Treating schools to a new administration

The impact of South Africa's 2005 provincial boundary changes on school performance

MARTIN GUSTAFSSON AND STEPHEN TAYLOR

KEYWORDS: SOUTH AFRICA, SCHOOL IMPROVEMENT, MATHEMATICS EDUCATION, IMPACT EVALUATION

JEL: C21, H11, I21

MARTIN GUSTAFSSON
STEPHEN TAYLOR

DEPARTMENT OF ECONOMICS
UNIVERSITY OF STELLENBOSCH
PRIVATE BAG X1, 7602
MATIELAND, SOUTH AFRICA
E-MAIL: MGUSTAFSSON@SUN.AC.ZA, STEPHEN@SUN.AC.ZA

AUGUST 2013

This paper is work in progress. The current version was submitted for presentation at the September 2013 conference of the Economic Society of South Africa

Both authors are researchers based at the Department of Basic Education, and are members of Research on Socio-Economic Policy (ReSEP), a group within in the Department of Economics at the University of Stellenbosch. The current paper emerged from work undertaken for the Department of Basic Education to support the development of a new national mathematics, science and technology (MST) strategy. Advice on the interpretation of the data by Rufus Poliah, Willie Venter, Mathanzima Mweli, Jennifer Kinnear and other colleagues in the DBE was much appreciated. However, the authors take responsibility for the views expressed in the paper.

ABSTRACT

The impact that the systems and practices of the education authorities, as opposed to the management at the school, have on school performance is usually difficult to quantify. Provincial boundary changes occurring in South Africa after 2005 appear to create a quasi-experiment that lends itself to impact evaluation techniques. A total of 158 secondary schools experienced a switch in provincial administration and at least two types of switches, one from Limpopo to Mpumalanga and another from North West to Gauteng, were sufficiently common to make statistically significant trends a possibility. Various indicators of Grade 12 mathematics performance are explored which take into account passes at a low threshold, achievement at an excellent level and selection into mathematics. Models used are critically discussed include a simple value-added school production function, a difference-in-difference model and a fixed effects panel data analysis. The data include annual Grade 12 examination results for the period 2005 to 2012, which allow for lags in the impact to be explored. Spatial analysis is used to identify schools close to switching schools to establish whether student commuting effects could have confounded the results. A key finding is that schools moving from North West to Gauteng appear to enjoy benefits associated with the treatment especially as far as the production of students excelling in mathematics is concerned. However, a strong caveat is that the finding depends heavily on just 2012 values and that 2013 examination data will have to be included in the analysis before the study can inform policy recommendations. A brief comparison of institutional aspects of the education authorities in the two provinces North West and Gauteng, drawing from publicly available plans and reports, is provided to help interpret the differences seen in the data. The paper ends with some tentative conclusions in relation to how governance responsibilities in education can be optimally spread across the national, provincial and local levels in South Africa.

1 Introduction

Understanding how to improve not only individual schools, but entire schooling systems, in sustainable and cost effective manners, is fraught with problems. Identifying cause and effect in education systems is inherently difficult, even when good data are available, and often such data are not available. Moreover, ideology often clouds reason in education research and policymaking. Psacharopoulos (1996: 343), a key figure in the emergence of economics of education as a field in its own right, has argued: '[i]n the field of education, perhaps more than in any other sector of the economy, politics are substituted for analysis'.

This paper makes a modest contribution to understanding South Africa's schooling system, and such systems generally, largely by confirming empirically that the effectiveness of the education administration is important, and that there seems to be some truth in the perception that a specific South African province is relatively well governed. But the paper is also intended to contribute towards the growing trend towards more innovative use of data in exploring causal relationships in education.

Section 2 explains the institutional background, including some specifics of the South African Grade 12 examinations, how educational improvement is viewed in South Africa, and the 2005 provincial boundary changes. Section 3 explains the economics of education framework used in the paper, including ways of measuring impact. Section 4 discusses the Grade 12 examinations data used for the analysis, including the derivation of eight indicators of improvement. Section 5 provides both some descriptive analysis of the data as well as the modelling of impact, using a simple value-added school production function, a difference-in-difference model and a fixed effects panel data analysis. Section 6 explores explanations from an institutional angle, using published plans and reports of two provinces as a point of departure. Section 7 concludes.

2 Institutional background

For several decades national examinations at the Grade 12 level have provided the only more or less reliable measure of school performance in South Africa and much behaviour has understandably been oriented towards Grade 12 indicators, in particular 'pass rates', the percentage of students successfully obtaining the certificate or surpassing minimum thresholds in individual subjects. Around 40% of youths have obtained the Grade 12 certificate in recent years (Gustafsson, 2011). Both public and private schools participate in the national public examinations. Around 90% of candidates in recent years writing the examinations at schools have been full-time students and it is these students that this paper focuses on. What is thus excluded are part-time students, nearly all repeaters, taking fewer than the full set of seven subjects. Around half of schools have part-time students.

Currently, the full menu of subjects in the Grade 12 examinations includes 27 non-language subjects. The system was changed rather fundamentally between 2007 and 2008. Subjects were redesigned, a distinction between standard grade and higher grade examination papers across all subjects was removed and it became compulsory for all students not taking mathematics to take mathematical literacy, a relatively easy subject. This paper focuses on improvements in mathematics, a subject that is widely taken and is of special importance for economic development. In 2005, 58% of examination candidates took mathematics and 8% of all candidates took mathematics at the higher grade. In 2011, 45% of candidates took mathematics. The percentage of schools with mathematics candidates was 99% in 2005 and 97% in 2011. In 2005, 59% of schools had students taking mathematics on the higher grade. Simkins (2010) provides an important account of the 2007 to 2008 transition with respect to mathematics. Van der Berg (2004) describes the legacy of race-based inequality in South Africa that continues to influence performance in mathematics.

Poor student performance, in particular in mathematics, is widely acknowledged as being a key hurdle to economic and social development in South Africa. The low numbers of black African students achieving sufficiently high scores in mathematics to enter university studies requiring minimum levels of mathematics competencies continues to worry policymakers. The data used for this paper, which include the race of students, and population data indicate that in 2011 only 0.5% of black African youths were obtaining 70 or more out of 100 in Grade 12 mathematics, against a figure of 3.5% for youths of other races[1]. The importance of increasing high-end mathematics performance influences the focus of the data analysis presented below.

It is often useful to think of interventions aimed at improving schools in terms of whether they follow a more school-specific interventionist approach, or a more system-wide structural approach. Both are clearly important and receive much attention in South Africa. The former would include interventions such as the Dinaledi programme, designed to provide capacity building to a set of 500 secondary schools in the areas of mathematics and physical science (World Bank, 2010; Blum, Krishnan, Legovini, 2010). The latter would include the Annual National Assessments programme, introduced in 2011 and aimed at improving accountability against national standards below the Grade 12 level (South Africa: Department of Basic Education, 2011).

An event not aimed at improving school quality, but which nonetheless can assist in understanding school improvement, and in particular system-wide structural solutions, were relatively minor changes to the boundaries of South Africa's nine provinces occurring after 2005. Most of the provinces were completely new entities established in 1994 as part of the dismantling of the apartheid system. Seven provinces saw their boundaries change in 2005, all except for Western Cape and Free State. A total of 710 schools, 158 of which offer Grade 12,

---

[1] The other races would be coloureds, whites and Indians, using to the official South African terms.

experienced a change in provincial administration. The distribution of the 158 schools is shown in Figure 1. All except for two of the 151 schools involved in the larger province-to-province school 'migrations' had Grade 12 groups that were between 90% and 100% black African (the two exception schools were in the Mpumalanga to Limpopo migration). The schools in question are thus interesting in terms of understanding questions that are critical for policymakers.

The boundary changes occurred largely to ensure that all municipalities fell into just one province and did not span two provinces. It must be noted that municipalities do not play a role in the administration of schools in South Africa. Instead, provincial governments are responsible for administering schools. Reporting directly to the provincial authorities are a number of 'education districts', whose boundaries often coincide with those of the municipalities, though institutionally they are independent of each other. Provincial education departments thus determine to a large degree how schools are resourced, how they are provided with capacity building and the manner in which they are accountable for their educational achievements. Crucially, provincial education departments are funded by the provincial department of finance, and not the national education authorities. Responsibilities left in the hands of the national authorities include maintenance of the national curriculum, the setting of the Grade 12 examinations and quality control in the marking process, the Annual National Assessments programme and central bargaining with teacher unions around teacher salaries.

