

Paying for Performance: The Effect of Teachers' Financial Incentives
on Students' Scholastic Outcomes*

Victor Lavy

The Hebrew University of Jerusalem

July 2002

* Special thanks go to Alex Levkov for outstanding research assistance and for Josh Angrist for helpful comments and suggestions. The 2001 Teachers' Incentive program was funded by the Israel Ministry of Education and administered by the Teaching Staff Division. The views expressed in this paper are those of the author alone and have not been endorsed by the program sponsors.

Paying For Performance: The Effect of Teachers' Financial Incentives on Students' Scholastic Outcomes

Abstract

Performance-related pay for teachers is being introduced in many countries, but there is little evaluation to date on the effects of such programs. This paper evaluates a particular incentive experiment. The incentive program is a rank-order tournament among teachers of English, Hebrew, and mathematics. Teachers were rewarded with cash bonuses for improvements in their students' performance on high-school matriculation exams. Since the schools in the program were not selected at random, the evaluation is based on comparison groups. Three alternative identification strategies are used to estimate the causal effect of the program: a natural experiment stemming from measurement error in the assignment variable, a regression discontinuity method, and propensity score matching. The results of all three methods tell a consistent story: teachers' monetary performance incentives have a significant effect on students' achievements in English and math. No spillover effect on untreated subjects is evident and the general equilibrium impact of the program is positive as well. The program is also more cost-effective than alternative forms of intervention such as extra instruction time and is as effective as cash bonuses for students.

Victor Lavy
The Hebrew University of Jerusalem, Department of Economics
msvictor@mscc.huji.ac.il

Performance-related pay for teachers is being introduced in many countries, amidst much controversy and opposition from teachers and unions alike.¹ The rationale for these programs is the notion that teachers may be motivated by incentive pay. However, there is little evidence on the effect of teachers' incentives at schools. Several dilemmas and challenges arise in the task of designing and evaluating teachers' performance incentives. How should teacher performance be measured? How can individual teachers' contributions be identified? How should the rewards be structured and how generous should they be? Do teachers' effort responds to financial incentives? Are teachers' performance incentives more effective than school-based performance rewards? How relevant and important are the spillover and general equilibrium effects of teachers' incentives? The evidence presented in this paper relates directly to these questions and is based on results of a pay-for-performance experiment among a sample of high-school teachers in Israel, designed to improve their students' achievements on matriculation exams in English, Hebrew, and Mathematics.

This paper evaluates an Israeli program where teachers were rewarded with cash bonuses for improvements in their students' performance on the high-school matriculation exams. The bonus program was structured as a rank order tournament among teachers, in each subject separately.² Thus, teachers were rewarded on the basis of their performance relative to other teachers of the same subjects. Relative performance was preferred over measurements based on absolute performance for two reasons: these awards would stay within budget and there were no obvious standards that could be used as a basis for absolute performance measures. The relative measurements were based on comparison of the achievements of each teacher's students with predicted values using regressions. Two measurements of students' achievements were used as indicators of teachers' performance: the passing rate and the average score on each matriculation exam. The total amount to be awarded in each tournament was predetermined and individual awards were determined on the basis of rank and a predetermined award scale.

The main questions of interest in this experiment relate to the effect of the program on teachers effort and the effect of the experiment on students' achievements. The paper attempts to answer the following key questions: did the program cause teachers to exert more effort, improve their preparation and teaching, and evaluate more effectively the students' need for additional instructional assistance? Did the students' outcomes improve as a result of the program? Did the

¹ Examples of such programs in the USA are performance pay plans in Dade County, Florida, Denver, Colorado, and Dallas, Texas, in the mid-1990s; statewide programs in Iowa and Arizona in 2002; programs in Cincinnati, Philadelphia, and Coventry (Rhode Island); and the Milken Foundation TAP program. In the UK, the government recently concluded an agreement with the main teachers' unions on a new teachers' performance pay scheme to start in 2002/2003, with a budget of nearly £150 million. In New Zealand, the government completed a system wide program of performance-related pay for teachers in 2001. For discussion and analysis of these programs see Clotfeller and Ladd, 1996; Conley and Odden, 1995; Elmore, Abelman and Fuhrman, 1996; Kelley and Protsik, 1996 and Sadowski and Miller, 1996.

² The theory of Individual and group incentives in rank order tournaments are discussed by Lazear and Rosen (1981), Green and Stokey (1983) and Prendergast (1999).

program have spillover effects on students' outcomes in untreated subjects? What general equilibrium effect did the program have?

Although the program was designed as a three-year experiment, schools were not assigned to it at random. Therefore, the search for answers to the foregoing questions is complicated by the possibility that the schools included in the program were a selective sample with attributes that might be related to students' outcomes for reasons other than those related directly to the intervention. Three alternative identification strategies were used to estimate the causal effect of the program. The first strategy is based on a measurement error in the assignment variable that was used to assign schools to the program. The assignment variable, the mean matriculation rate of the school in 1999, was compared with a given threshold (45 percent): any school with a rate equal or below the threshold was assigned to the program. The administrators of the program, unaware that the assignment variable used was measured erroneously, assigned some schools to the program mistakenly. Since the measurement error was random and unrelated to the potential outcome, as will be shown, assignment to the program was actually random in a subsample of the schools. This resulted in a natural experiment or an instrumental variable that could be used to identify the effect of the program in this subsample. This identification strategy is abetted by the use of panel data (before and after the program) that allow an estimation of differences in differences estimates in the "random" assignment subsample.

The second identification strategy is based on the assignment rule that determined program participation. This process was based on a threshold function of the school matriculation rate: schools with a rate equal to or lower than a critical value were included in the program; others were excluded. Since the school matriculation rate varies from year to year, to some extent due to random effects, it is very likely that untreated schools that were just above the threshold resembled schools that were at or below the threshold. The narrower the threshold value of the band, the more likely such a similarity is. We exploited this "regression discontinuity" feature of the assignment mechanism to estimate the effect of the program using the sample of schools within a relative narrow band around a sharply drawn threshold.

The third approach makes use of the very rich data available on all school and students, including many measures of lagged outcomes, to build a comparison group by matching. The methods used ranged from simple regressions that control for observed student characteristics and lagged outcomes, to matching on the propensity score and a non-parametric comparison of treated and control students within cells that are defined by observable school and student characteristics and outcomes.

Section I of this paper provides background information about the Israeli school system, describes the teachers' incentive program, and discusses the theoretical context for pay-for-performance programs. Section II discusses the evaluation strategy. Section III presents the three approaches used for causal identification of the program effect and the empirical results of the effect

of teachers' incentives on the mathematics and English performance of students exposed to treatment obtained with each approach. Section IV presents evidence on the effect of incentives on teachers' behavior in the classroom during the program. Section V discusses the correlation between teacher attributes—such as quantity and quality of schooling, teaching experience, age, gender and parental schooling—and performance in the tournament and the classroom. The results presented in the paper suggest that student achievement improves when teachers are offered financial incentives that reward this. The concluding section presents evidence on the relative effectiveness (cost-benefit) of paying teachers for performance and other interventions, such as school group incentive programs and monetary incentives for students.

I. Tournaments as a Performance Incentive

A. Theoretical Context

Formal economic theory usually justifies incentives to individuals as a motivation for efficient work. The underlying assumption is that individuals respond to contracts that reward performance. However, only a small proportion of jobs in the private sector base remuneration on explicit contracts that reward individual performance. The primary constraint in individual incentives is that their provision imposes additional risks on employees, which is costly to employers through higher wages (Prendergast, 1999). A second constraint is the incompleteness of contracts, which may lead to dysfunctional behavioral responses in which workers emphasize only those aspects of performance that are rewarded. These constraints may explain why private firms reward workers more through promotions and group-based merit systems than through individual merit rewards (Prendergast, 1999).

