Statistics Norway
Statistisk sentralbyrå

Torbjørn Hægeland and Lars J. Kirkebøen

Documents

School performance and value-added indicators - what is the effect of controlling for socioeconomic background?

A simple empirical illustration using Norwegian data

Statistics Norway/Research Department

Contents

1. Introduction

One of the main arguments for constructing school performance indicators is that raw, unadjusted results are influenced by factors beyond the schools' control. Pupil composition is maybe the most important of such external factors. Pupils differ with respect to ability and learning potential, and are not randomly allocated across schools. Differences in results between schools are therefore likely to be driven partly by differences in pupil composition.

There are two obvious ways - not mutually exclusive - of adjusting the results of schools for differences in pupil composition:

1: Control for pupils' prior attainment
2: Control for pupils' socioeconomic status (SES)

The choice of what to control for is in most cases restricted by data availability, both with respect to prior attainment and SES variables. In many settings, only a limited set of SES variables may be available. To what extent school performance indicators are affected by which (if any) SES variables that are controlled for, is an empirical question.

The purpose of this short paper is to give an empirical illustration of how school performance indicators are affected by differences in how SES variables and prior attainment is controlled for. The analysis is carried out on Norwegian data for lower and upper secondary schools. The dataset is quite rich in SES information by international standards. Although the illustrations provided may be relevant for indicators on other levels and in other countries, the relationships between students' scores, prior attainment and different sets of SES variables may be both level- and country-specific. The "true" relation between – and accordingly the appropriate model of - SES variables and attainment may differ across subjects, educational levels and countries, and there may be differences in how similar SES variables are defined and measured across countries. One should therefore be cautious in drawing general conclusions with respect to what SES variables that should be included in models for school performance indicators.

2. Model

Consider the following regression model:

(1) $A_{ij} = \beta F_{ij} + \gamma PA_{ij} \ \Sigma q_j S_{ij} + \varepsilon_{ij}$

Where A_{ij} is achievement for pupil *i* in school *j,* F_{ij} is a vector of variables describing socioeconomic background. PA_{ij} is a vector or single variable representing prior achievement. S_{ij} is an indicator variable which equals one if pupil *i* is in school *j* and zero otherwise. ε_{ij} is the error term.

In this model, school performance indicators are the estimated school fixed effects ($\hat{q}_j$), which may be interpreted as the average student achievement at school *j,* conditional on socioeconomic background, as captured by F, and prior achievement, as captured by *PA,* being equal to the sample mean (the model is without a constant term). These school fixed effects can be compared with unadjusted school results, computed as the average of A_{ij} within in each school (or equivalent, estimating (1) without including F and PA).

Within this simple model framework, we may consider the following four types of models:

- Unadjusted results (UR), which includes only school identifiers (S)
- Contextualised attainment models (CA), which includes SES variables (F) and school identifiers
- Value added models (VA), which includes measures of prior achievement (PA) and school identifiers (S)
- Contextualised value added models (CVA), which includes SES variables (F), measures of prior achievement (PA) and school identifiers

In this paper we compare the estimated school effects from the different types of models. We also look at how estimated school effects are affected by which sets of SES variables that are included in the F vector within the CA and CVA models. The Norwegian school performance indicators currently published[1] correspond to the CA model with SES specified as the "full F vector", see below.

The results in this paper are based on estimating the fixed effect model in (1) using ordinary least squares for all specifications. This is done for simplicity, but it should be noted that this may not be the appropriate method for all result measures, as some of them take on only a few values. However, methods more suited for discrete ordered outcomes, e.g. ordered probit models, yield very similar results to those presented in this paper.

3. Data

This section describes the data sources used in the estimation of school performance indicators, the construction of important variables and sample restrictions.

[1] Documented in Hægeland et al. (2005). For a description of data and methods in English, see Hægeland et al. (2004).

3.1 Data on student achievement

Information on individual marks is collected annually by the Directorate for Education and Training. The data contain a personal identification number, an identification number of the school from which the pupil graduated and information on grades and subjects. We consider two different datasets for pupil performance:

1: All students who completed compulsory education in Norway (10^{th} grade in the lower secondary school) in the years 2003 and 2004. Here, we consider one indicator based on the sum of marks in all (eleven) subjects and one indicator based on the final national written exam. Marks are awarded on a scale from 1 to 6, where mark 6 indicates that the pupil holds exceptionally high competence, and 1 indicates that the pupil has attained little competence in the subject. Hence, the first result measure ranges from 11 to 66, and the second from one to six.