Figure 1: Schools experiencing a province change in 2005

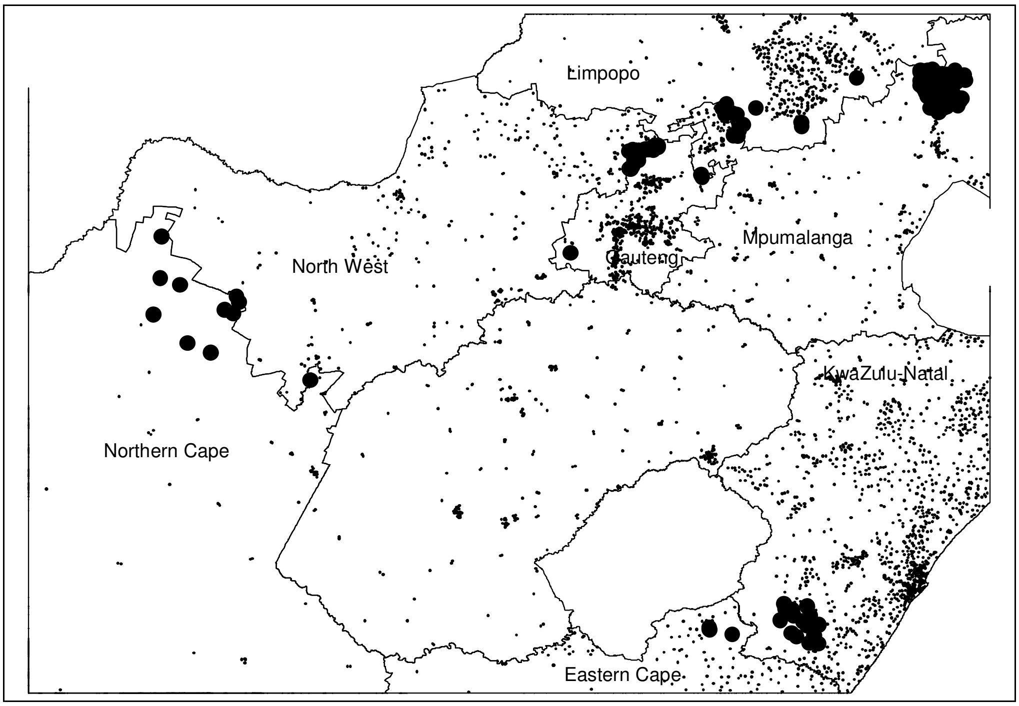

Source: Constructed in Stata using dataset analysed in this paper. The points refer only to schools offering Grade 12.
Note: Provincial boundaries are those created by the 2005 boundary changes. Schools which moved from a neighbouring province are marked by large points.

A key question for this paper is exactly when the change to a new province could be expected to change the dynamics within the affected schools. Though legislation changing the provincial boundaries was passed in December 2005, after the 2005 examinations had been written, the transfer of schools to their new provincial administrations was not immediate. In the official national master list of schools the new province appears by the start of the 2008 school year, in January 2008. Officially, the 'treatment' thus only began in 2008, though the fact that schools knew that this was coming before 2008 could have influenced behaviour in the schools at an earlier point in time. For instance, school principals may have felt invigorated or dejected by the fact that they were being transferred to what was perceived as a better or worse province. These kinds of concerns lay behind serious political protests during this period. In particular, there was a perception that Gauteng was a relatively good province to be in from a service delivery angle. To illustrate, the township of Khutsong, which was moved from Gauteng to North West following the 2005 boundary changes, saw protests that led to the widespread destruction of property, but no deaths, over the dissatisfaction of residents with the move (Centre for Development and Enterprise, 2007). In 2008, further legislation moved Khutsong back to Gauteng. No other such reversals of the 2005 boundary changes occurred, but the Khutsong events provide a telling indication of how much differences in the capacity of administrative units to deliver services matter to ordinary South Africans.

3 Theoretical underpinnings

An education system is internally efficient when it optimises the use of available resources, including, essentially, human resources in the form of teachers, in order to produce the educational outcomes it sets itself out to achieve. External efficiency is realised when the education system optimally contributes to the kind of human capital development required to advance development priorities such as raising incomes, reducing unemployment and promoting social cohesion. A key shift with regard to external efficiency has been the relative decline of manpower planning approaches, or the determination of specific vocational skills sets required in the economy, and a rising awareness of the importance of basic competencies, in language and subjects such mathematics, for economic development. Hanushek and Woessman (2009) have been particularly influential in promoting the shift towards the prioritisation of basic competencies in policymaking circles.

How best to advance the internal efficiency of schooling systems has been a hotly debated matter, but there appears to be a convergence of opinions, partly due to more empirical research in this area. Importantly, a greater use of impact evaluations, either evaluations planned upfront or evaluations undertaken opportunistically with the available data, have underlined that things are often not what they seem. Interventions that at face value appear to be doing the right thing often produce no or very disappointing results when subjected to rigorous impact evaluation. Bruns, Filmer and Patrinos (2011) identify a few intervention types that appear promising, if interventions are carefully designed and appropriately combined: strengthening parent involvement in schools by giving parents more reliable information on the performance of their children; giving school principals more autonomy; making teacher employment conditions and pay less insensitive to student performance; improving the in-service training of teachers. Taylor (2013) demonstrates how getting good educational materials to students in South Africa has made a positive difference to Grade 12 examination results.

As argued by Gustafsson and Mabogoane (2013) and others, the optimal bundle of interventions to improve the internal efficiency of a schooling system may differ rather fundamentally according to the level of development of that system, and society as a whole. For instance, very human capital-intensive interventions where teams of experts work with individual schools to bring about change, whilst appropriate in high income countries with limited pockets of dysfunctional schools, may be inefficient in developing countries where basic problems, such as teacher absenteeism, are widespread and school turnaround experts are scarce. In developing countries, there appears to be a stronger argument for the use of system-wide structural reforms, focussing for instance on overcoming information asymmetries and introducing better incentives for teachers and school managers.

One advantage with impact evaluations, in particular more rigorous ones which are planned upfront and make use of techniques associated with randomised control trials (RCTs), is that they help to combat publication bias, a bias that adversely affects the relationship between research and policymaking, according to Duflo, Glennester and Kremer (2006). Publication bias occurs where the finding of a zero impact results in non-publication of some research. This easily happens where researchers rely exclusively on production function-type analysis of relationships between inputs and outputs. Due to inherent problems with this technique, it becomes easy to attribute a zero impact finding to problems with the data, even if the problem may be that there is no impact. As a consequence, policymakers are denied access to information on interventions that do not work, when such information is clearly necessary. Impact evaluations tend to come with fewer methodological problems and are more likely to be taken seriously when no impact is found.

Yet typical impact evaluations need to be interpreted with care. They tend to focus on pockets of the system, as opposed to the system as a whole, and thus do not deal with general equilibrium dynamics such as reactions of teacher unions over time to certain interventions. Important system-wide changes, such as the remarkable improvements in Brazil's PISA results between 2000 and 2009 are difficult to decipher with the available data and methods, yet they are of huge importance. Theory needs to be informed by good knowledge of institutional factors. A drawback with a 2010 impact evaluation of the Dinaledi programme, which pointed to some evidence of positive impacts on the basis of patterns seen in historical data not originally designed for impact evaluation purposes, is an insufficient acknowledgement of institutional realities. Specifically, there is a possibility that improvements seen in Dinaledi schools were at least partially the result of migration to these schools by good students from neighbouring schools, given that South African policy allows for considerable migration of this type and that evidence points to this being widespread. In impact evaluation terms, the possibility of contamination of the treatment group was not taken into account. The analysis presented below attempts to deal with this risk through some simple spatial analysis techniques.

