardless of whether these controls are included.

minor component of school resources, differences in this type of funding will be practically insignificant. Second, schools may compensate lower government funding with higher fees. If this were the case, schools at either side of the quintile threshold would have smaller or no differences in resources compared with government funding alone, and this would bias our estimates toward zero. In addition, differences in fees at either side of the threshold could threaten the validity of our approach, since they would lead to a violation of the key assumption that $F34_{ij}$ is uncorrelated with the error term u_{ij} after controlling for a smooth function of the poverty score. In particular, differences in fees could imply that schools at either side of the threshold are different to start with, possibly through changes in the pool of students. We address these and other identification threats in the following section.

V. Validity of the Approach

This section addresses potential concerns with our empirical approach before proceeding to the results. We start by addressing the concerns associated with not observing overall school resources in our data. We do this by providing evidence on schools in the Gauteng province for which we have additional information. We then use our main data set from all provinces to assess the first stage of our model and to focus on the basic regression discontinuity assumptions of no manipulation of the running variable and of no jump at the relevant threshold exhibited by predetermined variables.

A. Fees and Government Funding

As mentioned above, an important challenge to our approach is that we cannot observe overall school resources, only government funding. To explore the implications of this limitation for our analysis, we use data from the Gauteng province. These data include information on fees and material resources, such as computers. We assess whether jumps in government funding are economically significant and whether it appears that fees are used to compensate for the discrete jump in funding. Since these issues apply to all schools and not only to secondary schools, we increase our sample size by using all schools in the province, although we also report results using secondary schools only.

Figure 2 shows average school fees by poverty score in Gauteng. The left panel includes all schools, while the right panel considers only secondary schools. Each data point corresponds to the average from 2008–12 (excluding 2010, for which no data are available) for the schools with a given poverty score. On average, each poverty score includes around 40 schools in the left panel and 12 high schools in the right panel. Vertical lines show the different

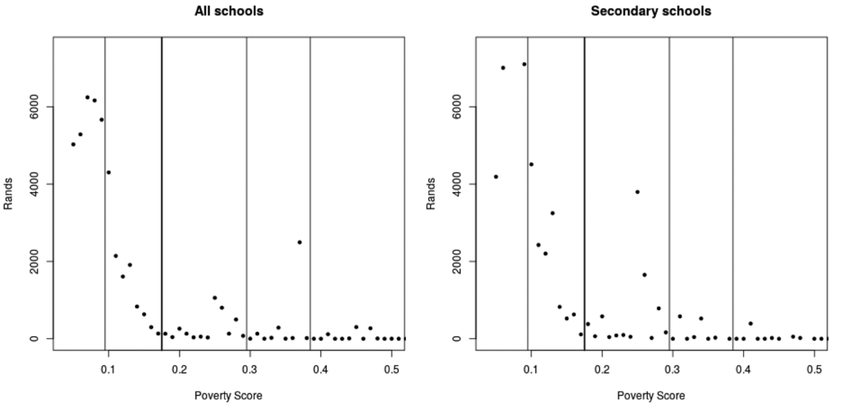

Figure 2. Average fees by poverty score, Gauteng province.

quintile cutoffs. Richer schools (in higher quintiles) appear at low levels of the poverty score on the left side of the distributions.

There are two important messages from figure 2. First, fees at richer schools are so large that they dwarf any nonpersonnel government funding. While pupils in quintile 5 schools pay around 6,000 rands (~600 USD) per year, nonpersonnel government funding in these schools is less than 200 rands (~20 USD) per pupil. This is also true around the threshold between quintiles 4 and 5. This implies that the change in government funding at that threshold may be barely noticeable for the schools concerned, rendering that jump less attractive for identification purposes. Second, the figure shows how fees decrease very rapidly with poverty scores for quintile 4 schools, leading to no discernible differences in fees for schools at the right boundary of quintile 4 compared with schools at the left end of quintile 3. Together with the very limited jumps in government funding on crossing quintiles 1–3 documented above, this indicates that the only jump with potential for identifying the impact of school resources is the one from quintile 3 to quintile 4. This threshold appears particularly promising for our purposes since the poorest schools in quintile 4 charge very low fees (close to zero), as do the richest schools in quintile 3, but receive significantly more government funding.

To formally assess the extent to which overall resources change on crossing the different quintile thresholds, we add fees and government funding for each school to construct a measure of overall resources that can be used for running nonpersonnel expenditures. We then estimate the first stage in equation (2) separately for each jump. Table 2 shows the results. Column 1 considers all schools and uses a third-order polynomial and a relatively large window (1.5 standard deviations) around the cutoff. Column 2 uses a linear function with

TABLE 2
CHANGE IN LOG TOTAL RESOURCES AT DIFFERENT QUINTILE JUMPS, GAUTENG PROVINCE

	(1)	(2)	(3)	(4)
Indicator jump:				
4–5	−.016	−.041	.156	.129
	(.121)	(.124)	(.154)	(.15)
3–4	.439	.423	.375	.384
	(.026)***	(.028)***	(.052)***	(.051)***
2–3	−.139	−.015	−.142	−.067
	(.039)***	(.021)	(.087)	(.081)
1–2	−.002	−.029	.089	.102
	(.038)	(.055)	(.101)	(.095)
Sample	All schools	All schools	Secondary schools	Secondary schools
Bandwidth (standard deviation)	1.5	.5	1.5	.5
Polynomial order	3	1	3	1
Spline	No	Yes	No	Yes
Number	826	241	213	57