2: All students who completed upper secondary school ("allmenn studieretning") in 2005. Here, the indicator is based on marks at the final National written exam in Norwegian, which is common for all students. This exam is also graded on a scale from 1 to 6. For this sample, we also have information of prior achievement (lower secondary school diplomas).

3.2 Data on socioeconomic status

By combining a large number of administrative data sources, we have assembled detailed information on pupil characteristics and family background for all pupils who completed compulsory education (10^{th} grade in the lower secondary school) in the years we study. This provides us with a wide range of family background characteristics. In addition to basic demographic information, we have information of parents' education, immigrant status, parents' wealth, income and (un)employment histories, disability status and receipt of social assistance. Below, we describe these variables in more detail.

Demographic information

We construct dummy variables for the pupil's gender, quarter of birth (given graduation in the year they turn 16) and graduation in years earlier than expected.

Family structure

The following sets of variables are included

- Parents' marital status - dummy variables reflecting whether they are married (to each other), cohabitants, separated, divorced or neither of them.
- Unknown parents - dummy variables indicating whether the father and/or mother is unknown

- The age of the mother and father at the birth of their first child - dummy variables reflecting age intervals
- The number of full siblings and the pupil's rank in the birth order - detailed set of dummies
- Half siblings - dummies indicating the number of half siblings

Parents' education

We have information of the highest completed level of education for each parent. Parents' education is classified into five categories: Lower secondary (up to 9 years of schooling), upper secondary (10-12 years), lower tertiary (13-16 years) and higher tertiary education (17 years or more). We also include an additional category for missing education information. Since missing education information typically appears for immigrants, missing information is also interacted with immigrant status. Based on this classification, we construct dummy variables for all combinations of father's and mother's education.

Immigrant status

Pupils who were born abroad by non-Norwegian parents and pupils whose parents were born abroad are classified as immigrants in our analysis. We construct a detailed set of dummy variables indicating country or region of origin (Scandinavia, ex-Yugoslavia, other Eastern Europe, Turkey, Western Europe (plus USA, Canada, New Zealand, and Australia) , Somalia, other Africa, Sri Lanka, Iran, Iraq, Vietnam, Pakistan, Vietnam, other Asia (plus Oceania), Latin America). We also control for the age of immigration *for the pupil*, with dummy variables distinguishing between those who born in Norway or immigrated before they were three years old, and those who immigrated when they were 3 to 5, 5 to 7, 7 to 9, 9 to 11, 11 to 13 or 13 years or more. We experimented with more detailed country classifications, and also with distinguishing between first- and second-generation immigrants on country/regional level, but the results were basically unaffected.

Economic resources

Based on information on individual taxable labour income, we calculate a measure of family income for the pupil as the sum of the father's and the mother's income during the last ten years (regardless of marital status). Although income tends to be quite persistent over time, we have included not only current but also previous income, to make it reflect family income over the period the pupil went to school, and not only around the time ofgraduation. We then construct dummy variables reflecting the position (quintile) in the family income distribution (for the parents of graduating pupils, not the population as a whole). Based on information of individual taxable wealth, we calculate a measure of family wealth for the pupil as the sum of the father's and the mother's wealth for the year prior to graduation (regardless of marital status). Wealth typically increases over the major part of the life cycle. We therefore construct age-specific wealth distributions, where we divide families into five-year age intervals based on the average age of the parents. We then construct dummy variables indicating whether the family belongs to one of

the upper four deciles of its respective age-specific wealth distribution. Note that a majority of families have negative taxable wealth, since their net financial wealth is negative, and the tax value of housing, which is most common non-financial asset, is far below market value. Negative taxable wealth is reported as zero.

Unemployment, disability pension and social assistance

We use information of parents' unemployment history to construct variables indicating the incidence of unemployment for the ten years prior to the pupil's graduation from lower secondary school. We define a person to be unemployed in a given year if he or she was registered as unemployed for at least three months of a calendar year. Based on this definition, we construct dummy variables, separately for mother and father, for unemployment in the graduation year, and for unemployment one, two, three, four and five or more years during the ten-year period prior to graduation. Similarly, we construct variables indicating the receipt of disability pension and social assistance. We define a person to be on disability pension in a given year if he or she received disability pensions for more than six months of the calendar year. Our criterion for defining a person as receiving social assistance is that he or she received at least Nkr 20,000 (approx € 2,500) in a given year. The dummy variables for disability pensions and social assistance are constructed in the same manner as for unemployment.