In education, too, group incentives are more prevalent than individual incentive schemes. The explanation for this pattern, it is argued, lies in the inherent nature of the educational process. Education involves teamwork, the efforts and attitudes of fellow teachers, multiple stakeholders, and complex and multitask jobs. In such a working environment, it is difficult to measure the contribution of an individual member; the group (of teachers) often has better information about its constituent individuals and their respective contributions, enabling it to monitor its members and encourage them to exert themselves or exhibit other appropriate behavior. It is also argued that individuals who have a common goal are more likely to help each other and exert greater effort when a member of the group is absent. On the other hand, standard free-rider arguments cast serious doubt on whether group-based plans provide a sufficiently powerful incentive, especially when the group is quite large.³

Tournaments as an incentive scheme has been suggested initially as appropriate in situations where individuals exert effort in order to get promoted to a better paid position, where the reward associated with that position is fixed and where there is competition between individuals for those

³ See Jenson and Murphy, 1990; Holmstrom and Milgrom, 1991; Milgrom and Roberts, 1992; Gaynor and Pauly, 1990; Kandel and Lazear, 1992; Gibbons, 1998; Malcomson, 1998 and Prendergast, 1999; for a discussion of these issues in the general context of incentives.

positions (Lazear and Rozen, 1982, Green and Stockey, 1983). All that matters for winning in such tournaments is not the absolute level of performance, but how well one does relative to others. Although promotion is not an important career feature among teachers, emphasize on relative rather than absolute performance measures is relevant for a teachers incentive scheme for two reasons. First, awards based on relative performance and a fixed set of prizes would stay within budget. Second, in a situation where there are no obvious standards that could be used, as a basis for absolute performance, relying on how well teachers do relative to others seems a preferred alternative. We therefore used the structure of a rank order tournament for the teachers incentive experiment described below.

B. The Israeli Secondary Schooling

Lavy (2002) presents the results of a group incentive experiment in Israel (1995–1999), in which schools competed on the basis of their average performance and the rewards were distributed equally to all teachers in the winning schools. The purpose of the program was to improve students achievement on the *bagrut* (matriculation) examinations, a series of national exams in core and elective subjects that begins in tenth grade, continues in eleventh grade, and concludes in twelfth grade, when most of the tests are taken. Pupils choose to be tested at various levels in each subject, each test awarding from one to five credit units (hereinafter: units) per subject. Some subjects are mandatory and many must be taken at the level of three units at least. Tests that award more units are more difficult. A minimum of twenty units is required to qualify for a matriculation certificate. About 52 percent of high-school seniors received matriculation certificates in 1999 and 2000, i.e., passed enough exams to be awarded twenty units by the time they graduated from high school or shortly thereafter (Israel Ministry of Education, 2001).⁴

In September 2001, the Ministry of Education inaugurated a new teachers' bonus program in forty-nine Israeli high schools. The main feature of the program was an individual performance bonus paid to teachers on the basis of their own students' achievements. The experiment included all English, Hebrew, Arabic, and Mathematics teachers who taught classes in grades ten through twelve in advance of matriculation exams in these subjects in June 2001. In December 2001, we conducted an orientation with principals and administrators of the forty-nine schools. The program was announced as a voluntary three-year experiment. All principals reacted very enthusiastically to the details of the program. One principal changed his mind later and removed his school from the program. A survey among all participating teachers showed us that 92% percent knew about the program and 80% percent were familiar with the details of how the winners and the size of the bonuses would be determined.

Three formal rules guided the assignment of schools to the program: only comprehensive high schools (having grades 7–12) were eligible, schools must have a recent history of relatively poor

⁴ For more information about the *bagrut*, see Lavy (2002) and Angrist and Lavy (2002).

performance in the Mathematics or English matriculation exams,⁵ and the most recent school matriculation rate must be equal to or lower than the national mean (45 percent). Ninety-seven schools met the first two criteria; forty-nine met the third one.⁶

The initial intention was to limit the program to math and English teachers. Under pressure from the teachers' union, teachers of Hebrew and Arabic were added. Schools were also allowed to replace the language (Hebrew and Arabic) teachers with teachers of other core matriculation subjects (Bible, history, literature, or democracy). The schools and teachers were informed on December 20, 2001, about their participation in the program.

C. The Israeli Policy Experiment

Each of the four tournaments (English, Hebrew and Arabic, math, and other subjects) included teachers of classes in grades 10–12 that were about to stand for a matriculation exam in one of these subjects in June 2001. Each teacher entered the tournament as many times as the number of classes he or she taught and was ranked each time on the basis of the mean performance of each of his/her classes. Teachers were ranked in view of their classes' passing rate and mean score. Ranking was based on the difference between the actual outcome and a value predicted on the basis of a regression that controlled for the students' socioeconomic characteristics, the level of proficiency in each subject, and a school fixed effect. Separate regressions were used to compute the predicted passing rate and mean score, and each teacher was ranked twice, once for each outcome. The school submitted student enrolment lists with itemizations by grades, subjects, and teachers. The reference population was the enrollment on January 1, 2001, the starting date of the program. All students who appeared in these lists (including dropouts and students who did not take the June 2001 exams, irrespective of the reason) were included in the class mean outcomes at a score of zero.

All teachers who had a positive residual (predicted outcome less actual outcome) in both outcomes were divided into four ranking groups, from first place to fourth. Points were accumulated according to ranking: 16 points for first place, 12 for second, 8 for third and 4 for fourth. The program administrators gave more weight to the passing rate outcome, awarding a 25 percent increase of points for each ranking (20, 15, 10, and 5, respectively). The total points in the two rankings were used to rank teachers in the tournament and to determine winners and awards, as follows: 30–36 points—\$7,500; 21–29 points—\$5,750; 10–20 points—\$3,500; and 9 points—\$1,750. These awards are significant relative to the mean gross annual income of high-school teachers (\$30,000) and the fact

⁵ Performance was measured by the average passing rate in the math and English matriculation tests during the last four years (1996-99). Two occurrences or more of any of these rates being lower than 70 percent was considered a poor performance. English and math were chosen because they are the subjects with the higher failing rate among the matriculation subjects.

⁶ A relatively large number of religious and Arab schools met all the three selection rules. To keep their proportion in the sample close to their population share, the matriculation threshold for these schools was set to 43 percent.

that a teacher could win several awards in one tournament if he or she prepared more than one class for a matriculation exam.⁷

The program included 629 teachers, of whom 207 competed in English, 237 in mathematics, 148 in Hebrew or Arabic, and thirty-seven in other subjects that schools preferred over Hebrew. Three hundred and two teachers won awards—94 English teachers, 124 math teachers, 67 Hebrew and Arabic teachers and 17 among the other subjects. Three English teacher won two awards each; twelve math teachers won two awards each, and one Hebrew teacher won two first place awards totaling \$15,000.

A follow-up survey of teachers in the program was conducted during the summer after the end of the school year. Seventy-four percent of teachers were interviewed and there were very few refusals. Most absences among potential interviewees were due to wrong phone numbers and the teachers' being overseas. The survey results show that 92 percent of the teachers knew about the program, 80 percent had been briefed about its details—almost all by the school principal and the program coordinator—and 75 percent thought that the information was complete and satisfactory. Almost 70 percent of the teachers were familiar with the award criteria and about 60 percent of them thought they would be among the award winners. Only 30 percent did not believe they would win; the rest were certain about their chances. Two-thirds of the teachers thought that the incentive program would lead to an improvement in students' achievements.

II. Evaluation Strategy

The first evaluation issue to address is the non-random selection of schools and, therefore, of teachers for the program. In the next section, we suggest three strategies meant to surmount this difficulty and help to identify the causal effect of the program. The evaluation may include English and math teachers only because school participation in Hebrew and Arabic was optional and all schools had the choice of replacing these subjects with other core matriculation subjects. Since some schools did exercise this option, the sample of schools that elected not to do so is endogenous.

A second issue of concern relates to the implications of the teachers' potential strategic behavior. In other words, the teachers' increased investment of time and effort, due to the incentives offered them, may prompt students to reallocate their time and effort toward the rewarded subjects at the expense of other subjects. Hence, the program may have an adverse effect on outcomes in subjects other than those rewarded. By implication, we should estimate the effect of the program on the students' overall outcomes and not only on the treated subjects. For the sake of simplicity, let us separate students' educational outcomes into awarded subjects (Y1) and other subjects (Y2). If students invest less time and effort in Y1, then the effect on Y2 may be negative. However, the additional instruction time in the treated subjects may free some of the students' time for other

⁷ For more details, see Ministry of Education, High School Division, "Individual Teacher Bonuses Based on the

subjects, so that the effect on Y2 may be positive. Estimating the effect of treatment by examining the change in the rewarded subjects only may overstate or understate the treatment effect of the program, if there are indeed negative or positive spillover effects due to change in time and effort allocation.

We may address this general equilibrium aspect of the evaluation by estimating the effect of the program on the outcomes of all other “untreated” subjects. However, the number of subjects that we may use as truly untreated is limited, for two reasons. First, students are tested in many different subjects at the end of twelfth grade and the sample size in some of the tests is very small. Second, as we will recall, schools were allowed to include in the program teachers of one other subject in lieu of Hebrew or Arabic. Some school exercised this option and included subjects such as Bible and social studies, which basically excluded these subjects from being considered untreated. Instead of estimating the effect of the program on each subject, we may generate a summary measurement that takes account of all the exams, e.g., the total number of credit units earned on all the exams. Another possibility is to confine the focus to untreated subjects that had the largest sample size. Using this criterion, two subjects, history and biology, stand out. A third alternative is to estimate the effect of the program on the students’ matriculation status. This is an overall high school achievement measurement that encompasses the outcomes of all matriculation tests from tenth grade through the end of twelfth grade. Below we report results using all three alternatives: evidence of the effect of the program on the biology and history outcomes, evidence of the effect on total credit units accumulated during the program period in untreated subjects, and evidence of matriculation status.

