

Beyond Incentives: Do Schools use Accountability Rewards Productively?*

Marigee Bacolod
University of California, Irvine

John DiNardo
University of Michigan

Mireille Jacobson
University of California, Irvine

October 9, 2007

Abstract

“Accountability mandates” – the explicit linking of school funding, resources, and autonomy to student performance on standardized exams – have proliferated in the last 10 years. In this paper, we examine California’s accountability system, which rewards schools (and subgroups within schools) based on a deterministic function of test scores. The sharp discontinuity in the assignment rule – schools that barely missed their target received no funding – generates “as good as random” assignment of awards for schools near their eligibility threshold and enables us to estimate the (local average) treatment effect of California’s award program

This design allows us to explore an understudied aspect of accountability systems – how schools use their financial rewards. Our findings indicate that California’s accountability system significantly increased resources allocated to some schools. In 1999, the average value of the award was 60 dollars per student and 50 dollars in 2000. Moreover, we find that the total resources flowing to districts with schools that received awards increased more than dollar for dollar. This resource shift was greatest for districts with schools that qualified for awards in 1999, the first year of the program, increasing total per pupil revenues by roughly 5 percent.

Despite the increase in revenues, we find no evidence that these resources were used productively. For example, schools that won awards did not purchase more instructional material, such as computers, which may be inputs into achievement. Although the awards were likely paid out as teacher bonuses, we cannot detect any effect of these bonuses on test scores or other measures of achievement. More worrisome, we also find a practical effect of assigning the award based in part on the performance of “numerically significant subgroups” within a school was to reduce the relative resources of schools attended by traditionally disadvantaged students.

*We wish to thank Tom Y. Chang, Julie Berry Cullen, Nora Gordon, Eric Hanushek, Justin McCrary, Doug Miller, Heather N. Royer, Diane Whitmore Schanzenbach, and participants at the 2006 SOLE meetings and the 2007 All-UC Labor Workshop for many helpful suggestions. We especially thank Nora Gordon for her help with the title and Tom Y. Chang for assistance in putting together the data on award apportionments. Financial support was provided by the Haynes Foundation. All errors are our own

1 Introduction

“Accountability mandates” – the explicit linking of a public school’s funding, resources, and autonomy to student performance on standardized tests – have proliferated in the last 10 years. A major impetus to this proliferation was the “No Child Left Behind Act” (NCLB) of 2001, which requires public school students be tested in reading and math annually from third to eighth grade. Schools whose students fail to make adequate yearly progress on these tests risk losing their Title 1 funding and face the possibility of intervention or an outright takeover.

A core rationale for such programs is that a system of sanctions and rewards—for example, the threat of funding cuts or less frequently the promise of more funding—will lead to higher student “achievement” (as measured by standardized exams). One possible unintended consequence of accountability mandates is that students, teachers, and administrators may narrowly tailor instruction and other practices to maximize *measured* performance – possibly at the expense of learning or other socially valuable goals.¹ Studies of the gaming and unintended consequences of accountability systems have found that they lead some teachers to focus only on tested subjects or even on just the relevant students near performance cutoffs.² Other studies suggest the possibility of teacher cheating (Jacob and Levitt 2003a and 2003b). Hanushek and Raymond (2003) observe that the recent literature has paid more attention to gaming responses to the enactment of school accountability measures than to the impact on student achievement itself. As a result, the empirical evidence on the impact of high-stakes testing on student learning is comparatively slim.³

An important challenge to evaluations of interventions based on student performance on stan-

¹Holmstrom and Milgrom (1991), for example, illustrate with a simple model that when faced with “incentive schemes” agents may focus on the most easily observable elements of a multi-dimensional task.

²See for instance, studies cited by Hanushek and Raymond (2003).