Sample attrition

Between 5 and 10 percent are excluded from the estimating samples,. Missing information of family background or prior achievement are the major reasons for excluding individuals from the sample. But also a significant number of pupils with too few marks and without information on which school they attended are excluded from the sample. Indicators for small schools (number of pupils below the limit for publication) are not reported in the figures. Indicators are published for around half of the schools, which have around three quarters of the pupils.

4. Results

Below, we present the results from estimation of the model types described in Section 2 on our datasets. First, we give an overview of the different specifications of the SES part of the model (the F-vector) that we use to illustrate to what extent “simpler” SES specifications produce different school performance indicators than more detailed specifications. Given the large number of SES variables that are available in the dataset, there are many combinations that may be tried out. We have estimated a large number of different specifications, but present results for only a subset of them to illustrate our main findings. These are:

- A: Unadjusted results.
- B: Basic specification: School year, gender and for upper secondary schooling also year of birth, plus years of schooling of highest educated parent (linear).
- C: Basic specification plus dummies for highest education levels of parents, dummy for parents living together, dummies for first and second generation immigrants, dummies for "region" of origin for immigrants (Western/non-Western).
- D: Full specification, as described in previous section.

In addition to how the SES part of the model is specified, there may also be differences in how we control for prior attainment, i.e. how the PA vector in (1) is specified (when data for this is available). We consider four different alternatives:

- No control for prior attainment
- Control for prior attainment in the same subject.
- Control for prior attainment (mark) in the same subject, plus the total sum of marks in all subjects at the prior level
- Control for prior attainment (mark) in the same subject, plus marks in core subjects at the prior level (entered as separate variables).

This leaves us with a set of specifications as shown in Table 1.

Table 1: Overview of specifications estimated

	Control for prior attainment			
SES specifications	No	Mark in same subject	Same subject, total sum of all marks at prior level	Mark in same subject, mark in other core subjects (entered separately)
A	100	110	120	130
B	101	111	121	131
C	102	112	122	132
D	103	113	123	133

For the lower secondary school sample, we can only estimate specifications 100, 101, 102 and 103, since we do not have access to prior attainment for the pupils in this sample. For upper secondary school, all specifications are estimated. Results for specifications 110-113 and 120-123 are not reported in detail, but only discussed briefly.

4.1 Estimated school performance indicators - lower secondary school

Table 2 shows some summary statistics for school performance indicators estimated with different sets of SES variables included. We see that adding more SES variables in the conditional attainment model narrows the distribution of the school performance indicators. This implies that the SES variables included on the margin are correlated with individual attainment. This is seen from the significantly higher explanatory power of models with more SES variables included. It also implies that pupils with these SES characteristics are not evenly distributed across schools.

Table 2: Key figures from different specifications, lower secondary school. Attainment measure: Sum of grade points

Specification	Mean of school performance indicator	Standard deviation	Min. value	Max. value	Fraction of individual variance explained
100	43.665	1.87	38.97	50.752	
101	43.801	1.479	39.472	48.964	0.223
102	43.733	1.413	39.623	49.116	0.265
103	43.775	1.386	40.264	49.131	0.338

In terms of explaining individual variation in pupil attainment, parental education is by far the most important set of SES variables. It also matters how parental education is included in the model. A linear specification does not seem to capture all the relevant variation, and it is important (with respect to increase the explanatory power of the model) to include the education of both parents, possibly with an interaction. Adding variables for family structure adds explanatory power on the margin, i.e. when parental education is already controlled for. The single most important factor (which is also included in specification C) is whether parents live together. Inclusion of immigrant status, contributes to the total explanatory power, but less so when parents' education is included. Specification D, in addition to more detailed specification for some variables, adds information on parental earnings and wealth, and on parental unemployment, disability pension and social assistance received by mother and father during the last five years. Of these, the earnings and wealth variables add most to the explanatory power.

Although interesting in its own right, the numbers in Table 2 do not give a complete answer to the question of how school performance indicators estimated from a contextual attainment model differ with respect to how much "context" that is included. The numbers are informative, but they provide no insight into how the results for individual schools are affected by different specifications. The impact of different specifications on the ranking of schools is a key question here. Table 3 reports correlations between estimated school effects from different specifications. We see that just by including parents' years of schooling linearly (101), we obtain school indicators that look significantly different from

unadjusted means (100), as the correlation is 0.85. Including a more detailed specification of parents' schooling, an indicator for parents living together and dummies reflecting immigrant status (102) does not change the picture significantly with respect to which schools that perform well and less well, see Figure 1. But when we use the full set of SES variables, and our most detailed specification, the picture of schools' performance changes significantly, as seen from Table 3 and Figures 2 and 3.