A third issue to address is how to measure the outcome variables. In each of the treated subjects, the requirements often include several exams. For example, a student who takes mathematics at the proficiency level of three units has to take two exams, one for the first unit and the second for the other two units. In some subjects there are additional exams, such as a lab test in science subjects and an oral test in languages (English, Hebrew, or other). The final score is a weighted average of all components, the weights reflecting the credit units of each component. In this study, we used three different but related measures of outcomes for each subject: the number of tests taken by a student in the given subject, the total number of units in the tests attempted, and the total units earned. The second and third measurements reflect the proficiency level of the curriculum. Below we estimate the effect of the program on these three outcomes in every treated subject and in biology and history. We will also use as an outcome the total number of units accumulated by each student in all untreated subjects.

III. Identification Strategies, Estimation and Results

A. Natural Experiment: Random Measurement Error in the Assignment Variable

The program rules limited assignment to schools with a 1999 matriculation rate equal to or lower than 45 percent (43 percent for religious and Arab schools). However, the matriculation rate used for assignment was an inaccurate measure of this variable. The matriculation-rate data given to administrators were culled from a preliminary and incomplete file of matriculation status. For many students, matriculation status was erroneous since it was based on missing or incorrect information. The Ministry later corrected this preliminary file, as they do every year.⁸ As a result, the matriculation rates used for assignment to the program were inaccurate in a majority of schools. The measurement error is useful for identification of the program effect. In particular, conditional on the true matriculation rate, program status is virtually randomly assigned by mistakes in the preliminary file.

Figure 1 presents the relationship between the correct matriculation rates and those erroneously measured for a sample of 507 high schools in Israel in 1999.⁹ Most (80 percent) measurement errors were negative, 17 percent were positive and the rest had no error. The deviations from the 45-degree line do not seem to correlate with the correct matriculation rate. This may be seen more clearly in Figure 2, which demonstrates that the measurement error and the matriculation rate do not co-move; their correlation coefficient is very low, at -0.085 , even though the p-value that it is different from zero is 0.055. However, if a few extreme values (five schools) are excluded, the correlation coefficient becomes basically zero. Although the figure may suggest that the variance of the measurement error is lower at low matriculation rates, this is most likely due to the floor effect that bounds the size of the negative errors: the lower the matriculation rate, the lower the absolute maximum size of the negative errors. Similar evidence arises when the sample is limited to the ninety-seven schools that were eligible for treatment, those from which forty-nine schools were assigned for treatment (Figures 3 and 4). If the two extreme values in Figure 4 are excluded from the sample, the estimated correlation coefficient between the correct 1999 matriculation rate and the measurement error rate, although negative, is practically zero. Similar evidence is observed when the sample is limited to schools with a matriculation rate higher than 40 percent. In this sample, the problem of the bound imposed on the size of the measurement error at schools with low matriculation rates is eliminated (Figure 4A).

The measurement error can be used for identification either as a natural experiment, where treatment is assigned randomly in a subsample of the ninety-seven-school sample, or as an instrumental variable. Seventeen of the forty-nine treated schools had a correct 1999 matriculation

⁸ Matriculation status depends on the fulfillment of many requirements (e.g., a minimum number of credit units and the coverage of compulsory subjects such as math, foreign language, and Bible) that tend to vary by school type (technical, agricultural or regular, Jewish or Arab, religious or nonreligious) and level of proficiency in each subject. The verification of information between the administration and the schools is a lengthy process. The first version of the student file that includes the results of the matriculation exams becomes available in October of every year (for the cohort that graduated in June of that year). However, it is updated continuously and the final version is not completed until late December of the same year.

⁹ The sample was limited to schools with positive ($> 5\%$) matriculation rates.

rate above the threshold line. Thus, these schools were “erroneously” chosen for the program. For each of them, there might have been a school with a similar matriculation rate but with a random measurement error not large (and negative) enough to drop it below the assignment threshold. Figure 6 shows this pairing. The drawn ellipse circles the treated schools and their matching counterparts. There are twelve such ellipses. Within this sample (twenty-nine schools) treatment assignment was random, as shown above. Therefore, the twelve untreated schools may be used as a control group that reflects the counterfactual for identification of the effect of the program. The treated schools in this sample, however, are not a random sample culled from the sample of all treated schools, as may be seen clearly in Figure 5. For example, the correct 1999 matriculation rate is 45 percent or higher for all schools in this sample, while many schools in the full sample have correct matriculation rates that is lower than 45 percent. Another point of concern is the small number of treated schools included in this sample, only 35 percent of all treated schools. We should bear this in mind when interpreting the results.

Table 1 presents the pre-program (2000) and post-program (2001) means of students (those graduating twelfth grade) and school characteristics for the seventeen treated schools and the twelve control schools. The student background characteristics include father’s and mother’s schooling, number of siblings, and dummies for gender and for immigration status. The treatment-control differences and standard errors in these variables (columns 3 and 6) reveal that the two groups are very similar in both years in all background characteristics and in no case statistically different. The only non-identical variable is number of siblings, and in 2001 the difference in the number of siblings was surprisingly large.

The second panel in Table 1 presents students’ outcomes in the form of units earned before twelfth grade in treated subjects (English, math) and untreated subjects (biology, history). For twelfth graders in the treatment year (2001), this measure reflects pre-program (or lagged) outcomes. No significant treatment-control differences are observed in English and math, in either year. Some differences are observed in history but not in biology or in total units. The differences in history are evident in both years, but they probably reflect differences among schools in the timing of the history exam (in eleventh or twelfth grade), which is left to the discretion of the school.

The third panel in Table 1 compares the school-level covariates. Treatment and control are balanced in terms of religious status but not in terms of nationality, since there are no Arab schools in the control group. The 1999 mean matriculation rate is almost identical in the two groups, an unsurprising result since this school-level outcome was used for matching. A similar balance is found in the groups’ 2000 matriculation rates.

The evidence in Table 1 suggests that, generally speaking, the treatment and control schools are well balanced in most individual and school characteristics, reflecting the basic similarity of the two sets of schools. These findings support the notion of a natural experiment as a strategy to identify the effect of the teacher bonuses program in this sample of treated schools. Nevertheless, it is still

necessary to control for all these variables in the estimation to net out the effect of any remaining differences. In particular, we should control for the true 1999 matriculation rate. The evidence about the effect of the program will be based on regressions for twelfth graders in 2000 and 2001 (stacked panel data) that include pre-program students, school covariates, and constant school effects. The estimated treatment effect will be equivalent to differences in differences estimate embedded in a natural experiment setting. The constant school effects will absorb any remaining permanent differences, observed and unobserved, between the treated and the control schools. The estimation framework is discussed at greater length in Section III.d below.

Estimation

The following model was used as the basis for regression estimates:

$$(1) \quad Y_{it} = \alpha_j + x_i' \beta + Z_{it}' \gamma + \delta T_{it} + \varepsilon_{it}$$

where i indexes pupils; t indexes years 2000 and 2001, and T_{it} is the assigned treatment status. The model includes a vector of student-level covariates x_{it} , and a vector of school-level covariates Z_{ijt} . We estimate the regressions on the basis of pooled data from both years (the two adjacent cohorts), stacked as school panel data. This data structure allows the inclusion of school fixed effects in the regression:

$$(2) \quad Y_{ijt} = \alpha_j + x_{ij}' \beta + Z_{ijt}' \gamma + \delta (T_{ijt} \times D_{2001}) + \varepsilon_{ijt}$$

where j indexes schools. This model also includes a constant effect for year 2001 (D_{2001}) and the treatment indicator in this model is equal to the interaction between a dummy for treated schools and a dummy for year 2001 ($T_{ijt} \times D_{2001}$). The estimated treatment effect in this model is a difference in differences estimate. Its advantage is that it nets out any correlation between the outcome variable and any school characteristic that did not change between 2000 and 2001. This model is estimated with two samples, the measurement-error sample and the regression discontinuity sample. The propensity score sample is used to estimate a model such as Equation (1) for year 2001.