³See Hanushek and Raymond (2004), Carnoy and Loeb (2002), Jacob (2006), Figlio and Rouse (2006) among others for examples of assessments of the impact of school accountability on achievement.

standardized tests is that these tests are necessarily imperfect. Kane and Staiger (2002), for example, note that since exam scores contain both “noise” and “signal,” measures of student performance at smaller schools will be noisier than those for similar students at larger schools. The presence of noise can have important consequences for naive before and after comparisons of schools facing accountability mandates. Chay, McEwan, and Urquiola (2005), for example, document that accounting for mean reversion in exam performance substantially reduces the estimated impact of a Chilean accountability-like program on test score gains. A related challenge is that since student bodies change, it can be difficult to assess what changes in exam performance are merely the “spurious” consequence of changing student characteristics rather than changes in school practices *per se*.⁴

Accountability mandates can be crudely divided into those that rely primarily on “carrots” – increased funding for high performing schools – and those that rely on “sticks” – withholding of funds for low performing schools.⁵ The accountability reforms in the California Public School system that we study are an example of the first type: schools that made adequate progress or attained a “passing grade” were rewarded financially. The deterministic nature of the awards program allows us to circumvent many of the difficult issues involved in evaluating the use of these resources and their impacts on student achievement. Schools that met a pre-determined threshold for improvement in exam performance were rewarded with a significant increase in funding. The sharp discontinuity in the assignment rule - schools that barely missed the performance target received no funding - generates quasi-random assignment in awards recipiency for schools that are

⁴Figlio and Rouse (2006), for example, find that much of the estimated test score gains associated with Florida’s school accountability system can be accounted for by changes in the characteristics of students in the state rather than noise.

⁵While deferrals and withholding of funds is stipulated as a possible intervention in NCLB, to our knowledge it has not been practiced in any state. Most systems that rely on “sticks” threaten state takeover and/or outright closure of a school.

close to their eligibility threshold. This quasi-random assignment allows us to generate credible estimates of the impact of the increased resources associated with this program on both test scores and a host of other outcomes. Moreover, our approach provides a large battery of over-identification tests, which allow us to evaluate the validity of the design.

While this research design has many advantages, it falls short of being a complete assessment of the program's effect on the level and distribution of measured achievement in California. Specifically, it cannot capture any effects of the program that occur uniformly to schools that received the awards ("treatment" schools) and schools that did not ("control" schools). For instance, if the mere presence of the awards program causes all schools to sabotage more important goals in favor of more narrowly tailoring curriculum to maximize test scores, our design will not capture such effects. Moreover, for a school that receives funds from the program, the counterfactual to which it is compared is an otherwise similar school within California that did not receive the funds. This counterfactual may not be the same as a school that was not subject to any accountability scheme.

Although such counterfactuals are certainly of interest, ours are more focused on those of immediate policy relevance: what happens when (through an accountability system) we increase a school's resources? How does the school spend these additional resources? And does the school use these resources productively, increasing student achievement?

Our findings are that the program had a significant impact on the financial resources allocated to some schools. In 1999, the average value of the award was 60 dollars per pupil and 50 dollars per pupil in 2000. While this amount represents only about 1 percent of per pupil resources, other funds were directed towards schools qualifying for awards.⁶ At its peak, the financial resources flowing to districts with schools that qualified for awards was about 5 percent of per pupil resources. The

⁶In fact, one argument for accountability programs is its low cost relative to other types of education reforms such as class-size reduction. For instance, see Hoxby (2002).

sizable amount of the resource increase notwithstanding, we find little measurable improvement in standard metrics of achievement, such as exam performance, for those schools that received the award compared to those schools that did not. Our examination of the fiscal data finds no evidence that schools used these awards resources to increase instructional materials. Rather our results are consistent with anecdotal evidence from the California Department of Education that, although the GPAP “was envisioned as money to be used for school site purposes,” it was in many cases “awarded to certificated staff in the way of bonuses or stipends...”⁷

In what follows, we first discuss California’s accountability system and the GPAP, with particular focus on the subgroup rules and the determinants of awards eligibility. Along with the institutional background, we also present a statistical portrayal of California schools by award receipt. We proceed in Section 3 with a discussion of our econometric framework for estimating the effect of the GPAP. Our findings are presented in Section 4. Finally, Section 5 offers a summary and concluding observations.