Table 3: Correlations between estimated school effects under different specifications, lower secondary school. Attainment measure: Sum of grade points

	100	101	102	103
100	1			
101	0.85	1		
102	0.83	0.97	1	
103	0.66	0.88	0.93	1

Figure 1: Specifications 101 and 102 compared, lower secondary school. Attainment measure: Sum of grade points

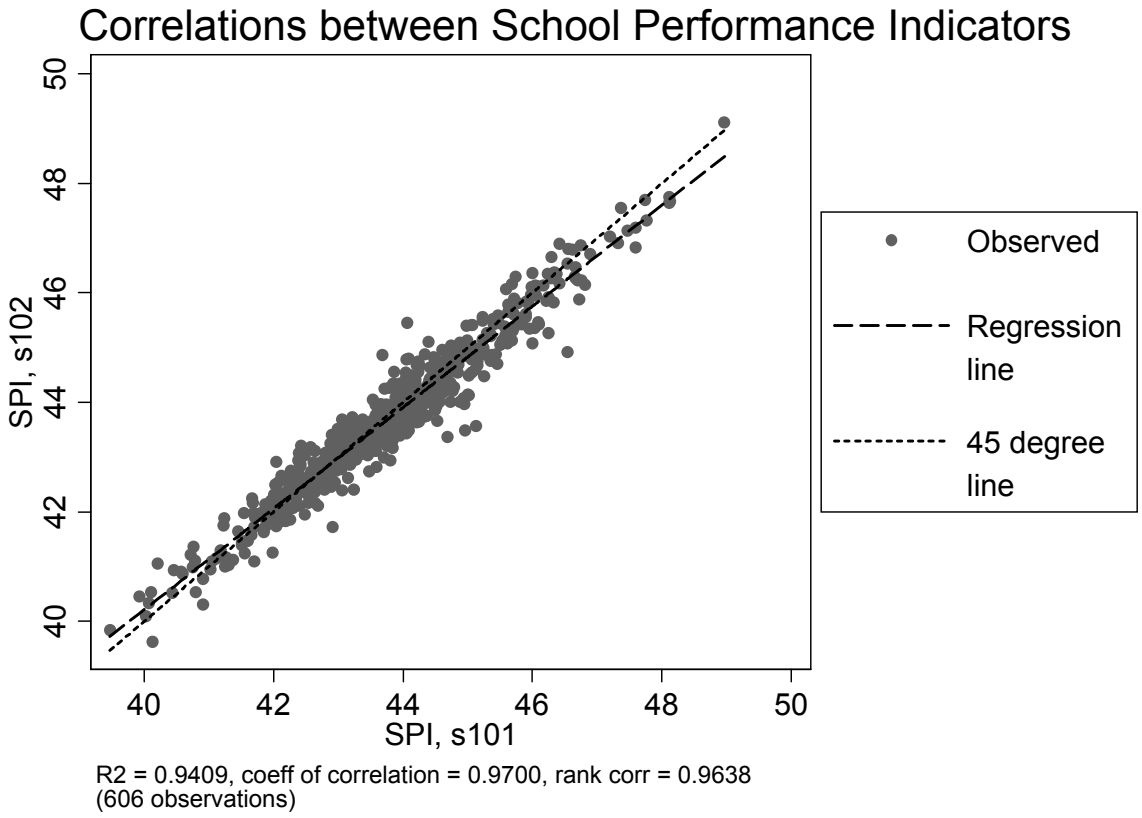

Figure 2: Specifications 101 and 103 compared, lower secondary school. Attainment measure: Sum of grade points