Results

The first panel rows in Table 2 present the evidence from regressions using the stack panel data for a sample of twelfth-grade students in 2000 and 2001. The treatment indicator equals 1 for year-2001 students in the seventeen treated schools. The regressions include, as controls, individual covariates (gender, father's and mother's education in terms of years of schooling, ethnicity, and subject-specific and total matriculation units earned before treatment), school covariates (a dummy for religious schools and a dummy for Arab schools and school fixed effects), and a dummy variable for year 2001. The treatment effect is estimated with two alternative specifications, with and without the correct 1999 matriculation rate. Treatment effect estimates are presented for three outcomes: number of exams attempted, number of units attempted, and number of units earned in English and math. The

standard errors reported in the table are adjusted for clustering, using formulas set forth in Liang and Zeger (1986).

The treatment effect in English and math, for all outcomes, are positive but vary in degree of precision. In math, all three outcomes are significantly different from zero but in English only the estimated treatment effect on awarded units is large relative to its estimated standard error. There is more room for improvement in units earned than in the first two outcomes because units earned can be increased by taking more exams or changing the curriculum, while the other outcomes may be improved by raising the score from fail to pass in any given test. Indeed, the estimated relative effects on units earned are greater than on the other two outcomes. The effect of treatment on units earned in math is 0.256, a 19 percent improvement relative to the mean of the control schools (1.35). The effect of treatment on awarded credit units in English is 0.361, a 21 percent improvement relative to the mean of the control schools (1.7). The relative improvement in units attempted is much lower—6.7 percent in math and 9.7 percent in English.

The causal interpretation of the above results are based on the program status being randomly assigned by the measurement error, conditional on the true 1999-matriculation rate. Indeed the treatment effect estimates are sensitive to the exclusion of the correct 1999 matriculation rate as a control. Without this control the math treatment estimate is much lower, 0.163 versus 0.256, and it is also less precisely estimated. The English treatment estimate is also lower but only marginally.

Table 2, lower panel, presents also the treatment effect estimates for untreated subjects, biology and history. None of the estimates is significantly different from zero. In fact, the estimates in all three outcomes are negative in math and positive in history, but practically speaking, given the large estimated standard errors, these point estimates reflect chance deviations from zero. We also estimated the treatment effect on all units attempted and earned *other than* those in the treated subjects. The effect on total units attempted was 0.063 (s.e.= 0.318); the effect on total units earned was 0.282 (s.e.=.257). Both estimates are not significantly different from zero. We may regard these results as a test of the quality of identification of the program's effects on math and English outcomes. We may also think of these results as evidence that there were no spillover effects.

To assess the overall general equilibrium effect of the program we also estimated the treatment effect on the matriculation rate. This estimated is 0.019 (s.e.=0.011). This effect is marginally significant and it represents a 5 percent improvement relative to the mean of the control schools.

B. Identification based on Discontinuity in the Assignment Variable

Since the rule governing selection to the program was based simply on a discontinuous function of a school observable (the erroneously measured 1999 matriculation rate), the probability of receiving treatment changes discontinuously as a function of this observable. This discontinuity in the treatment assignment mechanism may be exploited as a second source of exogenous identification information

for evaluation of the effects of the teachers' bonuses program.¹⁰ The discontinuity in our case is a sharp decrease (to zero) in the probability of treatment beyond a 45 percent school matriculation rate for nonreligious Jewish schools and beyond 43 percent for Jewish religious schools and Arab schools. The time series on school matriculation rates show that the rates fluctuate from year to year for reasons that transcend trends or changes in the composition of the student body. Some of these fluctuations are random. Therefore, schools with a rate above but close to threshold may be similar to treated schools with a rate close to but below this threshold. The degree of similarity probably depends on the width of the band around the threshold. Sample size considerations exclude the possibility of a bandwidth lower than 10 percent and a wider band implies fluctuations of a magnitude that is not likely to be related to random changes. Therefore, a bandwidth of about 10 percent seems to be a reasonable choice in our case.

There are twelve untreated schools with matriculation rates in the 0.46–0.52 range and fourteen treated schools in the 0.40–0.45 range. The 0.40–0.52 range may be too large, but we can control for the value of the assignment variable (the mean matriculation rate) in the analysis. We should also note that there is some overlap between this sample and the measurement-error sample: nine of the fourteen treated schools and five of the twelve control schools belong to the group of treatment and control schools, respectively, of the measurement-error sample.

Table 3 presents descriptive statistics for the treated and control groups in the discontinuity sample for the cohort that graduated in the year before treatment (2000) and after treatment (2001). The treatment-control differences and standard errors in the background variables (Columns 3 and 6) reveal that the two groups are very similar in both years in all background characteristics except the ethnicity variable. The proportion of treated students of African-Asian origin is lower in treated schools than in control schools; this difference is significant in 2000 but not in year 2001.

The second panel in Table 3 shows that before the program began students in treated schools earned fewer math units than students in control schools. This gap is evident in both years: the treatment-control difference is -0.339 in 2000 and -0.312 in 2001. The difference in 2000 is significantly different from zero; the difference in 2001 is marginally so. A discrepancy in the same direction, although not significant in either year, is observed in total lagged units. In English, however, the opposite is observed: a positive and significant difference in English units in favor of the treated schools in both years, although the difference is significant only in 2000. No significant treatment-control differences are observed in biology or history.

The third panel in Table 3 reveals a statistically significant treatment-control gap in the erroneously measured 1999 matriculation rate and a similar gap in the correct rate. The gap carries the expected sign, negative, because all treated schools had erroneously measured matriculation rates

¹⁰ Regression discontinuity designs are described by Campbell (1969), and were formally examined as an identification strategy recently by Hahn, Todd, and van der Klaauw (2001). For recent examples, see Angrist and Lavy (1999, 2002a), Lavy (2002), and van der Klaauw (1997).

below the threshold value and all control schools were above the threshold. Given that all measurement errors in the discontinuity sample were negative, we should expect the treatment-control difference in the correct matriculation rate to be negative as well. The two differences are of similar magnitudes—0.061 and -0.054—and both had low standard errors.

The evidence presented in Table 3 suggests that the treatment and control schools are balanced in some individual and school characteristics but differ in pre-program matriculation rates and in the achievement outcomes in the treated subjects. Although the differences are small, they should be controlled for directly in the empirical analysis. Since the measured differences may reflect other unmeasured differences, identification based on the discontinuity approach depends more than in the case of the natural experiment (measurement error) on the school constant effects model for estimating the treatment effect, since this model accounts for all unobserved but fixed correlates of potential outcome.

Results

We used the discontinuity sample to estimate models identical to those estimated with the measurement-error sample. The results are presented in Table 4. The treatment effect estimates are very similar, qualitatively and quantitatively, to the results obtained using the measurement error/natural experiment approach and sample. The treatment effect estimates are positive for all English and math outcomes. The estimated effect on earned units in math is 0.244 (s.e.=0.109), almost identical to the estimate obtained with the measurement-error sample. The estimated effect on English units earned is 0.177, about half the estimate derived from the measurement-error sample and quite imprecise (s.e. = 0.145). However, the effect on the number of exams and units attempted in English is positive and significant in this sample. The relative sizes of the effects resemble those derived from the measurement-error sample because the mean of the outcome variable for treatment and control is similar in both samples. The treatment estimates are not sensitive in this sample to the exclusion of the true 1999 matriculation rate as a control variable.

The evidence in the second panel of Table 4 indicates that treatment has no effect on history and biology outcomes in the discontinuity sample as well. The coefficients are all positive in history and all negative in biology, and in all cases the estimates are not significantly different from zero. We also estimated the effect of treatment on all units attempted and earned *other than* those in the treated subjects. The effect on total units attempted was 0.286 (s.e.= 0.334); the effect on total units earned was 0.166 (s.e.=.283). Neither estimate is significantly different from zero. The estimated treatment effect on the matriculation rate in this sample is 0.023 (s.e.= 0.012). These results may be viewed as supporting evidence for the causal interpretation of the estimated effects on math and English outcomes, no spillover effect on untreated subjects and overall positive effect on matriculation achievements of students.

C. Matching Observable Characteristics and Lagged Outcomes

The third method we use for identification is matching. It is based on the assumption that one may account for all differences between treated and untreated subjects by controlling for observable characteristics. Matching may be implemented non-parametrically by defining cells using discrete characteristics.¹¹ The more characteristics, however, the harder it becomes to find untreated individuals who are identical to treated individuals. Rosenbaum and Rubin (1985) suggest a solution to this dimensionality problem: a weighted index of each individual's characteristics, referred to as the "propensity score." In our case, the propensity score is calculated for each student using a regression of student and school characteristics on treatment status. Students are then matched according to their propensity score.¹² Students for whom no suitable match can be found are dropped to ensure that the comparisons between treated and control students take a place over a range of characteristics in which suitable comparison do exist (Dehejia and Wahba, 1998).