2 Background: California’s Academic Performance Index and Governor’s Performance Award Program

California’s accountability system, which predates NCLB, was established by the Public Schools Accountability Act (PSAA) of 1999. The PSAA was motivated by assessments indicating that California students were not progressing at the rate necessary “to achieve a high quality education.” Its goal is “to hold each of the state’s public schools accountable for the academic progress and achievement of its pupils within the resources available to schools.”⁸

To measure progress and rank public schools within the state, the PSAA created the “Academic

⁷See page 4 of the document by Patrick J. Chladek, <http://www.wcer.wisc.edu/CPRE/conference/nov02/chladek.pdf> As we explain further below, schools had discretion on how to spend these awards.

⁸See California Education Code 52050-52050.5 for a statement of Legislative Intent.

Performance Index” (API). The API is intended to be a summary measure of school-wide performance on various standardized tests. The index ranges from 200 to 1000 and combines test scores from students in grades 2 to 11. The tests (or other indicators) used and weights accorded to each API component vary from year to year. For the first two years of the program – the only years that the budget allocated funds for performance awards – the API was based solely on the nationally norm-referenced Stanford 9 exam. In broad outline, the API is a (noisy) weighted average of several different exams measured in terms of national percentile ranks, although its precise calculation is somewhat unclear.⁹ For middle and elementary schools, the API incorporated scores on reading, language arts, spelling, and math exams. For high schools, the API was based on reading, language arts, math, science and social studies exams.¹⁰

In 1999, two other awards programs were funded by the State – the Certificated Staff Performance Incentive Award (CSPIA) and the Schoolsite Employee Performance Bonus (SEPB). In contrast to the GPAP, which benefited schools directly, these awards targeted employees – certificated staff, i.e. teachers, in the first instance and both certificated and classified (i.e. paraprofessional, administrative and clerical) staff in the second. The SEPB did, however, grant half of its \$350 million award to the schools for unrestricted uses.¹¹

⁹It is impossible to do justice to how the API is calculated but the following very abbreviated summary in Rogosa (2003) may be useful: “For completeness, here’s a quick reminder of the calculations for the API metric. For a Stanford 9 test, transform the national percentile rank into quintiles: 1-19, 20-39, 40-59, 60-79, 80-99. The quintiles are assigned values 200, 500, 700, 875, 1000; an individual’s API score on a single test is the value for the attained quintile. For any collection of students, the API component score for a single test (e.g. Reading) is the average, over the individuals, of these values (any missing test scores are imputed by the mean of the group). For the battery of tests, API scores in grades 2-8 are a combination of Math (.4), Reading (.3), Language (.15), and Spelling (.15). API scores in grades 9-11 are a combination of Math (.2), Reading (.2), Language (.2), Science (.2) and Social Science (.2).”

¹⁰In 2001, the API was based on both the Stanford 9 and the California Standards Test in English-Language Arts (CST ELA). Additional test components have been added since then. Documentation suggests that other performance indicators such as graduation and attendance rates have also been incorporated into the API calculation but it is unclear how this is done. By law, however, test results must constitute at least 60 percent of the API. For additional details, see API information guides for each year.

¹¹The CDE provides very little information about the SEPB. Discussion of the sharing rules were found only in news reports, such as the one available here: <http://www.svcn.com/archives/lgwt/04.04.01/education-0114.html>

Like the GPAP, both awards were paid out based on API growth, although the \$100 million CSPIA was allocated only to staff at schools that demonstrated the greatest growth over *twice* their GPAP target and had shown growth in the 1998/9 school year. The SEPB was paid to all schools (and staff in schools) that received the GPAP. Because only schools or the employees of schools that received the GPAP could have received these other two awards, our GPAP analysis below is sufficient to capture the combined effect of these award programs on achievement. We discuss the implications of these additional award programs for our analysis of resources in section 4.3. Note, however, that because these programs were effectively suspended after the 1999/2000 school year, our analysis of the 2000 awards program captures the effect of the GPAP alone.

2.1 Award Eligibility – The Simple Case

One bit of complexity in California’s accountability system is that performance awards are based on API “growth” scores – the year to year change in API – for a school as well as for each “numerically significant subgroup.” Before describing what numerically significant subgroups are and how they affect eligibility, we first explain the award determination for a school without subgroups.