Note. Robust standard errors are in parentheses. Coefficients are from regressions corresponding to the reduced form from equations (1) and (2). Each coefficient corresponds to a separate regression where the outcome is always log total resources (fees plus government nonpersonnel funding) and the instrument is the indicator function for the corresponding jump in the row. Columns 1 and 2 use all schools, and cols. 3 and 4 use only secondary schools. For each type of school, the first column uses a large window of 1.5 standard deviation around the threshold, and the second column uses a small window of 0.5 standard deviation with spline.
*** Significant at 1%.

spline and a smaller window around the threshold (0.5 standard deviations). Columns 3 and 4 replicate the same estimations using only high schools. As expected, the only coefficients that are consistently large and statistically significant are those for the jump from quintile 3 to quintile 4.[9] Coefficients are in the order of 0.4 when all schools are used and slightly lower when only high schools are considered. This implies that quintile 3 schools close to this threshold obtain on average around 40% more resources than those at the other side of the threshold. This amount is thus both statistically and economically significant. As an illustration, figure 3 shows the jump in log total resources on crossing the different quintile thresholds fitting a linear function allowed to vary by quintile. Each point in the figure again corresponds to the average of log resources for the schools with a given poverty score. The figure clearly shows that only the jump from quintile 3 to quintile 4 represents a genuine increase in overall resources.

Table 3 explores in more detail the jump in resources from quintile 3 to quintile 4 in Gauteng. Different rows correspond to different outcome variables, while the columns correspond to different specifications (and to using all schools as opposed to high schools), as in the previous table. For reference

[9] The only other statistically significant estimate is a negative coefficient at the 2–3 threshold in col. 1. However, the estimate is not stable across specifications and appears on the whole to be both economically and statistically insignificant.

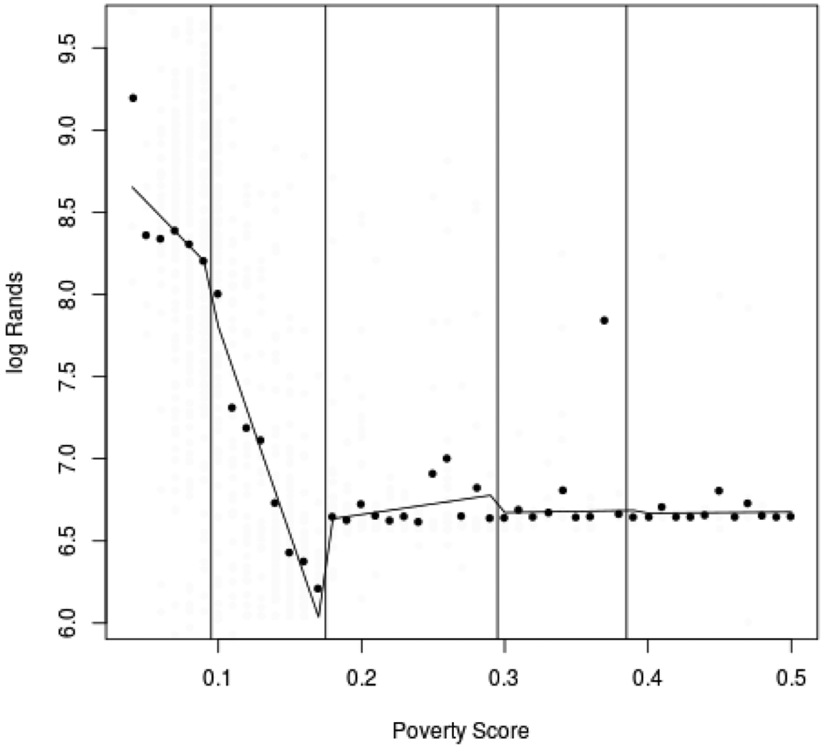

Figure 3. Log total resources by poverty score, Gauteng province.

purposes, the first row reproduces the results of table 2, using the log of total resources as the outcome variable. We wish to assess whether the jump in total resources posited by our administrative data is translated into a detectable jump in material resources owned by the school. The second row in the table thus uses computers per learner at the school as the outcome variable. While there seems to be an increase in computers per learner on crossing the threshold when considering all schools, no jump is discernible when using only high schools. This need not invalidate our approach because there can be reasons why high schools may not spend marginal resources on computers, but given that our analysis will focus on high schools, the fact that no jump in computers is apparent in these schools does raise questions on alternative uses of government funding and thus provides a word of caution when interpreting our results.