Matching by propensity score allows us to trade off different observable characteristics against each another according to their importance, in order to find the best match for matched students among the untreated students. A good match on observable characteristics does not necessarily ensure a good match on important unobservables such as ability and motivation. Our data, however, include many lagged achievement outcomes that, if included in the propensity equation, may improve the match on important unobserved individual attributes because they are likely to correlate. Therefore, our propensity equation uses individual and school characteristics and all lagged matriculation outcomes. In estimating the treatment effect by the matching method, we derive corrected standard errors by using numerical bootstrapping methods.

The advantage of the matching method over the first two methods described above is that it uses all forty-nine treated schools and searches for matches in all other 520 high schools countrywide. Before discussing the results of the match, however, we should address an important point. About 82 percent of twelfth-grade students in treated schools were enrolled in math classes during the treatment year and about 88 percent were enrolled in English classes. Not being enrolled in either subject means one of two alternatives: either the student completed his or her matriculation studies and exams in English or math by mid-semester of the senior year (by December 2000, before the program began), or the student dropped out of English or math class before December 2000.¹³ The data do not allow us to distinguish reliably the two alternatives. However, each alternative has different implications for the measurement of lagged outcomes. To avoid these complication and to focus on treated students, we can include as treated students only those who were enrolled in English or math classes and searched for their matches in all other schools countrywide. Alternatively, we can search for matches

¹¹ For an application of this method, see Angrist (1998).

¹² See Dehejia and Wahba (1998) for a recent application of the propensity score method to an evaluation in economics.

for all students irrespective of their enrollment in English or math. Therefore, in comparing the results of the matching method with those of the other two methods, we should bear in mind that the estimation samples we used in the two last-mentioned approaches included all twelfth graders irrespective of their enrollment in English or math.

Table 5 presents descriptive statistics for the treated and control groups in the propensity score sample that is based on students enrolled in English and math classes. The English and math samples are different, reflecting the fact that the treated samples in the two subjects were different. All variables that appear in the table were used in the matching equation. Matches were found for all treated students, from 330 schools in the English sample and 350 schools in the math sample.

The first panel in Table 5 shows that the matching leads to a perfect match in all demographic variables in the English and math samples except for immigrant status. None of treatment-control differences in the background variables is statistically different from zero except for immigrant status, as noted.

The second panel in Table 5 shows that the treated students and their matched samples in both subjects are perfectly balanced in all six measures of lagged outcomes. This is evident not only in each of the specific treated subjects, English and math, but also in the lagged outcomes of the untreated subjects, history and biology, and the lagged total units and mean score in all untreated subjects. These results concerning lagged outcomes reinforce our confidence that the two samples are well balanced in terms of unobserved student covariates. The third panel in Table 5 reveals no statistically significant treatment-control differences in the school covariates that were used in the propensity score equation.

Overall, the evidence presented in Table 5 suggests that the treatment populations and their matched samples are well balanced in all dimensions of comparison. The only covariate that was not perfectly balanced was immigrant status. In estimating the treatment effect, we will control for this variable as well as for all other variables reported in Table 5.

Table 5A presents the comparison between the treatment and the control group that is based on matching for all students irrespective of enrollment in English or math classes. The results are very similar to those presented in table 5A: the two samples are well balanced in all variables except the immigration status. We will estimate the program effect in the two matched samples, the one based only on treated students and the one based on all students.

Results

Table 6 presents the results of estimating the treatment effect with the propensity score matched sample of the students enrolled in English and math classes. The treatment effect estimates in English and math are positive; all are significantly different from zero and resemble in magnitude the

¹³ A third possibility is that the student was erroneously deleted from the English and math class rosters that the

estimated treatment effects obtained by use of the other two methods, as reported in Tables 2 and 4. Focusing for comparison on the effect of treatment on credits earned, the estimate for math in the match sample points to a 19.2 percent effect relative to the mean of the control group, which is identical to the relative estimated size effect in the measurement-error sample (19 percent) and the discontinuity sample (16.3 percent). The relative effect on English credits earned is 7 percent in the matched sample, close to the 8.9 percent estimate in the discontinuity sample but smaller than the 21 percent estimate in the measurement-error sample. However, the estimated relative effect on attempted units in English is identical in all three methods, at about 9 percent relative to the control mean.

The evidence in the second panel of Table 6 indicates that treatment has no significant effect on any of the history or biology outcomes. For example, the estimated treatment effect is 0.020 (s.e. = .046) on history units earned and 0.024 (s.e. = .019) on biology units. The effects on attempted exams and units in history and biology are estimated more precisely but are still not significantly different from zero at a conventional significance level. The effect on total untreated units attempted or earned, although not shown in Table 6, is also not significantly different from zero. The treatment effect on the matriculation rate in this sample is 0.032 (s.e.=.013).

The lower rows in the first panel in Table 6 presents the evidence from the matched sample that includes all students irrespective of their enrollment in English and math classes. The results are qualitatively similar to those obtained from the sample of students enrolled in English and math classes though the treatment effect estimates are somewhat lower. However, when the point estimates are weighted by the proportion of students who were treated (0.88 in English and 0.8 in math), the differences between the estimates in panel c and d are much smaller.

IV. Financial Incentives and Teachers' Effort

The evidence in the previous section shows clearly that the teachers' incentive program led to significant improvements in students' achievements in English and math. How closely do these improvements correspond to changes in teachers' behavior as a result of the program? Do they reflect greater effort on teachers' part or changes in teaching methods due to the program? To address these questions, a telephone survey was conducted among the English and math teachers who participated in the program. For comparison purposes, a similar survey was conducted with a similar number of nonparticipating English and math teachers. The comparison group was chosen for practical and logistical reasons and not necessarily because it was an appropriate comparison group vis-a-vis the treated schools. However, as Table A1 in the Appendix shows, the characteristics of the teachers in the two groups are very similar. For example, the treated and control teachers of English are identical

schools submitted before the program began. This possibility can be neither verified nor excluded.

in age (forty-five), gender (81 percent female), and schooling. Similar results are observed for math teachers.

Table 7 presents evidence about the effect of the incentive program on three behavioral outcomes of participating teachers: teaching methods, teachers' effort, and focusing of effort on weak or strong students. The evidence, shown for English and math teachers separately, points to two patterns: the program modified teaching methods and led to a major increase in teachers' effort, expressed in the form of time devoted to student instruction beyond regular classroom time, especially in the weeks before the matriculation exams.

The mean of English teachers who taught students in small groups is 63 percent in the sample and 8.5 percentage points higher among teachers who participated in the program. Fifty-eight percent of teachers in the sample used individualized instruction, as against 69 percent of teachers in the program. Program teachers used tracking by ability much more than the comparison group of teachers, 42 percent versus 62 percent. Ninety-three percent of the control teachers reported having adapted their teaching methods to their students' ability; 100 percent of the treated teachers so reported.

Among math teachers, the behavioral change focuses on additional instruction time as opposed to teaching methods. The only difference in math teaching methods is in the prevalence of tracking by ability: 53 percent of program teachers as against 40 percent of other teachers. Most math teachers apparently invest additional time in teaching, beyond regular scheduled classroom hours, throughout the year. However, treated teachers added much more—4.8 hours per week as against 2.7 hours by other teachers, almost double the effort. The induced and intensified effort of treated math teachers is even greater as the matriculation exam period approaches: treated teachers begin special preparations with their students seven weeks before the exam date as against 4.5 weeks in the case of other teachers. English teachers who participated in the program also made an additional effort, but in only one dimension: preparing students for matriculation exams. Fifty percent of them report making such an effort, as against 37 percent of teachers in the comparison group.

Most teachers target the additional effort to all their students and to their weakest students. However, targeting of effort to weakest students is more prevalent among English teachers: 33 percent of those in the program versus 20 percent of other teachers. Among program math teachers, we see more effort targeted toward average students.

V. Does Tournament Ranking Correlate with Teachers' Characteristics?

The results presented in this paper about the effect of teachers' incentives on students' achievements prompt us to conclude that individual teachers matter in improving schooling quality. Can we predict who the better teachers in our sample would be by some conventional measure of teacher quality? The ranking of teachers in the tournament was based on their students' average residuals, as determined from a regression of students and class characteristics on test scores. Teachers with positive residuals

in both mean score and pass/fail regressions won awards. We can try to characterize the good teachers by seeking correlations between conventional teacher characteristics (shown in Table A1) and the teachers' residuals. We examined these correlations within a regression framework for English and math teachers separately.