Award eligibility is based on a comparison of a school’s growth score with its “target.” In the simplest case, for schools without subgroups, the “API growth target” in a given year is five percent of the distance from the previous year’s API to the statewide target of 800 or a specified minimum. In 1999, the minimum gain was set to one point; in 2000, it was raised to five points. In other words, to receive an award based on 2000 performance, schools had to achieve the maximum of five percent of the distance to the statewide target of 800 or five points.

The 1999 and 2000 award decision rules can be expressed simply as

$$Target_{1999} = \max(40 - .05 * baseAPI_{99}, 1) \tag{1}$$

$$Target_{2000} = \max(40 - .05 * baseAPI_{00}, 5) \tag{2}$$

where $Target_t$ is the minimum gain score (or one year change in API) needed to qualify for an award in year t and $baseAPI_t$ is just the (adjusted) API from $t - 1$.¹²

Figure 1 plots the 1999 and 2000 award targets and demonstrates several noteworthy features of California’s awards program. First, although not made explicit in the official rules, gain scores are always rounded to the nearest integer and thus the awards eligibility thresholds are represented as a step function. Second, and perhaps most importantly, poor performing schools (i.e. schools with lower initial API scores) have to achieve larger test score gains to receive an award than do high achieving schools. Finally, the figure clarifies the effects of the minimum targets set in each year. In 1999, schools with base scores at or above 780 needed to gain only one point over their initial year score to qualify for an award. In 2000, award eligibility was contingent on a minimum gain score of 5 points. This change had the effect of uniformly raising the award threshold by 5 points for those at or above an API of 780 while increasing the target by the nearest integer value of $0.05 * baseAPI - 35$ for those with an API of 700 to 780.

2.2 Award Eligibility with Numerically Significant Subgroups

The PSAA also mandates that “numerically significant” subgroups make “comparable achievement,” defined as 80 percent of the school-wide growth target. Subgroups are defined by racial/ethnic categories (African American, American Indian, Asian, Filipino, Hispanic, Pacific Islander and Cau-

¹²The California Department of Education adjusts base scores to make them “psychometrically” equivalent to the growth scores in the following year. In other words, in principle the base score in a given year, $baseAPI_t$, can differ from the growth score in the previous year, API_{t-1} .

casian) and socioeconomic disadvantage.¹³ A disadvantaged student must either qualify for free or reduced-priced meals or come from a family where the highest level of education is below high school completion. Subgroups with fewer than 30 tested students are not numerically significant. To achieve “numerical significance” a subgroup must have between 30 and 99 tested students and constitute at least 15 percent of total enrollment or have 100 or more tested students.

Table 1 documents the award eligibility calculation for two elementary schools with multiple subgroups in the 2000/2001 academic year (AY).¹⁴ The first column of Table 1 indicates the overall size of the school and the number of students tested in each subgroup.¹⁵ Both schools have over 800 students enrolled, putting them above the 75th percentile of elementary school enrollments in the state. Both schools also have tested students in each of the state-defined subgroup categories but they differ in terms of which groups are sizable enough to be subject to performance targets. Neither school has tested numbers of American Indians, Filipinos, Asians, or Pacific Islanders above 30, exempting them from subgroup performance targets. African American students in Salida Union are also exempt since they number only 16. In contrast, since the tested numbers of Hispanics, whites, and socially disadvantaged students are greater than 100 in both schools, these groups must meet the subgroup performance targets. Tested African American students in Mission Elementary are also subject to subgroup rules as they number well over 30 and represent over 15 percent of tested students.

¹³While racial/ethnic subgroups are mutually exclusive, the socially disadvantaged category may contain students from the racial/ethnic subgroups.

¹⁴Data on academic performance for the 1999/2000 - 2003/4 academic year as well as the monetary awards apportioned to schools under the GPAP for performance in the 1999/2000 and 2000/2001 academic years come directly from the California Department of Education (CDE). School characteristics come from the CDE’s California Basic Educational Data System (CBEDS), an annual school-based census. From CBEDS, we collect data on student enrollment, the allocation of teachers across subjects, in addition to other demographic characteristics such as racial breakdown, parental education, and percent of students receiving free lunch. All of our data sources are described in more detail in the Data Appendix.