The third row in table 3 uses average fees (in rands) as the outcome variable in order to test whether schools compensate lower government funding with higher fees at this threshold. Coefficients are generally small and never significant, suggesting that the phenomenon of fee compensation may be irrelevant

TABLE 3
CHANGE IN RESOURCE-RELATED INDICATORS AT JUMP 3–4, GAUTENG PROVINCE

	(1)	(2)	(3)	(4)
log total resources	.439	.423	.375	.384
	(.026)***	(.028)***	(.052)***	(.051)***
Computers per student	.012	.007	−.001	−.004
	(.003)***	(.003)**	(.006)	(.005)
Fees	54.762	−97.174	250.169	75.225
	(158.588)	(120.66)	(225.055)	(94.514)
Percent African black	.004	.008	.009	.016
	(.012)	(.014)	(.019)	(.024)
Sample	All schools	All schools	Secondary schools	Secondary schools
Polynomial order	3	1	3	1
Spline	No	Yes	No	Yes
Number	1,637	922	484	254

Note. Robust standard errors are in parentheses. Coefficients are from regressions corresponding to the reduced form from equations (1) and (2). Each row uses a different outcome variable. Computers per student refer to the years 2010–12. Fees refer to the years 2008–12, except for 2010. Columns 1 and 2 use all schools, and cols. 3 and 4 use only secondary schools. For each type of school, the first column uses a large window of 1.5 standard deviation around the threshold, and the second column uses a small window of 0.5 standard deviation with spline.
** Significant at 5%.
*** Significant at 1%.

for the schools we are considering.[10] To further probe this issue, the fourth row in the table considers whether the demographic profile of the student body jumps at the threshold. If schools do raise fees when receiving less government funding, we should observe a higher proportion of low-income students on the quintile 3 side. Because African blacks remain the poorest population group in South Africa, we use—as an imperfect proxy for the socioeconomic status of the student body—the proportion of African blacks among students in the posttreatment period. As shown in the table, there is no jump in this variable on crossing the threshold. All in all, the patterns displayed in the bottom two rows of table 3 do not seem to indicate that schools compensate lower government funding with higher fees.

In sum, the data from Gauteng provide us with valuable guidance regarding the validity of our approach. The jump from quintile 3 to quintile 4 appears promising in the sense that it displays a substantial change of resources (although for high schools it does not translate into a higher number of computers). Moreover, fees at the schools around that threshold are close to zero, and, indeed, we do not observe evidence of higher fees compensating lower government funds in these schools. While we cannot ascertain the extent to which these results are generalizable to other provinces, there are reasons to believe

[10] If anything, the coefficients are mostly positive, which would be the opposite of fee compensation. However, the estimates are somewhat unstable possibly as a result of the fast exponential decline in fees in the relevant range (logs are not used because of the large amount of zeros).

that if this pattern is observed in Gauteng, it may also be observed in other South African provinces. First, we will show in our results that we find an effect of quintile assignment on student throughput from grades 10–12 but no effect on enrollment at grade 10. We believe this makes fee compensation in other provinces less probable, as there are few reasons to believe that cheaper schooling would result in higher throughput to grade 12 but not in higher enrollment at grade 10. Second, Gauteng is relatively more urban and, together with the Western Cape, one of the wealthier provinces, so fees may be more prevalent there. Thus, if schools at the 3–4 quintile threshold charge very low fees in Gauteng, it would appear unlikely that schools at a comparable quintile level in other provinces charge substantial fees. We will therefore proceed with our analysis for all provinces focusing on the jump from quintile 3 to quintile 4 and using government funding as the endogenous variable.

B. *First Stage*

Table 4 shows the jump in government funding from quintile 4 to quintile 3 using all provinces. The first row considers the jump in the actual quintile. If quintiles followed the poverty score perfectly, the coefficient for this variable would be −1. In our data the estimated coefficient is around −0.66, showing that about one-third of schools had their quintile reassigned relative to the original placement.

TABLE 4

FIRST STAGE: CHANGES IN QUINTILE AND GOVERNMENT RESOURCES AT JUMP 3–4, ALL PROVINCES

	(1)	(2)
Quintile	−.645	−.669
	(.045)***	(.05)***
Government funding	1.803	1.887
	(.099)***	(.108)***
log government funding	.356	.391
	(.024)***	(.025)***
log government funding 2008	.498	.472
	(.031)***	(.035)***
log government funding 2010	.439	.479
	(.035)***	(.033)***
log government funding 2012	.206	.298
	(.053)***	(.057)***
Bandwidth (standard deviation)	1.5	.5
Polynomial order	3	1
Spline	No	Yes
Number	2,174	868

Note. Robust standard errors are in parentheses. Coefficients are from regressions corresponding to the first stage from equation (2). The first row uses school quintile as outcome, while the remaining rows use different government funding variables. Column 1 uses a large window of 1.5 standard deviation around the threshold, and col. 2 uses a small window of 0.5 standard deviation with spline.

*** Significant at 1%.

The second row uses as outcome our endogenous variable: government funding averaged over the previous 3 years. The jump from quintile 4 to quintile 3 generates almost 200 additional rands per pupil from the government. In log terms, this implies an increase of around 0.35 to 0.4 points, very similar to the magnitudes observed for the jump in log income from the Gauteng data. This is reassuring, and it strengthens the idea that fees do not appear to be very relevant for schools at this margin. The last rows in the table show that this result applies to different years. The increase is smaller for the last year of the sample period, indicating that schools have increasingly seen their assigned quintile recategorized over time.