The results support the view that we do not know how to measure teaching quality on the basis of conventional teacher characteristics such as age, gender, education, teaching certification, and years of teaching experience.¹⁴ None of these variables was highly significant in the achievement residual regressions. Other variables, however, evinced significant correlations in the regressions. Being born and educated outside of Israel has a positive influence on English teachers' effectiveness. Among English teachers educated in Israel, those who attended universities with the best reputations (the Hebrew University of Jerusalem and Tel Aviv University) were significantly more effective than those who attended other universities or teachers' colleges. Among math teachers, the only attribute that had a significant effect on teaching effectiveness was mother's schooling: teachers whose mothers had completed high school or earned a higher academic degree were much more effective than other teachers. No similar effect was found for father's education.

We also correlated these teachers' attributes with the three measures of teachers' effort discussed in the previous section: whether the teacher added instruction time beyond regular classes during the program, how many weeks before matriculation exams the teacher exerted special effort, and how many hours of instruction per week he or she added during that time. No significant correlations were found between the personal attributes of participating teachers, in either subject, or these measures of effort.

VI. Conclusions

The evidence presented in this paper indicates clearly that pay-for-performance incentives "work" among schoolteachers as well as in other occupations. This result is evident despite the widely-held concern about the team nature of learning in school, i.e., the belief that a student's output is not the outcome of the inputs of a single teacher but the product of the joint contributions of many teachers. The magnitude of the estimated effects and the evidence concerning teachers' differential efforts under an incentive regime suggest that teachers' incentives are a very promising path toward the improvement of school quality. The evidence culled from this new experiment adds important evidence to the results concerning group school incentives, presented in Lavy (2002).

How effective is incentive intervention relative to other forms of intervention that are also meant to improve matriculation results? We may compare three interventions in terms of their effects on the matriculation rate and their per-student cost. This comparison should be treated as an initial approximation and not as a set of exact figures. We should also bear in mind that the student

¹⁴ See Heckman (2002) and Hanushek (2002) for discussion of this point.

populations treated in these programs were different, a fact that diminishes the validity of the comparison. The teacher incentive program cost \$170 per student and led to a 3.3 percentage point increase in the matriculation rate, from 42 percent to 45.3 percent. The student bonus program evaluated by Angrist and Lavy (2002) cost \$300 per student and elevated the matriculation rate by 6–8 percentage points, from 19 percent to about 26 percent. The group incentive program analyzed by Lavy (2002) cost \$270 per student and boosted the matriculation rate by 1-2 percentage points, from about 45 percent to 47 percent. Another intervention relevant for comparison is the “*Bagrut 2001* program” that the Ministry of Education initiated since year 2000. This intervention targets additional instruction time in small groups (2-6 students) in several matriculation subjects to weak students. The ministry evaluation results show that the program led to a 11-percentage point increase in the matriculation rate of the treated students at an average cost of \$1,100 per student (Ministry of Education, Evaluation Division, May 2002).¹⁵

Among the three incentive programs, the student bonus program was the most expensive in per-student terms but it is marginally more effective in cost-equivalence terms than the teacher bonus program when adjusted for its higher impact on the matriculation rate. The group school incentive program was the least effective in cost-equivalence terms among the three incentive based programs. The added and targeted instruction time program led to a sharp increase in the matriculation rate but at a high cost, almost half the annual expenditure per student. In terms of cost-equivalence it was the least effective among the four programs compared, similar to the teachers group incentive program.

¹⁵ The 2001 average expenditure per student in regular high schools in Israel was about \$2,200.

References

- Angrist, J. (1998). "Using Social Security Data on Military Applicants to Estimate the Effect of Military Service Earnings." *Econometrica* 66 (2): 249-288.
- Angrist, J. and Lavy, V. (1999). "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement." *Quarterly Journal of Economics*. 114 (2): 533-575.
- Angrist, J. and Lavy, V. (2002a). "New Evidence on Computers in the Classroom." *The Economic Journal*, October 2002.
- Angrist, J. and Lavy, V. (2002b). "Achievements Awards for High school Matriculation: Research Methods and Findings." Draft, June.
- Campbell, D. T., "Reforms as Experiments," *American Psychologist* 24 (1969), 409-429.
- Cohen, D. and R. Murnane (1985). "The Merits of Merit Pay." *Public Interest*, 80 summer: 3-30.
- Clotfeller, C. T., and H. F. Ladd. (1996). "Recognition and Rewarding Success in Public Schools." In: H. F. Ladd (ed.), *Holding Schools Accountable: Performance-Based Reform in Education*. Washington, D.C.: Brookings Institution.
- Conley, Sharon and Odden, Allen. (1995). "Linking Teacher Compensation to Teacher Career Development." *Education Evaluation and Policy Analysis, Summer 1995*. 219-237.
- Dehejia, Rajeev H. and S. Wahba. (1998). "Causal Effects in Non-Experimental Studies: Re-evaluating the Evaluation of Training Programs." *Journal of the American Statistical Association*, 2000 .
- Elmore R. F, C. H. Abelman and S. H. Fuhrman (1996). "The New Accountability in State Education Reform: From Process to Performance." In: H. F. Ladd (ed.), *Holding Schools Accountable: Performance-Based Reform in Education*. Washington D.C.: Brookings Institution.
- Gaynor, Martin, and Mark V. Pauly (1990). "Compensation and Productive Efficiency in Partnership: Evidence from Medical Group Practice." *Journal of Political Economy* 98(3): 544-73.
- Gibbons, Robert (1998): "Incentives in Organizations", *Journal of Economic Perspectives* 12(4): 115-132.
- Green, Jerry and Nancy L. Stokey (1983). "A Comparison of Tournaments and Contracts." *Journal of Political Economy* 91: 349-64
- Hahn, Jinyong, Petra Todd, and Wilbert van der Klaauw. "Identification and Estimation of treatment Effects with a Regression-discontinuity Design." *Econometrica* 69 (2001):201-209.
- Hanushek, E. (2002). "Publically provided Education," NBER Working Paper No. 8799.
- Hards, E. C. and T. M. Sheu (1992) "The South Carolina School Incentive Reward Program: A Policy Analysis." *Economics of Education Review*, Vol. 11, No. 1: 71-86.
- Heckman, J. J., (2002). "Human capital Policy", draft.

- Heckman, J. J., H. Ichimura and P. E. Todd (1997). "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Program," *Review of Economic Studies*, 64 (4): 605-54.
- Heckman, J. J., H. Ichimura and P. E. Todd (1998). "Matching as an Econometric Evaluation Estimator: Evidence," *Review of Economic Studies*, 65 (2): 261-294.
- Holmstrom, B. and P. Milgrom (1991). "Multitask Principal-Agent Analysis: Incentive Contracts, Asset Ownership and Job Design, *Journal of Law, Economics and Organization* 7 (Special Issue), 24-52.
- Israel Ministry of Education, *Bagrut Test Data 2000*, Jerusalem: Ministry of Education, Chief Scientist's Office, April 2001.
- Israel Ministry of Education, *The Bagrut 2001 program, an Evaluation*". Jerusalem: Ministry of Education, Evaluation Division, May 2002.
- Jensen C. Michael and Kevin J. Murphy (1990): "Performance Pay and Top-Management Incentives", *The Journal of Political Economy* 98(2): 225-264.
- Kandel, E. and E. Lazear (1992). "Peer Pressure and Partnership." *Journal of Political Economy* 100 (4):801-17.
- Kelley, Carolyn and Protsik, Jean. (1996). Risk and Reward: Perspectives on the Implementation of Kentucky's School-Based Performance Award Program. American Educational Research Association conference paper, April 8, 1996, New York City.
- Lavy, V. (2002). "Evaluating the Effect of Teacher Performance Incentives on Students Achievements." Forthcoming, *Journal of Political Economy*, December 2002.
- Lazear, E. and S. Rosen. (1981). "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy* 89: 841-64.
- Liang, Kung-ye, and Scott L. Zeger, "Longitudinal Data Analysis Using Generalized Linear Models," *Biometrika* 73 (1986), 13-22.
- Malcomson, James (1998): "Incentives Contracts in Labor Markets". In Ashenfelter, O. and D. Card, eds., *Handbook of Labor Economics* 3(B): 2291-2372.
- Milgrom, P. and J. Roberts (1992). *Economics, Organization and Management* , Prentice Hall, New Jersey.
- Moulton, Brent, "Random Group Effects and the Precision of Regression Estimates," *Journal of Econometrics* 32 (1986), pp. 385-97.
- Prendergast, Canice. (1999). "The provision of Incentives in Firms." *Journal of Economic Literature* 37: 7-63.
- Rosenbaum, P.R. and Rubin, D.B., (1985), "Constructing a Comparison Group using Multivariate Matched Sampling Methods That Incorporate the Propensity Score," *The American Statistician* 39: 33-38.