4 The data

The data used for the analysis were student records of results in the Grade 12 year-end examinations, for the years 2005 to 2006 and 2008 to 2012. These data were obtained from the Department of Basic Education and were collected over several years during work done for the Department. The key challenge was to link schools across years as the system of school identifiers had changed over time and there were many inconsistencies with respect to school identifiers even in years when provincial boundaries were not changing. Essentially the examinations database was not designed with school-level year-on-year comparison in mind. Data for 2007 were not easily obtainable, though this gap did not appear to limit the analysis seriously. The point of departure was to link the schools from the 2005 and 2011 datasets. The final school-level panel of data included all schools for which examinations data were available in both 2005 and 2011. For the other years, there were missing values for some of these schools. The extent of this is shown in Table 2. Table 1 indicates the number of students covered in the 2005 and 2011 datasets. The decline in the average number of mathematics students per school should be seen in the context of the introduction in 2008 of mathematical literacy, a subject not considered in the analysis. Simkins (2010: 19) has estimated that if mathematical literacy is considered, the 2007 to 2008 curriculum change resulted in an increase in the total amount of mathematics knowledge produced by the schooling system.

Table 1: 2005 and 2011 data coverage in terms of schools and students

			Mathematics examination takers			Average Grade 12 mathematics students per school	
	Total schools	Private schools	2005	2005 (HG)	2011	2005	2011
EC	824	24	41,989	2,390	37,482	51	45
EC>KN	15	0	756	28	741	50	49
FS	293	8	12,589	1,670	9,986	43	34
GP	585	112	47,644	8,995	28,309	81	48
KN	1,463	34	78,496	8,527	58,743	54	40
LP	1,180	29	41,039	9,901	31,914	35	27
LP>MP	83	1	2,886	897	3,362	35	41
MP>LP	13	2	578	52	356	44	27
MP	360	13	18,309	2,235	14,527	51	40
NC	95	2	3,269	502	2,454	34	26
NW>GP	29	0	2,326	191	1,297	80	45
NW>NC	11	0	426	4	302	39	27
NW	328	10	16,966	1,640	9,312	52	28
WC	342	30	21,053	4,736	13,662	62	40
Other	7	1	217	13	172	31	25
Total	5,628	266	288,543	41,781	212,619	51	38

Note: 'Other' refers to province switches involving fewer than ten schools within the switch. The 2005 student counts refer to all mathematics candidates and those candidates taking mathematics at the higher grade (HG).

There are many ways in which to measure a school's improvement with respect to mathematics over the years. A fundamental question is one of breadth versus depth, or what Hanushek and Woessman (2009: 22) refer to as 'rocket scientists versus education for all'. Is improvement occurring through the extension of mathematics skills to more students, or is the number of high-end achievers increasing? Of course both may be occurring simultaneously, though it is reasonable to assume that to some degree there is a trade-off between the two. If available teachers must teach more students they may have less time to support high-end achievers. Which is best for the external efficiency of the schooling system, in other words to advance economic growth, is not a straightforward matter.

The following table reflects eight indicators aimed at allowing a broad treatment of the question of improvements in mathematics. Clearly, values from the old examination system (2005 to 2006) and the new one (2008 to 2012) are not comparable in any simple sense. This should not cause problems for the analysis, however, as long as the analysis focuses on differences in the change over time. The number of mathematics passes per school, including or excluding standard grade passes from the old system, is one indicator of improvement, and one that is commonly used in official reports. To obtain the number of high-level passes per school, a threshold of 70 out of 100 was used, and only higher grade mathematics in the old system was counted. The use of absolute numbers such as these, in a context where school sizes differ, leads to results that must be interpreted with caution, yet there are good reasons for looking at absolute numbers, and not just means. Part of the reason for this is that numbers may increase exceptionally, not just because communities around schools grow, but also because successful schools attract students from outside their regular catchment areas, as discussed below. Unless the commuting students are exceptionally good performers, this kind of success would not be identified if only means are analysed.

The percentage of Grade 12 students taking mathematics is an important indicator of the level of opportunity afforded to students. The percentage of mathematics candidates passing the subject is an indicator that has been widely referred to in the official reports. The average mathematics score out of 100 obtained by students is a key indicator of improvements with respect to the quality of mathematics teaching and learning. To obtain the 2005 and 2006 averages, the standard grade score was multiplied by 0.75 to reflect the lower value attached by educationists to this score. This factor has been used by Foxcroft (2006: 70) and is roughly in line with the correspondences between standard grade and higher grade scores identified by Simkins (2010: 8) for physical science. The mark at the 95th percentile is based just on students who wrote the mathematics examination and for 2005 and 2006 again standard grade scores were multiplied by 0.75. In the case of the 2006 data, values for each student were not available. Instead, the number of students achieving within a band, for instance from 71 to 80, was provided. Thus for 2006 pseudo-records for individual students had to be created. Imputed student scores were distributed linearly within each band, so for instance four students in the band 70 to 80 would be given the scores 72.1, 74.4, 76.6 and 78.9. Scores were then adjusted so that they reproduced the actual mean, which was given separately for standard grade and higher grade within each school. Finally, the average spread across all candidates is the total score divided by all students sitting for the examination, even those not taking mathematics. The adjusted standard grade scores for 2005 and 2006 were used here too. This final indicator is the only one that is influenced directly by both the breadth of mathematics participation in the school and depth in the sense of the quality of mathematics learning that occurs.

Table 2: Mean values for the eight indicators for all years

	2005	2006	2008	2009	2010	2011	2012
Schools	5,628	5,621	5,600	5,585	5,616	5,628	5,602
Number of passes (SG included for 2005)	28.2	34.6	23.2	22.3	20.8	17.3	19.8
Number of passes (SG excluded for 2005)	4.9	7.7					
Number of high-level passes	1.3	5.9	4.4	3.1	3.0	2.2	2.7
% taking mathematics (SG included for 2005)	58.6	61.3	52.0	52.1	49.5	45.9	44.0
% of mathematics-takers passing mathematics	55.6	54.2	47.1	46.3	46.4	45.3	50.1
Average mark	25.4	24.1	31.1	29.8	30.7	30.6	32.4
Mark at the 95th percentile	46.2	50.1	58.9	56.6	57.7	56.0	59.2
Average mark spread across all candidates	14.6	14.1	15.5	15.1	14.8	13.7	14.1

5 Analysis

As pointed out by Simkins (2010), the Grade 12 examination performance of schools displays considerable year-on-year variability. This is to be expected given for instance that teachers, in Grade 12 and earlier grades, change over time. Moreover, national, provincial and local pressures around the criteria to use when promoting students into Grade 12 from Grade 11 change. The following graph, which uses data from the whole country, provides one way of viewing the year-on-year variability. The especially high degree of re-ranking of schools occurring between 2006 and 2008 can be probably be attributed to the introduction of the new examination system in 2008, which would have been more difficult for certain schools to adjust to than others, somewhat independently of the level of performance of schools in 2006. To illustrate further the high variability between years, only 42% of schools were in the same performance quintile in 2011 and 2012 with respect to the percentage of mathematics students passing the subject.