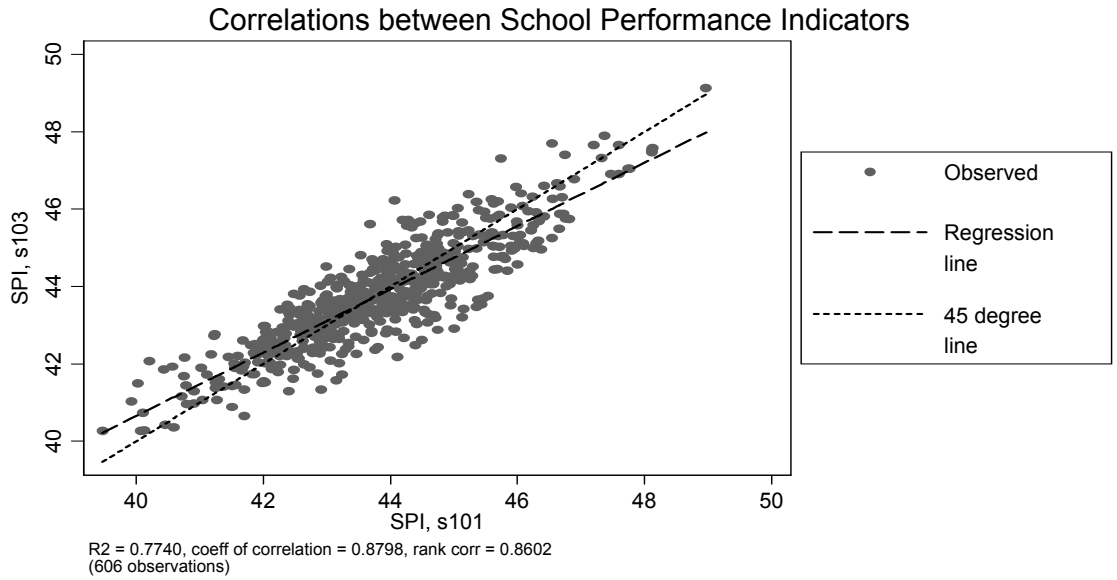

Figure 3: Specifications 102 and 103 compared, lower secondary school. Attainment measure: Sum of grade points