- Sadowski, Michael and Miller, Edward. (1996). "New Ideas Like Collective Incentives and Skill-Based Pay Raise the Same Old Questions." *The Harvard Education Letter*. January/February, 1996.
- Wakelyn, David J. (1996). The Politics of Compensation Reform: A Colorado Case Study. American Educational Finance Association conference paper, March 23, 1996.
- Van der Klaauw, W. (1996). "A Regression-Discontinuity Evaluation of the Effect of Financial Aid Offers on Enrollment," unpublished manuscript, New York University.

Figure 1: The Relationship Between the Correct and the Erroneously Measured 1999 Matriculation Rate
Sample=507 Schools




Figure 2: The Correct 1999 Matriculation Rate Versus The Measurement Error
Sample=507 Schools




Figure 3: The Relationship Between the Correct and the Erroneously Measured 1999 Matriculation Rate
Sample=97 Schools




Figure 4: The Correct 1999 Matriculation Rate Versus The Measurement Error
Sample=97 Schools




Figure 4A: The Correct 1999 Matriculation Rate Versus The Measurement Error
Sample=69 Schools




Figure 5: Determining the Sample of Schools That Were Randomly Assigned To Treatment or Control
Sample=97 Schools




Figure 6: Determining the Discontinuity Sample (Schools Close To the Threshold Value)
Sample=97 Schools




Table 1

Descriptive Statistics: The Measurement Error Sample

	Year 2000			Year 2001		
	Treatment	Control	Difference (s.e)	Treatment	Control	Difference (s.e)
Student's background						
Father's education	10.337	10.129	0.208 (1.007)	10.188	11.054	-0.865 (0.757)
Mother's education	10.315	10.340	-0.024 (1.061)	10.181	10.280	-0.099 (1.082)
Number of sibblings	3.058	2.406	0.653 (0.351)	3.053	1.993	1.061 (0.389)
Gender (Male=1)	0.494	0.505	-0.011 (0.058)	0.534	0.517	0.018 (0.057)
Immigrant	0.017	0.029	-0.011 (0.027)	0.015	0.014	0.001 (0.014)
Asia-Africa ethnicity	0.228	0.287	-0.059 (0.057)	0.216	0.260	-0.045 (0.047)
Lagged student's outcomes						
Math credits	0.375	0.557	-0.182 (0.178)	0.320	0.583	-0.264 (0.153)
English credits	0.175	0.148	0.026 (0.060)	0.138	0.123	0.015 (0.083)
History credits	0.131	0.403	-0.271 (0.084)	0.353	0.775	-0.422 (0.161)
Biology credits	0.000	0.000	0.000 (0.000)	0.000	0.000	0.000 (0.000)
Total credits	4.055	4.256	-0.201 (0.443)	4.111	4.420	-0.309 (0.413)
School characteristics						
Religious school	0.296	0.205	0.091 (0.158)	0.307	0.206	0.102 (0.160)
Arab school	0.165	0.000	0.165 (0.098)	0.164	0.000	0.164 (0.100)
Previous year Bagrut rate (1999,2000)	0.501	0.552	-0.051 (0.032)	0.479	0.498	-0.019 (0.041)
Number of observations	2405	1773		2350	1678	

Note: The table reports the mean of all variables by treatment and control, the differences of means and their standart errors adjusted for clustering using formulas in Liang and Zeger (1986).

Table 2

The Treatment Effect Estimated By the "Measurement Error" Sample

	Attempted exams	Attempted credits	Awarded credits	Attempted exams	Attempted credits	Awarded credits
	Math			English		
Sample mean	1.20	2.09	1.61	0.94	2.51	2.06
Control for correct 1999 matriculation rate	0.078 (0.041)	0.135 (0.080)	0.256 (0.106)	0.027 (0.052)	0.224 (0.142)	0.361 (0.154)
No control for correct 1999 matriculation rate	0.079 (0.044)	0.125 (0.071)	0.163 (0.094)	0.035 (0.043)	0.194 (0.121)	0.327 (0.145)
	History			Biology		
Sample mean	0.33	0.35	0.22	0.68	0.08	0.06
Control for correct 1999 matriculation rate	-0.084 (0.054)	-0.066 (0.051)	-0.077 (0.056)	0.111 (0.143)	0.251 (0.191)	0.136 (0.124)
No control for correct 1999 matriculation rate	-0.081 (0.065)	-0.064 (0.062)	-0.087 (0.069)	0.173 (0.123)	0.160 (0.171)	0.093 (0.107)

Note: The table reports treatment-control differences for different alternative strategic level samples. Standard errors are adjusted for clustering using formulas in Liang and Zeger (1986). The mean of each outcome in the 2001 treatment sample is presented in the squared brackets. Student level controls include the number of siblings, gender dummy, father's and mother's education, a dummy indicating an immigrant student, a set of dummy variables for ethnic background, a set of dummies for the number of credit units gained in the relevant subject before treatment, total credit units gained before treatment and the average score in the relevant tests.

Table 3

Descriptive Statistics: The Discontinuity Sample

	Year 2000			Year 2001		
	Treatment	Control	Difference (s.e)	Treatment	Control	Difference (s.e)
Student's background						
Father's education	11.055	10.424	0.631 (0.486)	10.889	10.455	0.434 (0.511)
Mother's education	11.124	10.733	0.391 (0.561)	11.088	10.882	0.206 (0.575)
Number of siblings	2.609	2.427	0.182 (0.344)	2.552	2.131	0.421 (0.393)
Gender (Male=1)	0.492	0.468	0.024 (0.066)	0.498	0.488	0.010 (0.062)
Immigrant	0.013	0.045	-0.032 (0.022)	0.013	0.008	0.005 (0.007)
Asia-Africa ethnicity	0.218	0.319	-0.101 (0.050)	0.211	0.287	-0.077 (0.054)
Lagged student's outcomes						
Math credits	0.229	0.568	-0.339 (0.147)	0.265	0.578	-0.312 (0.176)
English credits	0.208	0.088	0.120 (0.061)	0.185	0.125	0.060 (0.090)
History credits	0.155	0.176	-0.022 (0.084)	0.434	0.567	-0.133 (0.149)
Biology credits	0.000	0.000	0.000 (0.000)	0.000	0.000	0.000 (0.000)
Total credits	4.044	4.499	-0.455 (0.346)	4.230	4.594	-0.364 (0.388)
School characteristics						
Religious school	0.098	0.325	-0.227 (0.152)	0.092	0.309	-0.217 (0.150)
Arab school	0.128	0.000	0.128 (0.090)	0.128	0.000	0.128 (0.091)
Previous year Bagrut rate (1999,2000)	0.483	0.537	-0.054 (0.016)	0.482	0.507	-0.025 (0.042)
1999 measured with error Bagrut rate	0.426	0.488	-0.061 (0.011)	-	-	-
Number of observations	2523	1564		2535	1406	

Note: The table reports the mean of all variables by treatment and control, the differences of means and their standart errors adjusted for clustering using formulas in Liang and Zeger (1986).

Table 4

The Treatment Effect Estimated By the Discontinuity Sample

	Attempted exams	Attempted credits	Awarded credits	Attempted exams	Attempted credits	Awarded credits
	Math			English		
Sample mean	1.22	2.14	1.74	0.89	2.53	2.16
Control for correct 1999 matriculation rate	0.047 (0.042)	0.100 (0.078)	0.244 (0.109)	0.129 (0.042)	0.212 (0.112)	0.177 (0.145)
No control for correct 1999 matriculation rate	0.046 (0.045)	0.093 (0.080)	0.231 (0.110)	0.132 (0.042)	0.199 (0.109)	0.177 (0.150)
	History			Biology		
Sample mean	0.51	0.53	0.22	0.58	0.05	0.05
Control for correct 1999 matriculation rate	0.138 (0.118)	0.160 (0.117)	0.064 (0.112)	0.146 (0.125)	-0.114 (0.206)	-0.140 (0.107)
No control for correct 1999 matriculation rate	0.122 (0.126)	0.144 (0.125)	0.037 (0.124)	0.147 (0.123)	-0.127 (0.200)	-0.152 (0.100)

Note: The table reports treatment-control differences for different alternative strategic level samples. Standard errors are adjusted for clustering using formulas in Liang and Zeger (1986). The mean of each outcome in the 2001 treatment sample is presented in the squared brackets. Student level controls include the number of siblings, gender dummy, father's and mother's education, a dummy indicating an immigrant student, a set of dummy variables for ethnic background, a set of dummies for the number of credit units gained in the relevant subject before treatment, total credit units gained before treatment and the average score in the relevant tests.