Figure 2: Year-on-year change on school percentile rankings

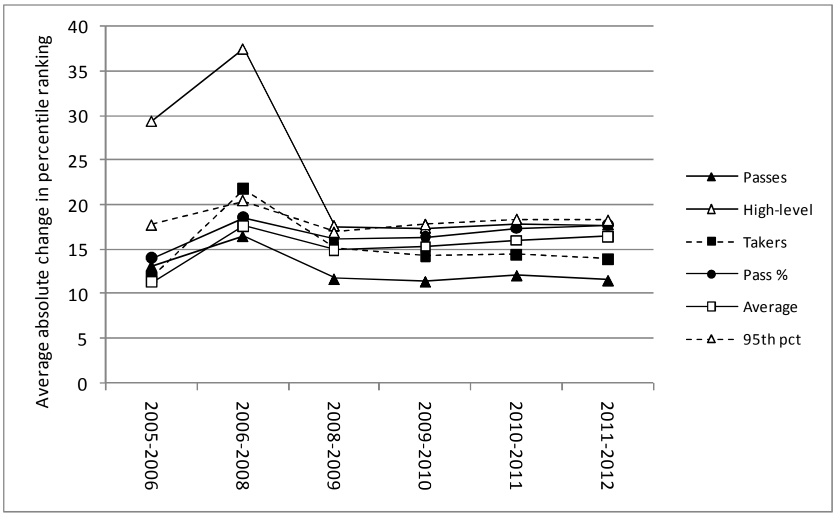

The data analysis described below borrows from the methods associated with randomised control trials (RCTs), given the quasi-experimental nature of the data. Schools across the country were not treated to an education improvement intervention as they might be in a proper RCT. Instead, a ‘treatment’ in the sense of a shift to an alternative provincial administration occurred, with schools being somewhat randomly selected for this. There was randomness in the sense that schools were not placed in new provinces because there was something special about them educationally. They were moved simply because there was a need to rationalise the system of local government. Of course there was not true randomness in the selection of schools as they all fell within specific geographical areas close to provincial borders where the redrawing of these borders was deemed necessary to align the provincial and local divisions. This non-randomness problem is dealt with below using two techniques. Firstly, a key variable that is strongly correlated with the status of being a moving school, namely school size, is controlled for in the analysis. One thus deals with the possibility that relatively large schools experienced changes that were different from those of other schools and that any indication of exceptional changes amongst switching schools was simply an artefact of their exceptional sizes. Secondly, changes within schools that did move are compared to changes in the general geographical area, which includes schools which did not move to a different province, to test the hypothesis that some different local dynamic unrelated to the province change lay behind any changes seen in the switching schools. In a sense, this involves a spatial version of what Duflo *et al* (2006) would refer to as a regression discontinuity analysis. To begin with, the focus is just on the start and end points in the panel of data, namely 2005 and 2012. But the analysis then explores various ways of using the full panel of data, with all years included.

The following empirical model, which is an extension of the basic regression model for an RCT put forward by Duflo *et al* (2006: 6), is our point of departure.

$$E_{i2} = \hat{\beta}_0 + \hat{\beta}_1 D_i + \hat{\beta}_2 E_{i1} + \hat{u}_i \qquad (1)$$

Score E_{i2} for school *i*, in terms of one of the eight indicators described in Table 2, in the final period 2, was predicted using E_{i1}, the score for the indicator in period 1, and *D*, a dummy variable indicating whether the school left the province. The year 2005 was considered period 1 and year 2012 period 2. The coefficient of interest is $\hat{\beta}_1$, the estimate of the impact of moving to a new province, specifically the extent to which moving schools would experience gains in the indicator in question, over and above gains experienced by non-moving schools. Importantly, a decline in the indicator value for moving schools could still represent a gain, and hence a positive coefficient, if the decline experienced by moving schools was smaller than declines experienced by non-moving schools. Such a scenario of arrested declines in treatment schools was found in Blum *et al*'s (2010) evaluation of the Dinaledi intervention.

Equation (1) was run separately for different sending provinces. For instance, it would be run for all the 839 schools that either remained in Eastern Cape (EC) or moved to KwaZulu-Natal (KN) from Eastern Cape (see Table 1). In the case of North West (NW), which 'sent' schools to two provinces, two separate dummy variables *D* would be used. In the case of Limpopo (LP) and Mpumalanga (MP), provinces that 'sent' schools to each other, regressions excluded schools that were received, so the Limpopo regression would include schools that stayed in Limpopo or moved to Mpumalanga, not schools originally in Mpumalanga which moved to Limpopo. Altogether 32 regressions were run, for the four sending provinces (Eastern Cape, Limpopo, Mpumalanga, North West) and the eight indicators. There were 40 coefficients for the dummy variable *D* given that movements out of North West occurred towards two provinces.

Table 3: Coefficients for 'treatment' dummies

Indicator	EC>KN	LP>MP	MP>LP	NW>NC	NW>GP
Number of passes (SG included for 2005)		5.3***			
Number of passes (SG excluded for 2005)		4.5**			
Number of high-level passes					2.1***
% taking mathematics (SG included for 2005)	-12.5**				-8.3**
% of mathematics-takers passing mathematics		-8.3***	10.2*		14.2***
Average mark		-3.6***			7.1***
Mark at the 95th percentile				-7.4*	11.0***

***** indicates that the estimate is significant at the 1% level of significance, ** at the 5% level and * at the 10% level.**

The above table reports on the slope coefficient for the province moving dummy *D*, wherever the coefficient was statistically significant at least at the 10% level. It is important to remember that the level of statistical significance would be lower for province switches involving fewer schools, such as the North West to Northern Cape switch. In fact, the two switches for which there are the largest number of schools, 'LP>MP' and 'NW>GP', have the largest number of statistically significant coefficients in Table 3. This should caution one against rejecting the possibility of educational gains associated with the other province switches. The eighth indicator has been excluded from the table as it produced no statistically significant values.

The Limpopo to Mpumalanga switch appears to be associated with a different kind of change to the North West to Gauteng switch. In the case of the former switch, the increase in the number of passes in conjunction with a decrease in the percentage of mathematics students passing mathematics and no change in the number of high-level passes suggests a scenario where more students with lower capabilities started taking mathematics, perhaps due to easier promotion of students from Grade 11 to 12, leading to both an increase in the number of passes and the number of failures. The change thus seemed more a question of quantitative expansion than qualitative improvement. The opposite seemed to occur in the schools moving from North West to Gauteng, where the percentage of students taking mathematics declined, the number of students passing did not seem to change significantly, but the number of students achieving a high-level pass increased. As the 'NW>GP' changes are the ones most in line with the need to increase mathematics excellence amongst black African learners, and thus most consistent with key South African policy aims, the rest of the paper focuses on this province switch. Figure 3 provides a map of the distribution of the schools in question.

Figure 3: Area around schools moving from North West to Gauteng

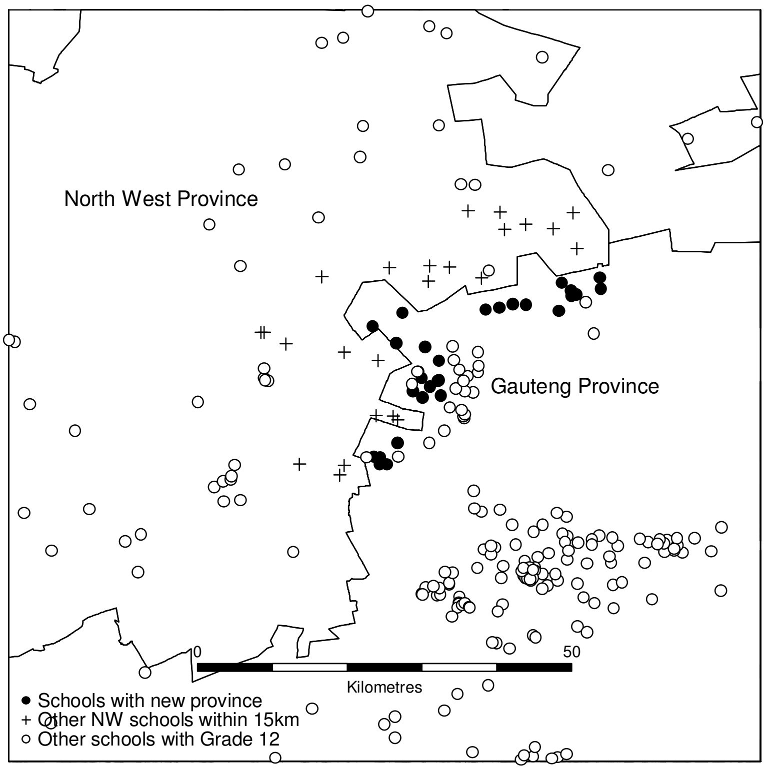

One possible specification error that should be considered is that equation (1) omits a variable capturing an inflow of better learners into certain schools from other schools. Specifically, it is conceivable that the shift in the provincial boundary between North West and Gauteng changed commuting patterns between home and school. There may have been ‘contamination’ of the treatment group as students from nearby North West schools which did not move to Gauteng decided to begin attending the province-switching schools, which now represented an opportunity for better schooling that had moved some kilometres closer. Students who decided to switch schools are likely to have been better performers and students with better educated parents, more aware of and able to make use of new opportunities for their children. This kind of ‘shopping around’ for better schools is common in South Africa, where school zoning is often non-existent or only weakly applied outside of high income urban neighbourhoods. Gustafsson (2011: 23) finds that at the secondary school level, around 17% of students attend a school which is not the closest school to home, the largest reason for this being that the closest school is not satisfactory in terms of the quality of schooling offered.