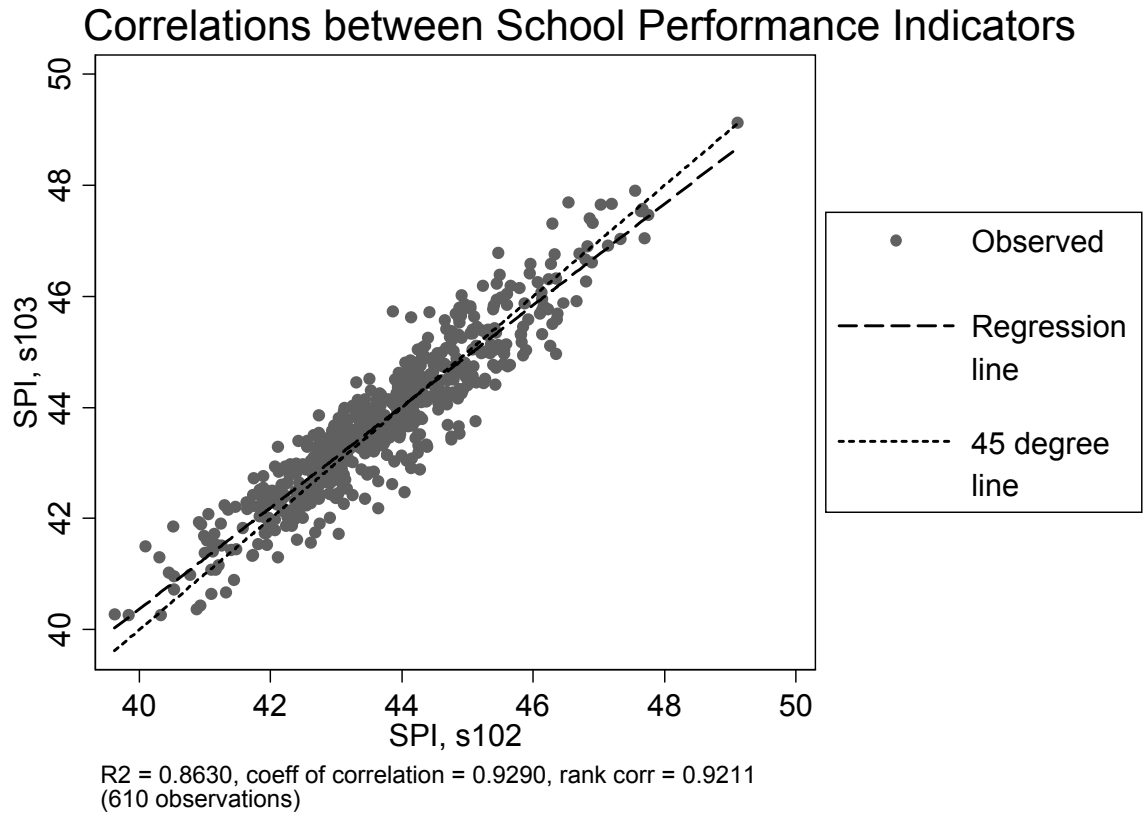

In the figures above, differences for specific schools between school performance indicators based on different models are seen as deviations from the plotted 45-degree line. We see that, even if including a SES variable (or a set of them) adds little on the margin to the total explanatory power of the model, and has little impact on the correlation between indicators estimated from different models, it may have a big impact on the indicators for specific schools. One should bear this in mind when choosing variables to include in the model, and when evaluating results from a model with a limited set of variables included.

Decomposing the adjustment of school results by specification 103 into contributions from different sets of family characteristics, we find that among schools with large downward adjustments, it is parents' education and economic resources that contributes most to the adjustment. These schools have pupils whose parents are particularly well educated and have high income and wealth. The contributions from the other sets of family characteristics are rather modest. Among schools with a large upward adjustment, the pattern is somewhat different. Parental educational attainment is important for the adjustment here as well, but it is interesting to note that the effects of parental unemployment, disability pension and social assistance all have a quite large impact on the results for these schools. An important insight emerges from this exercise: Even if some sets of family characteristics on the margin contribute little or nothing to explaining the *overall* individual variation, they may still be important factors in explaining why some schools perform particularly well or badly according to unadjusted school means.

Tables 4 and 5 show the results of the same exercise for an alternative performance measure, the written final exam. Apart from the explanatory power of SES variables being lower than for the sum of grade points, the pattern is quite similar

Table 4: Key figures from different specifications, lower secondary school. Attainment measure: Written exam

Specification	Mean of school performance indicator	Standard deviation	Min. value	Max. value	Fraction of individual variance explained
100	3.439	0.268	2.746	4.222	
101	3.452	0.218	2.782	4.177	0.163
102	3.447	0.213	2.777	4.2	0.181
103	3.452	0.204	2.855	4.277	0.221

Table 5: Correlations between estimated school effects under different specifications, lower secondary school. Attainment measure: Written exam

	100	101	102	103
100	1			
101	0.93	1		
102	0.93	0.99	1	
103	0.85	0.96	0.97	1

4.2 Estimated school performance indicators - upper secondary school

We now turn to indicators for upper secondary school, where we also have measures of prior attainment. Table 6 shows summary statistics for school performance indicators estimated with different models and different sets of SES variables included. Correlations between the corresponding school performance measures are reported in Table 7. Similar to lower secondary school, adding more SES variables in the conditional attainment model narrows the distribution of the school performance indicators. However, a simple value added model (130) has much larger explanatory power than even the most detailed contextualized attainment model (103). Adding SES variables in the value added models (131,132,133) have only minor effects on explanatory power of the model, and on the estimated school indicators[2]. In fact, adding more indicators on prior attainment, i.e. going from specifications 11x to 12x/13x, has a much larger impact on the estimated indicators. (Marginal R-squared around 0.05, correlations around 0.92.)

Table 6: Key figures from different specifications, upper secondary school. Attainment measure: Written final exam in Norwegian

Specification	Mean of school performance indicator	Standard deviation	Min. value	Max. value	Fraction of individual variance explained
100	3.26	0.321	2.455	4.052	
101	3.278	0.271	2.585	3.888	0.141
102	3.276	0.269	2.595	3.865	0.145
103	3.28	0.258	2.597	3.85	0.161
130	3.365	0.21	2.859	3.889	0.386
131	3.369	0.206	2.86	3.902	0.389
132	3.367	0.205	2.87	3.902	0.389
133	3.369	0.202	2.866	3.885	0.391

Table 7: Correlations between estimated school effects under different specifications, upper secondary school. Attainment measure: Written final exam in Norwegian

	100	101	102	103	130	131	132	133
100	1							
101	0.97	1						
102	0.97	1	1					
103	0.96	0.99	0.99	1				
130	0.7	0.73	0.73	0.72	1			
131	0.7	0.73	0.73	0.72	1	1		
132	0.69	0.73	0.73	0.73	1	1	1	
133	0.69	0.73	0.73	0.73	0.99	1	1	1

Some graphical comparisons between different specifications are provided below. They confirm the pattern from the tables: Given the availability of prior attainment, adding SES variables have only minor

[2] The same result is also obtained when using a smaller sample with more comprehensive performance measures (more subjects) taking on more values than our 1-6 measure.

effects on estimated school performance indicators. However, for some schools the "adjustment effect" from the SES part of the model may still be significant. Using specification 133, the largest adjustment due to SES variables for a single school corresponds one half of a standard deviation in the estimated school indicators. Thus, the lesson from the previous section still holds, although with a more limited impact: Even if some sets of family characteristics on the margin contribute little or nothing to explaining *overall* individual variation, they may still be important fin explaining why some schools perform particularly well or.