C. Poverty Score Manipulation

In the process of designing the new poverty scores in 2005–6, schools were allowed to contest their assigned poverty score. This introduces a risk of manipulation in our running variable. Wildeman (2008) documents findings from anonymous interviews with provincial officials involved in the creation of poverty scores. He reports that lobbying was quite limited in most provinces and that ad hoc adjustments were made only for schools near the boundaries of electoral wards (to account for possible discrepancies between the students' socioeconomic characteristics and those of the electoral ward). Garlick (2013) notes that the number of schools to be treated was chosen after the poverty scores had already been assigned, so that precise manipulation of the poverty scores in the neighborhood of quintile thresholds was not possible. It is also apparent from the first stage results that provinces amended the initial quintile assignments for a significant number of schools over time, suggesting that much tampering was done ex post. For these reasons, manipulation during the early poverty score assignments (our instrument) appears unlikely, although it cannot be completely ruled out, particularly for KwaZulu-Natal and Northern Cape, for which our poverty score variable dates from 2008.

We thus follow standard practice in regression discontinuity designs and check whether there is a higher density of schools at the beneficial side of our threshold of interest (i.e., the quintile 3 side) relative to the costly side (the quintile 4 side). If schools have successfully manipulated poverty scores to their advantage, we would expect such asymmetry. Figure 4 shows the distribution of poverty scores for schools in the neighborhood of the quintile 3 to quintile 4 jump. The figure shows no clear indication of an unusual number of schools at the rightward side of the threshold (which corresponds to quintile 3) relative to the leftward side.

We also perform a formal test of poverty score manipulation, taken from McCrary (2008). The test delivers a borderline result, with a p-value of .05.

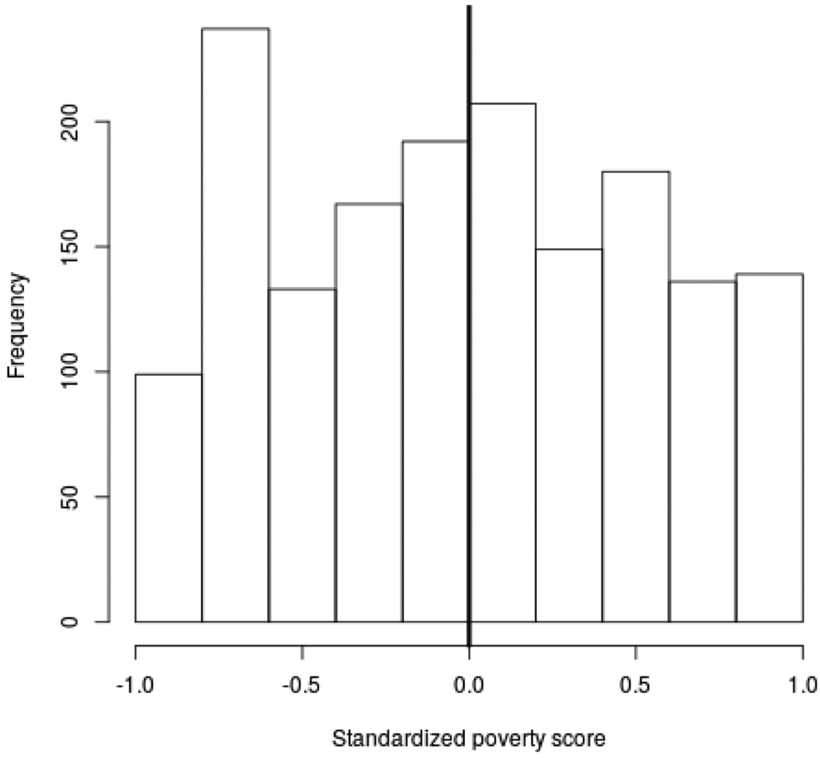

Figure 4. Check of manipulation of running variable: distribution of schools around the threshold, all provinces.

Since our poverty scores are assigned by province, we check if there are specific provinces that drive this result or whether it is the outcome of aggregating the provinces. Performing the test by province we find that one of them, KwaZulu-Natal, has a relatively low *p*-value at .06. This is perhaps unsurprising since KwaZulu-Natal is one of the provinces for which poverty scores are observed in a later year (i.e., 2008), and by that time some manipulation might have taken place. Rerunning the test with all provinces except KwaZulu-Natal yields a *p*-value of .2, a more reassuring result. We will therefore perform our analysis with and without KwaZulu-Natal to check the robustness of our estimates.

D. Balance at Different Sides of the Threshold

The critical assumption of our approach is that no variable relevant for school outcomes jumps at the quintile threshold (aside from nonpersonnel funding). Having addressed the possibility that fees may jump at the threshold, we now consider the question in a more general way, checking whether pretreatment outcomes (i.e., outcomes before the current quintile system was introduced) are discontinuous at the threshold. If our approach is valid, the poverty score threshold determining quintiles ought to have been meaningless before the establishment of the current system, and variables predating 2007 should not increase at that threshold.