A distance of 15km was considered an upper threshold for commuting. Non-switching schools from the sending province which were at most this distance from the switching schools were identified as a new control group. The result was just one nearby non-switching school in Limpopo and a more interesting 24 schools in the case of North West. This meant it was only meaningful to control for selection effects in the case of the North West analysis. A significant negative coefficient for the dummy variable denoting the group of 24 schools in North West might be interpreted as evidence that gains in the switching schools were the result of the depletion of better students from these non-switching schools, as these students chose to move to the switching schools. But the same dummy variable also serves another purpose. If combined with the dummy for the switching schools, it might indicate whether there was some factor describing this part of North West province, independently of the any provincial boundary change, which might have contributed to an improvement in learning outcomes. For instance, new economic opportunities in the area could have stimulated an interest in schooling and made it easier for households to access and subsidise educational activities. This would have affected not just switching schools, but even nearby schools.

A further enhancement to the value-added model of equation (1) that seemed necessary was the inclusion of some indicator of population density, urban status or school size. Experimentation with measures of population density, using not just the size of each school's Grade 12 enrolment but also the student counts of surrounding schools, with a weighting factor that diminished with distance, revealed a few interesting patterns. In particular, population density emerged as being a statistically significant predictor in some of the North West models referred to in Table 3. But what seemed an even stronger candidate for inclusion in the model was simply the size of the school's Grade 12 group. Using data for the entire country from different years, it was found that if the percentage of all Grade 12 students becoming high-level mathematics achievers was regressed on the number of Grade 12 students, the dependent variable increased by 0.5 percentage points for around every 100 Grade 12 students. This was found whether one used data for all schools, or just schools above the median Grade 12 group size, of around 70 students (the figure changed depending on which year's data one analysed). The coefficient on the Grade 12 student count was always significant at the 1% level, even if dummy variables for the official poverty quintiles of schools were introduced as additional explanatory variables. To illustrate what this means, using the Grade 12 group size in 2012 as one's ranking variable one obtains group sizes of 36 and 112 at the 25th and 75th percentiles respectively. Using the global average of 2.0%, one would expect the number of high-level achievers in the two types of schools to be on average 0.7 and 2.2 students. However, the size effect would add an extra 0.6 students to the last figure, making it 2.8 students. Why would larger schools have proportionally more high-level achievers? It could be the result of an economy of scale dynamic. For instance, larger schools would have more mathematics teachers, who could assist each other professionally. But the explanation could also be that larger schools are in urban areas where students have better access to resources such as libraries and the internet. What is clear from the Figure 3 map and the values in Table 1 is that the schools switching from North West to Gauteng are in densely populated areas and especially large. Controlling for school size in the model seemed essential.

To arrive at equation (2) below, the additions to equation (1) were the dummy variable δ, indicating that a school was one of the 24 nearby non-switching schools, and the number of Grade examination candidates T, entered in the quadratic form.

$$E_{i2} = \hat{\beta}_0 + \hat{\beta}_1 D_i + \hat{\beta}_2 E_{i1} + \hat{\beta}_3 \delta_i + \hat{\beta}_4 T_i + \hat{\beta}_5 T_i^2 + \hat{u}_i \qquad (2)$$

The next two tables provide descriptive statistics for the at most 357 schools run through the above model as well as model outputs for five different versions of the variable E. For E, three of the indicators described in Table 2 were used, as well as a variant of one of them, namely high-level passes expressed as a percentage of all Grade 12 students. The model outputs in all four models are consistent with the hypothesis that province-switching schools experienced changes that were significantly better than what was seen in other schools, even nearby non-switching schools, after controlling for the size of the Grade 12 group and excluding the possibility that switching schools experienced an influx of better students from nearby schools. In all models, the coefficient for the dummy variable indicating that a school moved from North West to Gauteng is positive and statistically significant, at least at the 5% level. Moreover, the coefficient for being a nearby school is not statistically significant. How large are the coefficients on the switching dummy? The coefficients dealing with the average and 95th percentile are over half as large as they were in the value-added regression models with no controls (see Table 3). One can think of an improvement of around 0.08 standard deviations per year in the mean score as being roughly the fastest change possible in an education system. The coefficient in Model 1 of Table 5 for the province-switching dummy is 0.09 of a standard deviation, when expressed in annual terms. In other words, if the coefficient does indeed refer to an impact, it is as large an impact as one might expect. The corresponding ratios for the other three models are lower, but at least 0.04.

Table 4: Descriptive statistics for North West model data

Variable	Mean	Min.	p10	Median	p90	Max.	Std. dev.
Average mark 2012	35.19	13.90	22.00	34.12	50.10	69.05	10.5
High-level passes 2012	2.04	0.00	0.00	0.00	6.00	36.00	4.5
95th percentile 2012	61.66	20.00	41.00	63.00	81.00	100.00	15.6
High-level passes as a % 2012	1.91	0.00	0.00	0.00	5.71	27.78	3.4
Dummy for nearby	0.07	0	0	0	0	1	0.3
Dummy for NW>GP	0.08	0	0	0	0	1	0.3
All examination takers 2012	81.63	6.00	21.00	69.00	156.00	316.00	56.0
Average mark 2005	24.86	5.18	12.32	22.72	42.47	66.62	11.3
High-level passes 2005	0.94	0.00	0.00	0.00	2.00	32.00	3.2
95th percentile 2005	46.15	11.00	27.75	44.00	68.00	98.00	15.3
High-level passes as a % 2005	0.81	0.00	0.00	0.00	2.44	13.89	2.3

Table 5: Regression outputs for North West models

Variable	Coefficient (t stat)			
Dependent *E* used for 2012:	1. Average mark	2. Number high-level passes	3. 95th percentile	4. High-level passes as a % of all candidates
Constant	22.06*** (13.46)	0.00 (0.01)	34.40*** (13.00)	0.68* (1.92)
Dummy for nearby	0.42 (0.21)	0.07 (0.15)	3.55 (1.24)	0.19 (0.35)
Dummy for NW>GP	6.64*** (3.56)	1.21*** (2.65)	7.49*** (2.75)	1.11** (2.09)
All examination takers 2012	0.03 (1.10)	0.00 (0.57)	0.16*** (4.01)	0.00 (0.57)
All examination takers 2012 squared	0.00 (-0.87)	0.00** (2.48)	0.00*** (-2.86)	0.00 (0.03)
E in 2005	0.45*** (9.97)	1.05*** (25.95)	0.39*** (7.79)	0.96*** (15.67)
Exam takers jointly signif. at 5% level?	No	Yes	Yes	No
N	351	357	351	357
R^2	0.260	0.757	0.286	0.449
Adjusted R^2	0.249	0.753	0.275	0.441

*Note: *** indicates that the estimate is significant at the 1% level of significance, ** at the 5% level and * at the 10% level. Missing values resulting in a reduction of the sample was always in non-switching schools. This latter group remained the same 29 schools in all models.*

To test the hypothesis that there was a local phenomenon, independent of the boundary change, which stimulated educational performance in the general area of the switching schools, the dummy δ from equation (2) was recalculated so that it took on the value 1 if a school was either a switching school or one of the non-switching nearby schools. The new dummy variable would thus refer to 53 schools (29 plus 24). Had the local phenomenon existed, this recalculated dummy would have emerged as a statistically significant predictor of performance and ‘overridden’ the switching dummy. Details on just the two dummy variables for the four models described above with the new δ are provided below. The equation remains the same, the only difference being the recalculation of δ. The results are mixed. The first two models are consistent with the hypothesis that the boundary change impacted positively on affected schools. The coefficients on the new dummy are not significant, whilst the coefficients for the switching dummy remain statistically significant, at least at the 10% level. In the third and fourth models, the coefficients for both dummies are not significant.