Figure 4: Specifications 101 and 103 compared. Upper secondary school. Attainment measure: Written final exam in Norwegian

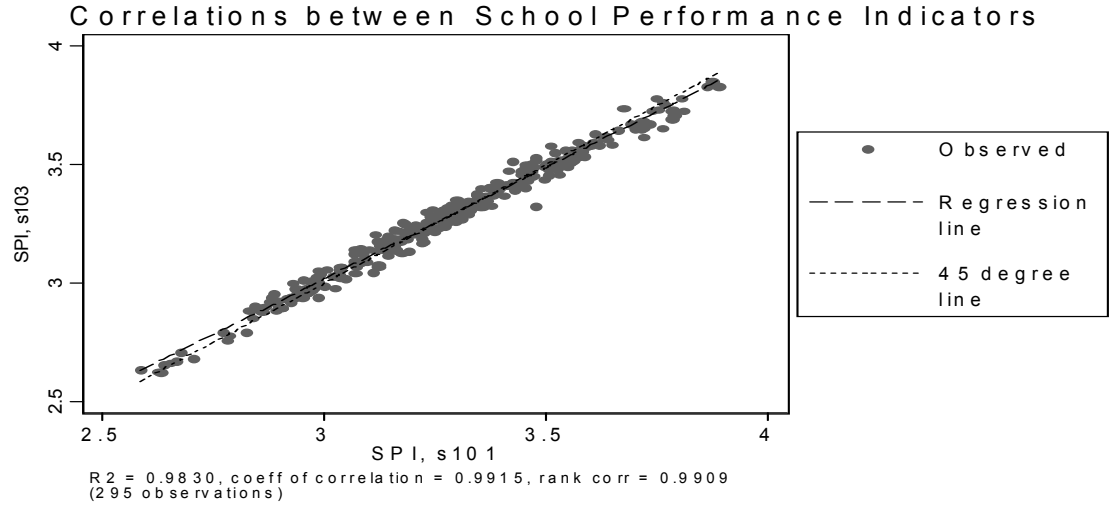

Figure 5: Specifications 130 and 133 compared. Upper secondary school. Attainment measure: Written final exam in Norwegian

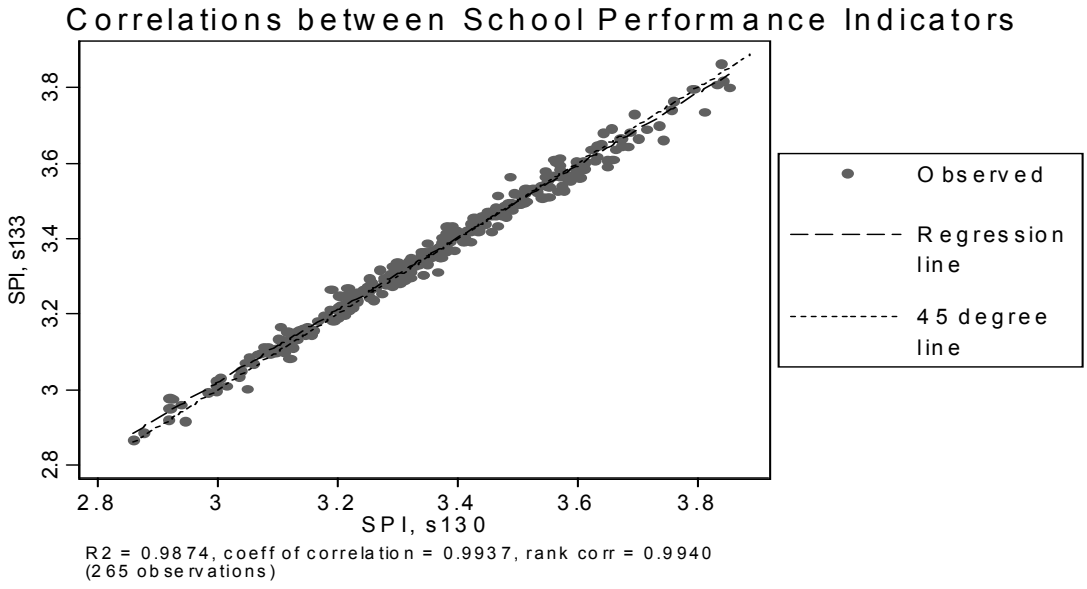

Figure 6: Specifications 101 and 133 compared. Upper secondary school. Attainment measure: Written final exam in Norwegian