Table 5 shows the jumps in pretreatment outcomes at the 3–4 quintile cutoff. As in table 4, each row corresponds to an outcome variable and each column to a different specification: either third-order polynomial with large window or linear with spline and small window. All regressions correspond to the reduced form, where outcomes are directly regressed on the instrument. The first variable considered is the school quintile as it was in 2005. As mentioned above, poverty scores and quintiles existed and determined funding already before 2006, but they were calculated in a very different way. Our first goal is thus to test whether the former poverty score system generated a quintile jump in the same neighborhood as the current one. As it is clear from the ta-

TABLE 5
BALANCE OF PRETREATMENT VARIABLES AND OF COVARIATES CONCURRENT WITH TREATMENT AROUND JUMP 3–4, ALL PROVINCES

	(1)	(2)
Predetermined variables:		
Quintile 2005	−.025	−.107
	(.1)	(.125)
log grade 10	−.051	−.018
	(.058)	(.067)
Throughput grades 10–12	.015	.025
	(.019)	(.024)
Matric attendance rate	−.008	.005
	(.011)	(.014)
Matric pass rate	−.032	−.028
	(.022)	(.023)
Matric attendance rate over grade 10	−.002	.01
	(.021)	(.023)
Matric pass rate over grade 10	−.003	.001
	(.016)	(.018)
Covariates concurrent with treatment:		
log educators	−.033	−.017
	(.049)	(.058)
Share staff hired by school governing body	.002	0
	(.009)	(.01)
Bandwidth (standard deviation)	1.5	.5
Polynomial order	3	1
Spline	No	Yes
Number	2,174	868

Note. Robust standard errors are in parentheses. Coefficients are from regressions corresponding to the reduced form from equations (1) and (2). Each row uses a different outcome variable. Predetermined variables are measured before the establishment of the new quintile system in 2006–7. The quintile variable and matric-related variables refer to 2005. Enrollment variables use the average from 2003–5. Posttreatment covariates are measured from 2008–12. Column 1 uses a large window of 1.5 standard deviation around the threshold, and col. 2 uses a small window of 0.5 standard deviation with spline.

ble, this is not the case, suggesting that the current quintile thresholds were indeed meaningless before 2007. The following rows use as outcome variables the pretreatment outcomes. No coefficient is significantly different from 0 at conventional levels in either specification.

A second set of balance tests we consider regards posttreatment variables that capture the possibility of funds not being used as intended or leading to crowding out in ways other than raising fees. To address these issues, we consider the number of teachers and the share of staff hired by the school governing body. Since the type of resources we are studying are meant to be used for nonpersonnel expenditures, these two variables should be balanced at either side of the threshold if resources have indeed been used as intended and have not led to human resource adjustments by school managers. The last two rows of table 5 show the results. All coefficients are small and statistically indistinguishable from 0, implying that resources appear to have been used as intended and that resources appear not to have altered hiring incentives. Perhaps more importantly, they provide further reassurance on the validity of our approach, that is, that schools at either side of the 3–4 quintile cutoff have indeed similar characteristics.

VI. Regression Discontinuity Estimates

Table 6 shows the main results of our exercise: the estimates of the change in school outcomes on crossing from quintile 4 to quintile 3. Column 1 shows the third-order polynomial specification with basic controls.[11] Column 2 uses the same specification but adds the pre-2007 outcome as a further control. Column 3 preforms the latter analysis but excluding KwaZulu-Natal, given the concerns that manipulation of our observed poverty score might have taken place in this province. Columns 4–6 replicate the same models but use a linear specification with spline for the smooth function of the poverty score.

The first relevant result from the table is that coefficients are relatively small. No coefficient is larger than 0.025, suggesting that the increase of about 200 rands in funding per pupil resulting from the jump from quintile 4 to quintile 3 (approximately 30%–45%) has little impact on student outcomes. The second important result is that there are no positive effects of funding on the pass rate, a common indicator of student learning in the literature and a widely used measure of high school performance in South Africa. Actually, the coefficients for matric pass rates appear to be quite reliably negative when using grade 12 enrollment as the basis for the ratio (row 4). The estimates are virtually zero when grade 10 enrollment is used in the denominator (row 6).

[11] Apartheid era education department and quintile jump dummies other than the one from 3 to 4.

TABLE 6
INSTRUMENTAL VARIABLES RESULTS OF POSTTREATMENT OUTCOMES FOR JUMP 3–4, ALL PROVINCES

	(1)	(2)	(3)	(4)	(5)	(6)
log grade 10	−.016	.008	.018	−.008	0	0
	(.034)	(.019)	(.025)	(.037)	(.022)	(.028)
Throughput grades 10–12	.024	.021	.027	.022	.016	.016
	(.012)**	(.011)**	(.015)*	(.012)*	(.01)*	(.012)
Matric attendance rate	−.003	−.002	−.003	−.003	−.003	−.007
	(.004)	(.004)	(.005)	(.004)	(.004)	(.004)
Matric pass rate	−.02	−.013	−.016	−.022	−.018	−.021
	(.009)**	(.008)	(.009)*	(.01)**	(.009)**	(.009)**
Matric attendance rate over grade 10	.018	.019	.028	.015	.015	.013
	(.01)*	(.009)**	(.014)**	(.011)	(.009)	(.011)
Matric pass rate over grade 10	0	.002	.008	−.003	−.002	−.004
	(.009)	(.006)	(.007)	(.009)	(.006)	(.007)
Bandwidth (standard deviation)	1.5	1.5	1.5	.5	.5	.5
Polynomial order	3	3	3	1	1	1
Spline	No	No	No	Yes	Yes	Yes
Control predetermined	No	Yes	Yes	No	Yes	Yes
Including KwaZulu-Natal	Yes	Yes	No	Yes	Yes	No
Number	2,163	2,138	1,265	866	852	610

Note. Robust standard errors are in parentheses. Coefficients are from instrumental variables regressions from equations (1) and (2). Each row uses a different outcome variable measured after the new quintile system was established in 2006–7. Variables are measured as averages from 2008–12. Columns 1–3 use a large window of 1.5 standard deviation around the threshold, and cols. 4–6 use a small window of 0.5 standard deviation with spline. For each of these windows, the second column adds as further control the corresponding pretreatment outcome (e.g., 2005 matric pass rate in row 4), and the third undertakes the analysis without KwaZulu-Natal schools, for which poverty score manipulation may have occurred.
* Significant at 10%.
** Significant at 5%.