Table 6: Regression outputs for recalculated North West models

Variable	Coefficient (t stat)			
Dependent *E* used for 2012:	1. Average mark	2. Number high-level passes	3. 95th percentile	4. High-level passes as a % of all candidates
Dummy for either nearby or NW>GP	0.42 (0.21)	0.07 (0.15)	3.55 (1.24)	0.19 (0.35)
Dummy for NW>GP	6.22** (2.45)	1.14* (1.84)	3.95 (1.06)	0.91 (1.26)

Had data only for the years 2005 and 2012 existed, it would probably have been acceptable to conclude that switching from an administration such as that of North West to an administration such as that found in Gauteng had a clearly positive impact on at least certain critical indicators of mathematics performance. However, we do have data for the intervening years and closer analysis using those data point to a less clear-cut picture.

The following graph illustrates the mean number of high-level passes per school across various categories of schools which were in North West in 2005. It is clear that much of the correlation seen in the models discussed above is a product of what occurred in just one year, 2012, when the indicator value for the switching ('NW>GP') schools jumped. A comparison between switching schools and the category 'Rest NW large' is instructive. The latter category includes large schools in the province other than the switching and nearby schools, where selection into the category occurred from larger to smaller schools until a mean of 129 was reached, which was the mean Grade 12 size of the switching schools in 2012. A comparison between the two groups reveals how exceptionally weak the switching schools have been over most of the period, relative to other similarly large schools, and that even if the 2012 improvement in the switching schools is sustained, this would be more a matter of bringing these schools up to the norm, as opposed to the realisation of exceptionally good results. Similar patterns emerge if one examines the trends over the entire period for the mark at the 95th percentile and high-level passes expressed as a percentage of all Grade 12 students.

Figure 4: Trends for high-level passes in original North West

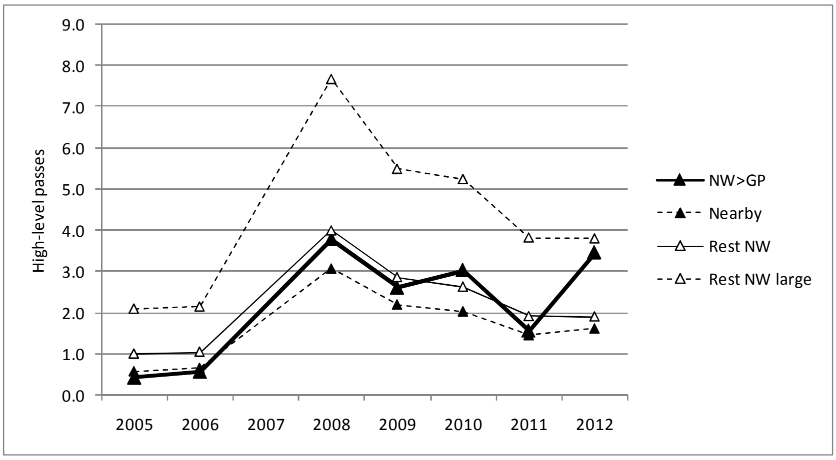

Note: In this graph and the graphs that follow 'Rest NW' is all schools in North West in 2012 except the 'Nearby' group, and 'Rest of NW large' is a subset of 'Rest NW, where the subset is the 139 schools needed, starting from the largest school, to make the mean Grade 12 group size in the subset 129, using 2012 data.

The next graph illustrates that with respect to the average mark, again what occurred in just 2012 would influence to a large degree the analysis results. There was a noticeable narrowing of the gap between switching schools and the rest of the province between 2006 and 2008, but if one compares this to the similar catch-up phenomenon seen in nearby schools, the possibility of early province switching effects seems not to be supported.

Figure 5: Trends for average school mark in North West

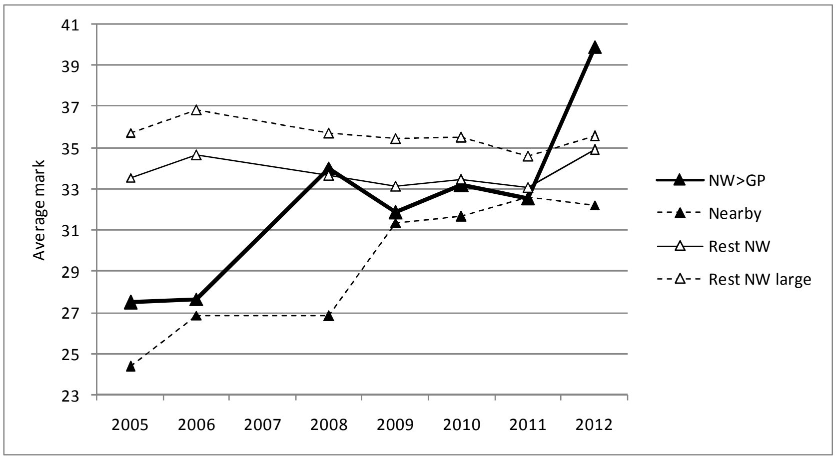

Note: 2005 and 2006 school averages were standardised so that they displayed the same overall means and standard deviations as in 2008.

There is one indicator, not discussed yet, that does seem to point to a sustained and exceptional positive trend in switching schools. The next graph illustrates the proportion of schools with at least one high-level pass. This indicator would reflect the distribution of excellence across schools and could been seen as an indicator of potential for further improvements. The emergence of single high achieving students would often be due to randomly distributed talent amongst students, but if schools systematically display a greater presence of such students, one might conclude that some institutional factor, for instance a more dedicated teacher, is at play. As seen in Figure 6, switching schools started off with a particularly low indicator value. Only 20% of schools had at least one high-level mathematics achiever, against 40% for other similarly large schools. By 2008, or the end of the first school year during which the new province, Gauteng, would have administered the school, 80% of switching schools had at least one high-level achiever and the switching schools remained at the top or almost the top amongst comparator groups for all years through to 2012, when the 'jump' occurred.

Figure 6: Trends for distribution of high-level passes in North West

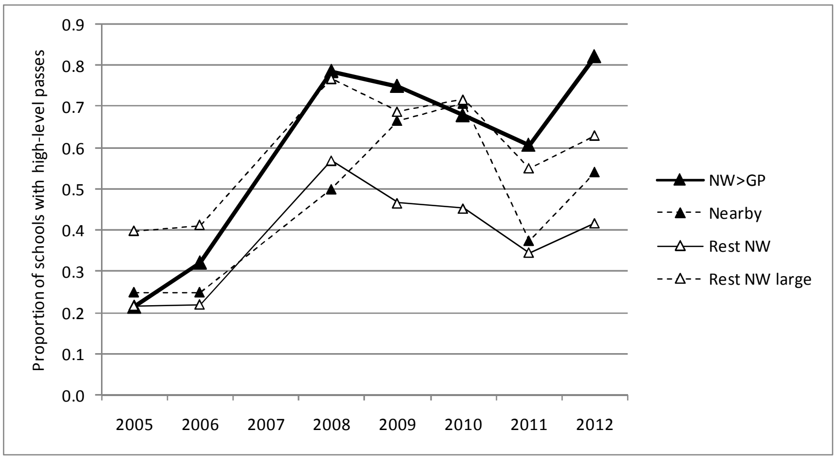

Figure 7 provides linear trendlines for the trends seen in Figure 6 with 90% confidence intervals. Whilst switching and nearby schools were not distinguishable in terms of the indicator in 2005, the general trend suggests that by 2012 they were, on the basis of the non-overlapping confidence intervals, and that switching schools had moved substantially ahead of their neighbours. However, a similarly clear divergence is not seen when the trendline with confidence interval for the ‘Rest NW large’ group is compared to the group of switching schools (this comparison is not illustrated in Figure 7).