¹⁵Small schools, defined as having between 11 and 99 valid Stanford 9 test scores, as well as “very small schools” (fewer than 11 valid scores) were evaluated under a separate, “Alternative Accountability System”.

To gauge award eligibility, columns (2) and (3) show growth and base year API scores, respectively. Column (4) takes the difference between the two and column (5) calculates the score needed to qualify for an award. A school is award eligible if the gain score (column 4) equals or exceeds the target (column 5) for the school as a whole and for each numerically significant subgroup.

Based on school performance alone, both schools would have qualified for an award. This can be seen in the first row (labeled “Overall”) in each panel. Specifically, Mission Elementary had a growth year score of 692 and a base year score of 682. Thus, it achieved growth of 10 API points, 4 points above the school-wide target of 6. Salida Union Elementary had a growth year score of 696, gaining 32 points on its base year score of 664 and 25 points on its target of 7. However, only Salida Union met both the school target and all of its subgroup performance requirements. Two (out of four) of Mission Elementary’s numerically significant subgroups failed to meet their performance target of 5 API points (80 percent of the school target). The API for Hispanic students actually fell nine points and that for socially disadvantaged students (a category that may contain students from the racial/ethnic subgroups) stayed the same.

Since a school cannot qualify for an award unless all performance targets are met, for the purposes of the regression discontinuity design employed subsequently, we characterize each school by the *minimum* of the difference between the gain scores and targets for the school overall and each of its numerically significant subgroups. Thus, Mission Elementary, despite a passing performance overall, is assigned the Hispanic award gap of -14 points, its highest barrier to award eligibility. In contrast, all subgroups at Salida Union Elementary had gain scores that exceed their performance targets. Socially disadvantaged students had the smallest “improvement” in scores, 18 points, and the minimum difference in gain score and performance target. Thus, we characterize Salida Union Elementary by an “award gap” of 12 API points, which qualified the school for an award.

2.3 GPAP Award Allocations

Table 2 describes GPAP allocations and performance overall and by award receipt status. The overall means, in the first column of the table, are based on all elementary, middle and high schools that met the testing participation requirements for the program (95 percent in elementary and middle schools and 90 percent in high schools) and had valid API scores for both base and growth years in the 1999/2000, 2000/2001 or both school years. The first column gives the overall means. The next four columns separate the data into schools that never won an award over the sample period, schools that won an award only in 1999, schools that won an award only in 2000, and schools that won an award in both the 1999 and 2000 school years,

As shown in the first row, about 23 percent of the sample is composed of schools that never won an award. About 31 percent of schools won awards in 1999 alone. Due at least in part to the state's raising the award eligibility threshold, the share winning awards in 2000 alone is only 14.7 percent or about half of the corresponding 1999 figure. In contrast, almost 32 percent of schools won awards in both the 1999 and 2000 academic years.¹⁶ The average per pupil payment was about \$63 across both years, a few dollars less than the average \$69 per pupil payment indicated by the state.¹⁷ To put this in perspective, K-12 public school expenditures in California in the 1999/2000 school year were roughly \$6000 per student (Carroll et al. 2005). Thus, in principle, these awards raised per pupil spending by just over 1 percent. More importantly, perhaps, awardees have considerable discretion (requiring only local school board approval) in how they use these funds. And, as we will show in section 4.3, additional resources flowed to districts and thus presumably schools that

¹⁶Note that these figures are not strictly correct since the observations in the table are at the school-year not the school level. But, since there are roughly the same number of schools with valid testing data in both years, this is a good approximation. Expressed by school, 19.2 percent never won an award, 34.4 percent won in 1999 alone, 14.5 percent won in 2000 alone, and 31.9 percent won in both years.

¹⁷This small discrepancy between our and the official figures, however, is likely due to differences in the quality of enrollment data.

qualified for awards.