This discrepancy appears to be driven by the positive effect of funding on student throughput from grade 10 to grade 12 (row 2) and, to a lesser extent, by the positive impact on writing the matric exam (row 5). At the same time, no discernible effect is apparent on grade 10 enrollment (row 1). In short, additional nonpersonnel funding appears to have no impact on grade 10 enrollment but increases throughput from grade 10 to 12. The increase in throughput, however, does not translate into higher pass rates.

It is worth noting that these results appear robust across specifications, with coefficients being relatively stable regardless of the polynomial/window used and whether pretreatment outcomes are controlled for. The fact that the pretreatment outcomes do not significantly affect the main coefficients provides further assurance that the sample is indeed balanced on the prereform characteristics, as discussed in Section V.[12] The coefficients with and without

[12] The size of coefficients in table 6 for throughput from grades 10–12 and for matric pass rates are not very different from the corresponding pretreatment variables in table 5. This could lead to concerns that our results are simply due to an increase in power from the inclusion of covariate controls in table 6. Note, however, that the validity tests in table 5 are reduced-form coefficients, while the results in table 6 are the second-stage results from the instrumental variables estimation. Since the

KwaZulu-Natal are also similar, implying that results are not driven by the possible score manipulation in this province.

Finally, in order to further ensure that our results are not driven by specific functional forms, we re-estimate our models with a larger set of polynomial and window specifications with and without controlling for pretreatment outcomes. The results are presented in table A1. We use polynomials of degrees 1–5 with the full window and additionally the optimal window bandwidth from Imbens and Kalyanaraman (2011) for the linear case. We observe that coefficients are stable across the various specifications, implying that results are not driven by a specific functional form or bandwidth. Significance levels decrease as the polynomial order reaches 4 and 5. This is to be expected, given the large degree of nonlinearity introduced (Jacob et al. 2012). In fact, such high levels of nonlinearity may be considered excessive (Gelman and Imbens 2014). Overall, our main conclusions appear to be robust: additional funding leads to a small but significant increase in student throughput, but this does not translate into higher student learning (as measured by matric pass rates).

VII. Discussion

Our empirical analysis rests on a similar conceptual framework as that of much of the applied literature on the impact of additional resources on educational outcomes (Hanushek 2006; Gibbons and McNally 2013). Student learning is the outcome of an education production function, which includes a variety of factors/inputs. The causal relationships underlying the education production function can vary considerably, reflecting different assumptions about the role and level of each input and the interaction between them. Reduced-form specifications of the education production function (such as those usually estimated in randomized evaluations) deliver the overall impact of changes in a given input on student learning. These are the estimates policy makers are typically interested in, which reflect not only the direct (partial) effect of the input of interest but also the interaction with other inputs and the optimizing responses of parents and teachers (Glewwe et al. 2013).

We estimated the total (reduced-form) effect of changes in nonpersonnel school funding on educational outcomes at the end of high school in South Africa. Our reduced-form estimates may reflect a combination of factors. First, it is difficult to say a priori whether nonpersonnel resources are substitutes or complements of other inputs in the education production function. For example, an increase in material input can improve teachers' marginal productiv-

first stage has a magnitude of approximately two, this implies that comparable coefficients for the posttreatment estimates would have to be about double the ones in table 6.

ity by providing better tools and/or infrastructures. At the same time, some nonpersonnel items (such as computers or tablets) could serve as substitutes for teacher's effort, especially in contexts where teacher motivation is low. These possible complementarities/substitutabilities would imply that nonpersonnel resources further affect education production as they alter the incentives to hire additional staff. Similarly, the overall effect of additional nonpersonnel resources is influenced by the interaction of school and household inputs. Optimizing parents may respond to an increase in school resources from the government by decreasing their provision of educational inputs if household and school inputs are technical substitutes (Das et al. 2013) or even by adjusting their school choice. Finally, the effect of nonpersonnel resources also depends on the management and incentive structures faced by principals. These determine how effective schools are at using resources and the extent of misappropriation of funds.

In light of these considerations, how can we interpret the different sets of results shown in table 6? Although our reduced-form results may be open to alternative interpretations, the zero effect of funding on enrollment at grade 10 (which is the first year of noncompulsory schooling in South Africa) may render some of these explanations less plausible. For instance, interpretations relating to school misuse of funds would seem to be at odds with the result of a positive effect on student throughput but no effect on grade 10 enrollment. That is, one would expect that changes in these inputs would affect enrollment and throughput in the same direction. Similarly, the absence of an effect on student enrollment at grade 10 would seem to be at odds with the possibility that our results reflect a migration from untreated to treated schools in response to a change in the perceived quality of education (due to the increased available resources). In addition, our finding that the number of staff hired by the school governing body is balanced at the two sides of the threshold suggests that nonpersonnel resources have not altered hiring incentives.