Figure 7: Trendlines for distribution of high-level passes in North West

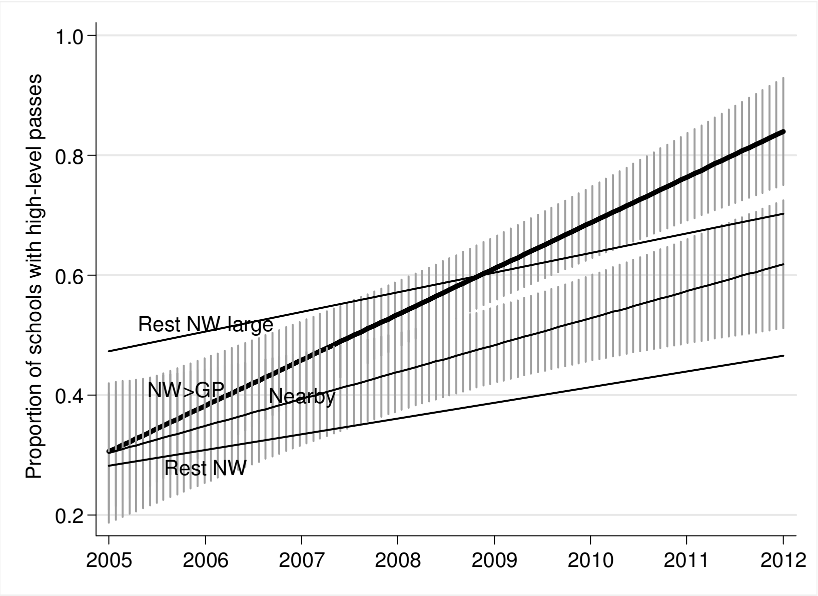

Note: Shaded areas are the 90% confidence intervals for the trendlines ‘NW>GP’ and ‘Nearby’.

The difference-in-difference (DiD) model offers a more rigorous method than the above graph for examining trends in a manner that takes into account values for all years. Schlotter, Schwerdt and Woessman (2010: 19) provide a discussion of difference-in-difference analysis aimed at a less technical audience. Equation (3) below is based on a form recommended by Khandker, Koolwal and Samad (2010: 72) for data containing more than two points in time.

$$E_{it} = \hat{\beta}_0 + \hat{\beta}_1 D_{it} P_t + \hat{\beta}_2 D_{it} + \hat{\beta}_3 P_t + \hat{u}_i \qquad (3)$$

For the dependent variable all measures of one outcome, for instance number of high-level passes, are included, for all schools *i* and all periods in time *t*. Explanatory variables are, firstly, the dummy *D*, for switching schools and, secondly, P, the period and, thirdly, the interaction of the two. For the period, values 1 to 8 were used, with value 3, for 2007, missing. Results for the two difference-in-difference models that supported the hypothesis of the boundary change having a positive impact are provided below. In each model, the coefficient of interest is the one for the interaction of the switching dummy and period. It is positive and statistically significant in both models. The indication is thus, for instance, that the mark at the 95th percentile increased in each period by 1.5 percentage points more in switching schools than in other schools. Here the negative coefficients for the province switching dummy simply reflect the fact that over the periods mean values in the switching schools have been lower than in other schools. Other models tested used the number of high-level passes and high-level passes as a percentage of all candidates. In these other models, not only was the coefficient of interest not significant, the R^2 values were below 0.01.

Table 7: Regression outputs for North West difference-in-difference models

Variable	Coefficient (t stat)	
Dependent *E*:	Average mark	95th percentile
Constant	23.99*** (42.73)	47.67*** (63.18)
Dummy for NW>GP (*D*)	-6.02*** (-3.06)	-3.61 (-1.37)
Period (*P*)	1.49*** (13.92)	1.90*** (13.25)
Interaction of *D* and *P*	1.12*** (3.01)	1.51*** (3.02)
N	2541	2541
R^2	0.089	0.085
Adjusted R^2	0.088	0.084

*Note: *** indicates that the estimate is significant at the 1% level of significance.*

Even the results for the two models shown in Table 7 constitute weak evidence in favour of a provincial boundary effect, as the introduction of explanatory variables from a localised effect operating independently of the boundary change render the coefficients of interest in Table 7 insignificant. Specifically, the introduction of the dummy for either being a switching or nearby school, plus the interaction of this dummy with *P*, as seen in Table 5, renders the two coefficients for the interaction between being a switcher and *P* (the last variable in Table 7) insignificant.

There is one difference-in-difference model where the switching of the province gets very close to emerging as a sufficiently significant predictor of the outcome even when one controls for a broader localised effect in the model. If one examines the patterns for the distribution of high-level passes seen in Figure 6 within a DiD model, one obtains the results seen in Table 8. Because the outcome, whether a school has at least one high-level mathematics achiever, is binary, a logit model is required. Here the coefficient for the interaction between being a switching school and period is significant at the 11% level (see p-value of 0.109), in other words a hairsbreadth away from the 10% threshold at which statistical significance would often be declared. Moreover, the coefficient for the interaction of δ and *P* can be considered non-significant. The key finding is then that being a switching school is associated with an annual increase of an additional 4 percentage points (see the last column) in the probability of having at least one high-level mathematics achiever.

Table 8: Logit model outputs for North West DiD model

	Coefficients	p-value	Change in probability
Constant	-1.04***	0.000	0.03
Period (*P*)	0.11***	0.000	0.04
Dummy for NW>GP (*D*)	-0.19	0.707	-0.04
Interaction of *D* and *P*	0.15	0.109	0.04
Dummy for either nearby or NW>GP (δ)	0.03	0.929	0.01
Interaction of δ and *P*	0.07	0.304	0.02
N	2585		
Pseudo R^2	0.032		

*Note: The dependent variable is whether a school produced at least one high-level mathematics achiever. *** indicates that the estimate is significant at the 1% level of significance. The change in probability is obtained using the* mfx compute *command in Stata.*

One disadvantage with the difference-in-difference model, and the previous graph, is that these analyses do not maintain the link between schools across time periods. They therefore concentrate on comparing means from different points in time, without considering the slopes or trends displayed by individual schools. This loss of information reduces the degree of certainty around the results. To illustrate the problem, if one reruns the second model in Table

5, but first scrambles the baseline values for high-level passes within each of the three groups (nearby schools, switching schools and other schools), meaning the baseline values are randomly reassigned to other schools within the group, the dummy variable for switching schools emerges as statistically insignificant in several repeated runs of the model.

A way of overcoming the loss of data from intervening years, a problem associated with equation (2), as well as the loss of school links across time periods, a problem observed with equation (3), is to use a panel model. Equation (4) below illustrates a panel model with fixed effects for each individual, or school. The variables S in brackets refer to dummy variables for each of the schools. Each school thus carries its own intercept λ. Within the dummy variable for switching schools, D, the value only became 1 from 2008 onwards, in other words from period 4 onwards, reflecting the official change in the handover of the school to the new administration. In a panel model, there needs to be variation within each individual over time with respect to each variable. The period, ranging from 1 to 8 with 3 missing, was included as P. Period was interacted with the switching dummy. Finally, the size of the Grade 12 group, T, was included in a quadratic form.

$$E_{it} = \hat{\lambda}_0 + \left(\hat{\lambda}_1 S_{i=2,} + \ldots + \hat{\lambda}_n S_{i=n}\right) + \hat{\beta}_1 D_{it} + \hat{\beta}_5 P_t + \hat{\beta}_6 D_{it} P_t + \hat{\beta}_3 T_{it} + \hat{\beta}_4 T_{it}^2 + \hat{u}_i \qquad (4)$$

The Table 9 results indicate that for all four of the dependent variables of the Table 5 models, the fixed effect models yield a significant coefficient for the school switching dummy, at least at the 5% level. Here the key variable of interest is the dummy variable D, whose coefficient captures the difference in output associated with the treatment having been implemented. The interaction term is also important insofar as it reflects the trend found after the treatment has been introduced, in other words during the 2008 to 2012 period. The coefficient here is significant in all four models, and positive in three of them, all except for the 95th percentile model.