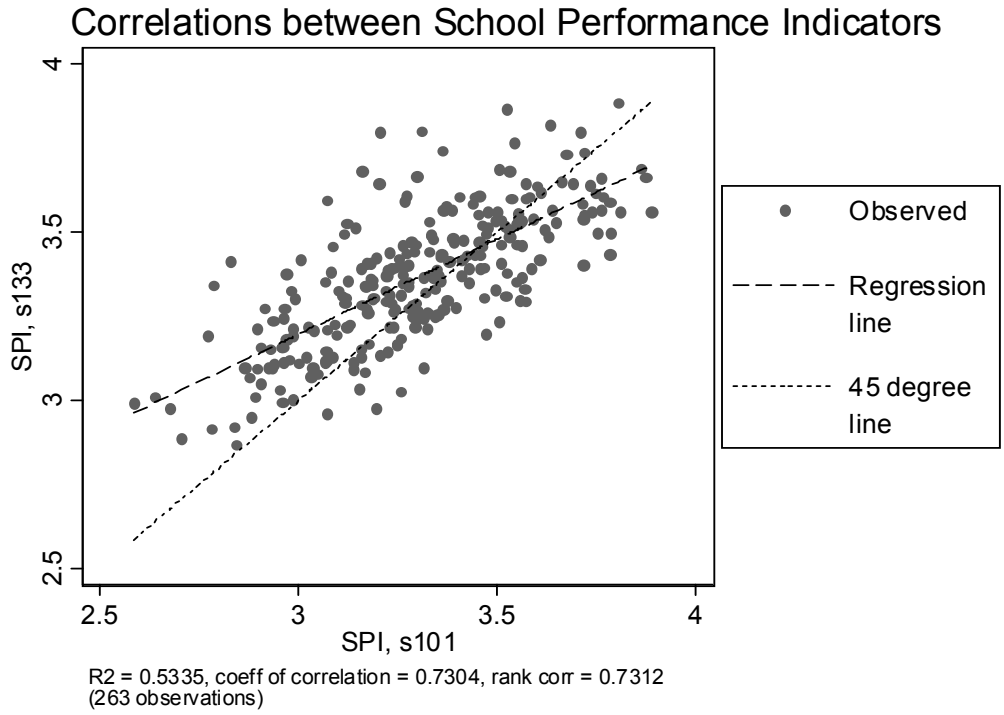

5. Conclusion

This paper has illustrated how school performance indicators may be affected by differences in how socioeconomic status (SES) variables and prior attainment is controlled for. Our findings - on the dataset at hand - indicate that in a contextual attainment model, without control for prior attainment, how SES is included in the model may have large impacts on school performance indicators. Even given a quite rich specification, inclusion of additional SES variables may have quite large impacts on school performance indicators. This implies that the absence of prior attainment measures places great demands on the "contextual part" of the model. In the value-added model, the effects of including (more) SES variables are limited, though they still may have impact for individual schools.

Though indicative, one cannot draw general conclusions from this exercise. The effects of including (more) SES variables in (contextual) value-added models, and of including more SES variables in a conditional attainment model, may vary across levels, years and countries. Given available data, one should always investigate these relations. However, one lesson from this paper may be that, if one has to make a priority, (more) data on prior attainment should be preferred to (more) data on socioeconomic background.

References

Hægeland, T., L.J. Kirkebøen, O. Raaum and K.G. Salvanes (2004): Marks across lower secondary schools in Norway: What can be explained by the composition of pupils and school resources? Rapport 2004/11, Statistisk sentralbyrå.

Hægeland, T., L.J. Kirkebøen, O. Raaum and K.G. Salvanes (2005): Skolebidragsindikatorer. Beregnet for avgangskarakterer for skoleårene 2002-03 og 2003-04. Rapport 2005/33, Statistisk sentralbyrå