A possible explanation for the observed differential effect on enrollment and throughput is that schools may react to the per-pupil nature of the funding policy, which increases the incentive to retain students past the compulsory grade/age. They may do this, for example, by being more lenient in allowing students to progress across grades. Gustafsson (2011) shows that failing a grade is one of the main correlates of student dropout past grade 9 in South Africa. As mentioned above, there have been reports of conscious control of the progression of weak students at grade 10 (Borkum 2012).[13] This would explain

[13] The incentives of school principals and government officials with respect to student grade progression have also been noted in the media (http://mg.co.za/article/2011-03-11-matric-quality-vs-quantity).

the positive effect of funds on throughput from grades 10–12. To the extent that these marginal students would have dropped out in the absence of additional funding, this may also explain the negative effect of resources on pass rates for grade 12 enrollees. One can assume that at the dropout margin, students are below average in terms of academic preparedness, which would be consistent with the zero effect of funding on pass rates when grade 10 is used as the basis for the ratio.

Taken together, our results raise doubts on the effectiveness of nonpersonnel funding in improving student learning in our setting. Although additional years of schooling are generally considered to be beneficial for youth, irrespective of final graduation, the South African literature on returns to education reports virtually no earnings returns to completing grades 10 and 11 (Keswell and Poswell 2004). This would imply little or no positive effect of the funding policy on student outcomes through additional years of education. Moreover, the marginal students who benefit from more years at school may be significantly different from the inframarginal students, leading to the possibility of negative externalities from changes in class size and peer composition.[14]

Nonetheless, we note that it would be inappropriate to conclude from our results that allocating resources to relatively poor schools—in contexts such as the one in South Africa—cannot and generally does not improve learning. In particular, our results regard the effects of an increase of 200 rands (~20 USD) in per-pupil resources for secondary schools around the quintile 3–4 threshold. There are three major reasons why resources in similar contexts could have a positive effect on learning. First, the productivity of additional inputs may be higher in the bottom quintile, where the lack of resources is more severe and the baseline learning is lower. Second, results could be different for primary schools in which students, by virtue of being younger, may be more amenable to acquiring additional cognitive skills (Heckman 2000). Third, learning may be nonlinear in resources. It is possible that complementarities between material inputs and indivisibilities imply that larger increases in resources have a proportionally higher impact than smaller increases.

Having said this, while our paper does not conclusively show that nonpersonnel resources do not matter, it does provide cautionary evidence regarding the potential for these resources to close educational attainment gaps, such as the one observed in South Africa. In addition, our paper highlights the possibility that funding disbursed on a per-pupil basis may alter incentives on the subsidized margin.

[14] We are grateful to Rob Garlick for pointing out these possibilities to us.

Appendix

TABLE A1

SPECIFICATION ROBUSTNESS CHECKS: INSTRUMENTAL VARIABLES RESULTS WITH DIFFERENT POLYNOMIAL AND WINDOW SPECIFICATIONS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
log grade 10	−.034	.011	−.018	.015	−.016	.008	−.025	.013	.032	.019	−.005	.006
	(.03)	(.016)	(.031)	(.017)	(.034)	(.019)	(.034)	(.021)	(.036)	(.02)	(.04)	(.021)
Throughput grades 10–12	.021	.019	.028	.022	.024	.021	.018	.014	.019	.016	.02	.015
	(.011)*	(.009)**	(.012)**	(.01)**	(.012)**	(.011)**	(.011)*	(.009)	(.012)	(.01)*	(.012)*	(.01)
Matric attendance rate	.001	.001	.001	.001	−.003	−.002	−.002	−.002	−.002	−.002	0	0
	(.003)	(.003)	(.003)	(.003)	(.004)	(.004)	(.003)	(.003)	(.003)	(.003)	(.003)	(.003)
Matric pass rate	−.014	−.008	−.014	−.009	−.02	−.013	−.021	−.014	−.018	−.014	−.015	−.011
	(.008)*	(.008)	(.008)*	(.008)	(.009)**	(.008)	(.01)**	(.008)*	(.01)*	(.008)	(.009)*	(.008)
Matric attendance rate over grade 10	.018	.018	.024	.021	.018	.019	.011	.012	.014	.016	.016	.016
	(.01)*	(.008)**	(.01)**	(.008)**	(.01)*	(.009)**	(.01)	(.008)	(.011)	(.009)*	(.01)	(.008)*
Matric pass rate over grade 10	.003	.005	.009	.005	0	.002	−.007	−.004	−.001	.002	.001	.004
	(.008)	(.006)	(.008)	(.005)	(.009)	(.006)	(.009)	(.006)	(.009)	(.006)	(.009)	(.006)
Bandwidth (standard deviation)	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	IK	IK
Polynomial order	1	1	2	2	3	3	4	4	5	5	1	1
Spline	No	No	No	No	No	No	No	No	No	No	Yes	Yes
Control predetermined	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Number	2,163	2,138	2,163	2,138	2,163	2,138	2,163	2,138	2,163	2,138		

Note. Robust standard errors are in parentheses. Coefficients are from instrumental variables regressions from equations (1) and (2). Each row uses a different outcome variable measured after the new quintile system was established in 2006–7. Variables are measured as averages from 2008–12. Columns 1–6 use a large window (1.5 standard deviation) around the threshold and different polynomial orders. Columns 11 and 12 use a linear specification with spline and with the Imbens and Kalyanaramana (IK) optimal bandwidth. For each of these windows/polynomial specifications, the second column adds as an additional control the corresponding pretreatment outcome (e.g., 2005 matric pass rate in row 4).