Table 9: Regression outputs for North West fixed effects models

Variable	Coefficient (t stat)			
Dependent *E:*	1. Average mark	2. Number high-level passes	3. 95th percentile	4. High-level passes as a % of all candidates
Constant	31.49*** (33.24)	0.47 (1.37)	52.69*** (36.18)	1.13*** (2.96)
Dummy for NW>GP (*D*)	7.74*** (2.66)	3.89*** (3.59)	19.21*** (4.30)	2.72** (2.26)
Period (*P*)	0.00 (-0.01)	0.16*** (6.54)	-1.28* (-1.86)	-0.28 (-1.53)
Interaction of *D* and *P*	1.28*** (18.84)	-0.35** (-2.13)	1.73*** (16.54)	0.14*** (5.11)
All examination takers 2012	-0.10*** (-7.22)	0.01 (1.10)	-0.06*** (-2.86)	0.00 (0.40)
All examination takers 2012 squared	0.00*** (3.88)	0.00 (1.46)	0.00 (1.16)	0.00 (-0.69)
N	2541	2585	2541	2585
Number of schools	371	372	371	372
R^2 overall	0.030	0.110	0.024	0.002

*Note: *** indicates that the estimate is significant at the 1% level of significance, ** at the 5% level and * at the 10% level.*

All four models are consistent with the hypothesis of a boundary change effect. The predictions of the average mark model are illustrated in Figure 8 below. In a sense, Figure 8 can be seen as a smoothing of the trends already seen in Figure 5, plus the fixed effects process. To illustrate the meaning of this latter process, if the same analysis occurred after 10 had been subtracted from the mean score of every switching school, the curves in the graph would display exactly the same pattern and would intersect at the same point. The only difference would be that the vertical axis would have shifted upwards slightly to reflect a small drop in the overall average, brought about by the downward adjustment in the averages of the 29 switching schools. In interpreting Figure 8, what is key is that in periods 1 and 2 the performance of switching schools was clearly below that of non-switching schools, whilst from period 4 onwards the opposite was true, using throughout 95% confidence intervals for both curves.

Figure 8: Plot of fixed effects model for average mark

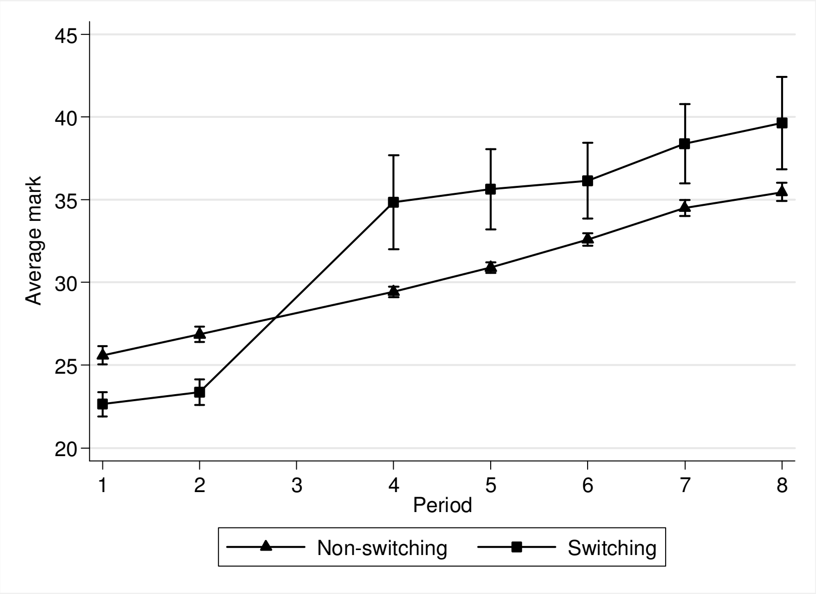

Though the models in Table 9 include school size as a control, they exclude the dummy for switching plus nearby schools. Including this dummy variable as was done in Table 6 renders the coefficients for the switching effect insignificant. First appearances thus deceive. The possibility that a local effect other than the boundary change lies behind the improvements seen in the switching schools cannot be discarded.

6 Institutional and policy explanations

Text still to come.

7 Conclusion

Of five ‘migrations’ of schools from one province to another resulting from the 2005 redrawing of provincial boundaries in South Africa, only one migration could be properly analysed with the given data, with a view to establishing whether the change in province resulted in improved mathematics results. This migration was the one from North West to Gauteng, which involved 29 schools. This was the only migration with a sufficient number of switching schools and with enough nearby non-switching schools to allow for statistically significant results and for the testing of the possibility that rather than the switch, changes in commuting patterns by students lay behind performance improvements.

The results for the 29 schools are compatible with the hypothesis that moving to Gauteng resulted in performance improvements, but the results cannot be considered sufficiently conclusive to inform policymaking, specifically to assert beyond reasonable doubt that emulating the practices of the Gauteng administration in other countries will improve mathematics performance. Yet with an additional year’s data, such a recommendation may be supported.

Standard models focussing on just two points in time, 2007 and 2012, support the notion that switching to Gauteng improved mathematics results, with respect to both average scores and performance at the top end of the spectrum. This finding is upheld even when two alternative explanations are taken into account, namely that performance in the 29 schools improved because high performing students from neighbouring schools moved to the province-switching schools and that some local phenomenon not connected to the provincial boundary change stimulated performance in the general area of the 29 schools.

Yet the finding is strongly dependent on a surge in performance in 2012. If models using several points in time are considered, then the possibility cannot be excluded that the apparent impact of the switch in province is masking a increase in the performance of schools in the general geographical area, on both sides of the new provincial boundary. Yet the results seem strongly suggestive that there was a province switching effect. If the 2013 examination results continue to reflect the exceptionally high levels of performance seen in 2012 in the 29 schools, the evidence can probably considered conclusive.

References

Blum, J., Krishnan, N. & Legovini, A. (2010). *Expanding opportunities for South African youth through math and science: The impact of the Dinaledi programme*. Washington: World Bank. [Unpublished report]

Bruns, B., Filmer, D. & Patrinos, H.A. (2011). *Making schools work: New evidence on accountability reforms.* Washington: World Bank.

Centre for Development and Enterprise (2007). *Voices of anger: Protest and conflict in two municipalities*. Johannesburg: CDE. Available from: <http://www.cde.org.za> [Accessed July 2012].

Duflo, E., Glennester, R. & Kremer, M. (2006). *Using randomization in development economics research: A toolkit*. Washington: NBER.

Foxcroft, C.D. (2006). Evaluating the school-leaving examination against measurement principles and methods. In Vijay Reddy (ed.), *Marking matric*. Pretoria: HSRC: 58-71.

Gustafsson, M. (2011). *The when and how of leaving school: The policy implications of new evidence on secondary schooling in South Africa.* Stellenbosch: University of Stellenbosch.

Gustafsson, M. & Mabogoane, T. (2013). Where to from here? From fact to act. In Nick Taylor, Servaas Van der Berg and Thabo Mabogoane (ed.), *Creating effective schools*. Cape Town: Pearson: 256-282. [Forthcoming]

Khandker, S.R., Koolwal, G.B. & Samad, H.A. (2010). *Handbook on impact evaluation: Quantitative methods and practices*. Washington: World Bank. Available from: <https://openknowledge.worldbank.org/bitstream/handle/10986/2693/520990PUB0EPI1101Official0Use0Only1.pdf?sequence=1> [Accessed June 2013].

Psacharopoulos, G. (1996). Economics of education: A research agenda. *Economics of Education Review*, 15(4): 339-344.

Schlotter, M., Schwerdt, G. & Woessman, L. (2010). *Methods for causal evaluation of education policies and practices: An econometric toolbox*. Bonn: IZA. Available from: <http://ftp.iza.org/dp4725.pdf> [Accessed April 2013].

Simkins, C. (2010). *Interpreting the mathematics, science and English performance of candidates in the November 2008 National Senior Certificate examination*. Johannesburg: Centre for Development and Enterprise.

South Africa: Department of Basic Education (2011). *Action Plan to 2014: Towards the realisation of Schooling 2025*. Pretoria.

Taylor, S. (2013). *The impact of study guides on matric performance: evidence from a randomised experiment*. Pretoria. [Forthcoming]

Van der Berg, S. (2004). *Education: The crisis in schooling*. In S Brown and A Fölscher (ed.), Taking power in the economy: Gains and directions. Cape Town: Institute for Justice and Reconciliation.

World Bank (2010). *DIME Brief: Expanding opportunities for South African youth: Dinaledi's math and science*. Washington.