Table 5

Descriptive statistics: the Propensity score sample: treated students

	Math matching			English matching		
	treatment	control	Difference (s.e)	treatment	control	Difference (s.e)
Student's background						
Father's education	9.282	9.297	-0.015 (0.550)	9.471	9.296	0.175 (0.541)
Mother's education	8.831	8.733	0.097 (0.697)	9.149	8.918	0.231 (0.673)
Number of siblings	3.342	3.483	-0.140 (0.394)	3.217	3.296	-0.079 (0.387)
Gender (Male=1)	0.488	0.471	0.017 (0.025)	0.483	0.471	0.012 (0.024)
Immigrant	0.014	0.001	0.014 (0.004)	0.014	0.001	0.014 (0.004)
Asia-Africa ethnicity	0.165	0.185	-0.021 (0.028)	0.179	0.210	-0.031 (0.027)
Lagged student's outcomes						
Math credits	0.362	0.304	0.058 (0.081)	0.531	0.535	-0.004 (0.109)
English credits	0.109	0.160	-0.051 (0.047)	0.060	0.051	0.009 (0.023)
History credits	0.442	0.409	0.033 (0.069)	0.453	0.419	0.034 (0.069)
Biology credits	0.000	0.002	-0.002 (0.001)	0.000	0.002	-0.002 (0.001)
Total credits	4.158	4.101	0.057 (0.288)	4.311	4.256	0.055 (0.259)
Average Score	64.225	63.416	0.809 (2.045)	65.047	63.559	1.488 (1.857)
School characteristics						
Religious school	0.179	0.203	-0.024 (0.064)	0.171	0.185	-0.015 (0.062)
Arab school	0.286	0.326	-0.041 (0.090)	0.232	0.253	-0.020 (0.081)
Previous year Bagrut rate (2000)	0.409	0.423	-0.014 (0.024)	0.422	0.426	-0.004 (0.023)
Number of observations	4490	4490		4865	4865	

Note: The table reports the mean of all variables by treatment and control, the differences of means and their standard errors adjusted for clustering using formulas in Liang and Zeger (1986).

Table 5A

Descriptive statistics: the Propensity score sample: all students

	All matching		
	treatment	control	Difference (s.e)
Student's background			
Father's education	9.362	9.248	0.114 (0.528)
Mother's education	9.005	8.774	0.232 (0.658)
Number of siblings	3.220	3.342	-0.122 (0.363)
Gender (Male=1)	0.483	0.468	0.015 (0.022)
Immigrant	0.013	0.001	0.013 (0.004)
Asia-Africa ethnicity	0.182	0.215	-0.033 (0.027)
Lagged student's outcomes			
Math credits	0.498	0.448	0.050 (0.100)
English credits	0.101	0.090	0.012 (0.038)
History credits	0.442	0.431	0.011 (0.066)
Biology credits	0.000	0.000	0.000 (0.000)
Total credits	4.196	4.256	-0.060 (0.262)
Average Score	63.921	64.132	-0.210 (1.797)
School characteristics			
Religious school	0.178	0.199	-0.021 (0.061)
Arab school	0.247	0.284	-0.037 (0.081)
Previous year Bagrut rate (2000)	0.417	0.425	-0.009 (0.023)
Number of observations	5512	5512	

Note: The table reports the mean of all variables by treatment and control, the differences of means and their standart errors adjusted for clustering using formulas in Liang and Zeger (1986).

Table 6

The Treatment Effect Estimated By the Propensity Sample

	Attempted exams	Attempted credits	Awarded credits	Attempted exams	Attempted credits	Awarded credits
	Math			English		
Treated students						
Sample mean	1.34	2.33	1.82	1.08	2.78	2.24
	0.242 (0.032)	0.398 (0.059)	0.293 (0.068)	0.088 (0.033)	0.230 (0.068)	0.145 (0.079)
All students						
Sample mean	1.19	2.07	1.60	1.05	2.66	2.13
	0.149 (0.039)	0.230 (0.063)	0.151 (0.065)	0.091 (0.033)	0.177 (0.070)	0.120 (0.076)
	History			Biology		
Treated students						
Sample mean	0.51	0.52	0.25	0.72	0.11	0.05
	0.158 (0.084)	0.145 (0.084)	0.020 (0.046)	0.156 (0.090)	0.062 (0.034)	0.024 (0.019)
All students						
Sample mean	0.50	0.51	0.26	0.73	0.11	0.05
	0.145 (0.080)	0.125 (0.080)	0.028 (0.045)	0.192 (0.081)	0.060 (0.032)	0.030 (0.019)

Note: The table reports treatment-control differences for different alternative strategic level samples. Standard errors are adjusted for clustering using formulas in Liang and Zeger (1986). The mean of each outcome in the 2001 treatment sample is presented in the squared brackets. Student level controls include the number of siblings, gender dummy, father's and mother's education, a dummy indicating an immigrant student, a set of dummy variables for ethnic background, a set of dummies for the number of credit units gained in the relevant subject before treatment, total credit units gained before treatment and the average score in the relevant tests.

Table 7

The Effect of Pay For Performance on Teaching Methods and Teacher's Effort

	English teachers		Math teachers	
	Sample mean	Treatment-control difference	Sample mean	Treatment-control difference
Teaching methods:				
Teaching in small groups	0.668	0.085 (0.052)	0.665	0.007 (0.050)
Individualized instruction	0.631	0.112* (0.054)	0.601	-0.028 (0.052)
Tracking by ability	0.512	0.221* (0.055)	0.458	0.13* (0.052)
Adapting teaching methods to students ability	0.954	0.068* (0.023)	0.947	0.011 (0.024)
Teacher's effort:				
Added instruction time during the year	0.662	0.220* (0.052)	0.838	0.015* (0.039)
Added instruction time before Bagrut exam	0.412	0.146* (0.067)	0.390	0.070 (0.056)
Number of weeks before Bagrut exam with added instruction time	4.808	0.564 (0.877)	5.176	2.159* (0.964)
Number of additional instruction hours	3.068	0.829 (0.657)	3.918	2.102* (0.769)
Teacher's effort targeted towards:				
All students	0.328	-0.004 (0.052)	0.575	-0.025 (0.052)
Weak students	0.252	0.129* (0.048)	0.184	-0.058 (0.041)
Average students	0.024	0.019 (0.017)	0.031	0.043* (0.018)
Strong students	0.012	0.028* (0.012)	0.003	0.006 (0.006)
Number of observations	329		358	

Note: Standard errors in parantheses. Astrixes denote estimates which are significantly different from zero at 5% significance level. The English sample includes 141 of the 198 12th grade English teachers that participated in the program. The Math sample includes 169 of the 203 12th grade Math teachers that participated in the program.

Table A1

Teacher's Education And Demographic Characteristics

	English teachers		Math teachers	
	Sample mean	Treatment-control difference	Sample mean	Treatment-control difference
Teacher demographics:				
Age	45.00	-0.697 (1.023)	44.20	0.301 (1.004)
Gender (Female=1)	0.81	-0.005 (0.044)	0.59	-0.024 (0.052)
Born abroad	0.62	-0.160* (0.053)	0.48	0.014 (0.053)
Teacher education:				
Teacher certificate	0.02	0.025 (0.016)	0.03	0.049 (0.019)
B.A in education	0.09	0.012 (0.031)	0.08	0.070 (0.028)
B.A	0.46	-0.066 (0.056)	0.41	0.024 (0.052)
M.A + Ph.d	0.43	0.034 (0.055)	0.47	-0.143 (0.052)
Teaching experience (years)	18.60	-1.470 (0.982)	19.01	-0.139 (1.009)
Education quality:				
Degree from elite universities	0.18	0.089* (0.042)	0.20	0.040 (0.043)
Degree from other universities	0.33	0.004 (0.053)	0.33	0.048 (0.050)
Degree from teacher colleges	0.08	-0.020 (0.031)	0.10	-0.045 (0.032)
Degree from overseas universities	0.41	-0.067 (0.055)	0.36	-0.050 (0.051)
Number of observations	329		358	

Note: Standard errors in parantheses. Astrixes denote estimates which are significantly different from zero at 5% significance level. The English sample includes 141 of the 198 12th grade English teachers that participated in the program. The Math sample includes 169 of the 203 12th grade Math teachers that participated in the program. Elite universities: Hebrew University in Jerusalem, Tel-Aviv, Technion and the Weizman Institute. Other universities: Bar-Ilan, Ben Gurion and Haifa university.