The fourth row of Table 2 shows school enrollments. The enrollment figures speak to a fundamental problem with relying on mean test scores to measure school performance (Kane and Staiger 2001; Chay et al. 2005). Specifically, due to large error variances, test scores from any given year provide a poor measure of school rankings. This problem is particularly acute for small schools, since, all else equal, their mean score will have higher sampling variation. The implication of this higher sampling variation is that small schools are more likely to have a particularly lucky year and win a performance award. Consistent with this, Table 2 shows that schools that won awards in both years are smaller ($p < 0.001$) and schools that never won an award larger ($p < 0.001$) than the average school. Furthermore elementary schools, which are the smallest type of school in our sample, with an average enrollment of 636 pupils, are underrepresented in the category of schools that never won awards and overrepresented among those winning awards in both years. While they make up 70 percent of schools in our sample, elementary schools represent only 49 percent of schools that never won an award and about 85 percent of those winning in both years. At the other extreme, high schools, which are by far the largest type of school in our sample, represent about 13 percent of the sample but almost 30 percent of schools that never won an award and only 3 percent of schools that won awards in both years.

The eighth and ninth rows of Table 2 explore the implications of the state's subgroup rules. The typical school has only one subgroup; 15 percent of schools have no subgroups. The most common subgroups are socially disadvantaged, followed by Hispanic and white. Not surprisingly, since schools with subgroups face additional eligibility criteria, schools that never won an award have more subgroups and those that won awards in both years have fewer subgroups than the average school. The next row shows the share of schools in each category that lost an award due to

subgroup rules. Put differently, this row shows the share of schools that would have won an award absent the subgroup rules. Overall, about 18 percent of schools would have won an award based on school performance alone but were ineligible because of subgroup rules. About 45 percent of schools that never won an award would have won without the subgroup rules. This average over the two award years masks differences across the two years attributable to the evolving awards eligibility criteria. Whereas 53 percent of schools in the never group would have won awards based on the school criteria alone in 1999, only about 38 percent would have won in 2000 were it not for the subgroup rules. In other words, raising the standards for schools had the effect of minimizing the bite of the subgroup rules.¹⁸.

The next set of nine rows in Table 2 show the average API score across award years overall and separately for all numerically significant subgroups. The mean API score in the sample is 652. White, Asian and Filipino subgroups perform well above average. Socially disadvantaged subgroups, which are numerically significant in 9082 of our 10720 school-year observations, have an API of 581, almost two thirds of a standard deviation below the overall average. Hispanics are the next most common subgroup. They are numerically significant in 8462 school-years and have an API of 574, also well below the overall mean. Not surprisingly, since those with lower API scores need to make larger gains in order to qualify for an award, mean API scores overall and within each subgroup are lowest for the set of schools that never won awards. Below the API scores are nine rows showing the average gain scores, the basis for award eligibility, across years for schools overall as well as for each subgroup. The average gain score is about 26 points but is less than 10 for schools that never won an award and just over 40 for schools that won in both years.

¹⁸This effect can also be seen by comparing the share of 1999 awards winners that lost awards in 2000 with the share of 2000 awards winners that lost awards in 1999 due to subgroup rules. While 20 percent of 2000 awards winner would have won awards in 1999 based on the school performance alone, only about 15 percent of 1999 winners would have done so in 2000

While these summary statistics shed light on the differences in characteristics of schools winning awards across treatment years, the primary focus of this paper is the impact of award receipt on school behavior, *ceteris paribus*. Consequently, in the work that follows we will first demonstrate that there was in fact a treatment generated by the award program. We then try to determine whether award receipt had any effect on either the level of API scores (overall and for each numerically significant subgroups) or school resources.

3 Econometric Framework

To identify the causal impact of California’s awards program on outcomes such as API scores or resource allocations, we employ a regression discontinuity (RD) design where we essentially compare the behavior of schools that just barely won an award to those that just barely missed winning an award.¹⁹ Our design is similar to the original RD approach used by Thistlewaite and Campbell (1960) to estimate the impact of a test-based scholarship award program on future academic outcomes except our unit of analysis is the school. Although individual student data would be appealing, particularly for estimating the impact of the awards program on achievement, these data are difficult, if not impossible, to obtain. Moreover, since the PSAA is ultimately based on average school performance, school-level data is sufficient for characterizing the program.