* Significant at 10%.

** Significant at 5%.

References

Borkum, E. 2012. “Can Eliminating School Fees in Poor Districts Boost Enrollment? Evidence from South Africa.” *Economic Development and Cultural Change* 60, no. 2:359–98.

Branson, N., D. Kekana, and D. Lam. 2013. *Educational Expenditure in South Africa: Evidence from the National Income Dynamics Study*. SALDRU Working Paper no. 124, Southern Africa Labour and Development Research Unit, Rondebosch.

Das, J., S. Dercon, J. Habyarimana, P. Krishnan, K. Muralidharan, and V. Sundararaman. 2013. “School Inputs, Household Substitution, and Test Scores.” *American Economic Journal: Applied Economics* 5, no. 2:29–57.

Department of Basic Education, Republic of South Africa. 2003. *Review of the Financing, Resourcing and Costs of Education in Public Schools*. Technical Report. Pretoria: Department of Basic Education.

Du, Y., and Y. Hu. 2008. “Student Academic Performance and the Allocation of School Resources: Results from a Survey of Junior Secondary Schools.” *Chinese Education and Society* 41, no. 5:8–20.

Garlick, R. 2013. “How Price Sensitive Is Primary and Secondary School Enrollment? Evidence from Nationwide Tuition Fee Reforms in South Africa.” Unpublished manuscript, Duke University.

Gelman, A., and G. Imbens. 2014. *Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs*. NBER Working Paper no. 20405 (August), National Bureau of Economic Research, Cambridge, MA.

Gibbons, S., and S. McNally. 2013. *The Effects of Resources across School Phases: A Summary of Recent Evidence*. CEP Discussion Paper no. 1226 (June), Centre for Economic Performance, London.

Glewwe, P., E. A. Hanushek, S. D. Humpage, and R. Ravina. 2013. “School Resources and Educational Outcomes in Developing Countries: A Review of the Literature from 1990 to 2010.” In *Education Policy in Developing Countries*, ed. P. Glewwe, 13–64. Chicago: University of Chicago Press.

Gustafsson, M. 2011. *The When and How of Leaving School: The Policy Implications of New Evidence on Secondary Schooling in South Africa*. Stellenbosch Working Paper no. 9, Stellenbosch University.

Hanushek, E. A. 2006. “School Resources.” In *Handbook of the Economics of Education*, Vol. 2, ed. E. A. Hanushek and F. Welch, 865–908. Amsterdam: North Holland.

Heckman, J. J. 2000. “Policies to Foster Human Capital.” *Research in Economics* 54, no. 1:3–56.

Imbens, G., and K. Kalyanaraman. 2011. “Optimal Bandwidth Choice for the Regression Discontinuity Estimator.” *Review of Economic Studies* 79, no. 3:933–59.

Jacob, R., P. Zhu, M.-A. Somers, and H. Bloom. 2012. *A Practical Guide to Regression Discontinuity*. New York: MDRC.

Keswell, M., and L. Poswell. 2004. “Returns to Education in South Africa: A Retrospective Analysis of the Available Evidence.” *South African Journal of Economics* 72, no. 4:834–60.

Leibbrandt, M., I. Woolard, A. Finn, and J. Argent. 2010. *Trends in South African Income Distribution and Poverty since the Fall of Apartheid*. Paris: OECD.

McCrary, J. 2008. "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test." *Journal of Econometrics* 142, no. 2:698–714.

Nannyonjo, H. 2007. *Education Inputs In Uganda: An Analysis of Factors Influencing Learning Achievement in Grade Six*. World Bank Working Paper no. 98, World Bank, Washington, DC.

Reinikka, R., and J. Svensson. 2004. "Local Capture: Evidence from a Central Government Transfer Program in Uganda." *Quarterly Journal of Economics* 119, no. 2:678–704.

Snilstveit, B., J. Stevenson, R. Menon, D. Phillips, E. Gallagher, M. Geleen, H. Jobse, T. Schmidt, and E. Jimenez. 2016. *The Impact of Education Programmes on Learning and School Participation in Low- and Middle-Income Countries*. 3ie Systematic Review Summary 7, International Initiative for Impact Education (3ie), London.

UNESCO. 2005. *Decentralization in Education: National Policies and Practices*. Technical Report, United Nations Educational, Scientific and Cultural Organization, Paris.

———. 2014. *Education for All Global Monitoring Report 2013/4: Teaching and Learning, Achieving Quality for All*. Technical Report, United Nations Educational, Scientific and Cultural Organization, Paris.

Weidman, J. C., and R. DePietro-Jurand. 2009. *EQUIP2 State-of-the-Art Knowledge in Education: Decentralization*. Washington, DC: US Agency for International Development.

Wildeman, R. A. 2008. *Reviewing Eight Years of the Implementation of School Funding Norms, 2000 to 2008*. Pretoria: IDASA.