Suppose that the relationship between schools’ average performance and resources and their receipt of an award is given by the constant treatment effects model:

$$Y_i = \alpha + \beta T_i + \gamma X_i + \epsilon_i \tag{3}$$

¹⁹More precisely, as shown in Lee (2005), we will estimate a weighted average of the population treatment effects, where the weights are positively related to each observation’s distance to their awards target. Thus, schools closest to their awards threshold will contribute the most and those farthest away the least to the estimated treatment effect.

$$T_i = \mathbf{1}(D_i \geq 0) \tag{4}$$

where Y_i is school i 's achievement score or measure of resources; α is a constant; and T_i is an indicator equal to 1 if school i received an award. In addition, let D_i equal the school's distance to its award eligibility target $((API_{it} - API_{it-1}) - Target_i)$, so that zero corresponds to having just met the target.

The primary challenge to estimating the effect of the award program β is that awards are not randomly allocated across schools. A simple comparison of schools that receive the treatment to those that do not will be biased because treated schools differ greatly from untreated schools for reasons other than the treatment. As our discussion above and Table 2 show, for example, schools with more minorities or subgroups are significantly less likely to receive an award.

To overcome this problem, we exploit the rules assigning treatment. Schools for whom $D_i \geq 0$ win the award; schools for whom $D_i < 0$ do not. The sharp discontinuity in the rules that translate test scores into award eligibility generates quasi-random assignment of award receipt near the eligibility threshold. To see how this works, consider the average outcome for treated schools with a specific value for their score ($D_i = \Delta > 0$):

$$E[Y_i|T = 1, D_i = \Delta] = \alpha + \beta + \gamma E[X|D = \Delta] + E[\epsilon|D_i = \Delta].$$

Likewise consider a school that does not receive the treatment $D_i = -\Delta < 0$ whose score is Δ below the threshold:

$$E[Y_i|T = 0, D_i = -\Delta] = \alpha + \gamma E[X|D = -\Delta] + E[\epsilon|D_i = -\Delta].$$

A naive comparison – the average outcome for treated schools above the threshold by an amount Δ to those untreated schools who are below the threshold by an amount Δ to the average outcome for the untreated is merely:

$$\begin{aligned}
E[Y_i|T = 1, D_i = \Delta] - E[Y_i|T = 0, D_i = -\Delta] &= \beta + \underbrace{\gamma E[X|D = \Delta] - \gamma E[X|D = -\Delta]}_{\text{Difference in Observables}} \\
&\quad + \underbrace{E[\epsilon|D_i = \Delta] - E[\epsilon|D_i = -\Delta]}_{\text{Difference in Unobservables}}
\end{aligned}$$

Now consider choosing Δ to be small so that we are comparing schools just above the threshold to those just below. In the limit, as $\Delta \rightarrow 0$ the above expression reduces to:

$$\begin{aligned}
\lim_{\Delta \rightarrow 0} E[Y_i|T = 1, D_i = \Delta] - E[Y_i|T = 0, D_i = -\Delta] &= \beta + \gamma E[X|D = 0] - \gamma E[X|D = -0] \\
&\quad + E[\epsilon|D_i = 0] - E[\epsilon|D_i = -0] \\
&= \beta
\end{aligned}$$

That is, as we approach the threshold from the left and the right, both the unobservable and observable differences of treated schools become smaller and smaller. This research design provides a large number of *testable* restrictions, which are similar to those available in a randomized controlled trial (RCT), and add to the credibility of the design. In particular, in an RCT a basic specification test is to compare the average baseline characteristics of treated group to the control group. In a proper RCT these will be the same on average. Likewise, in this research design, schools just to the left and just to the right of the threshold should look similar.

As a practical matter, it is not necessary to limit the comparison to just the few schools to the left and the right of the threshold. One can recast the problem as estimation of the following relationship:

$$Y_i = \alpha + \beta T_i + g(D_i) + \epsilon_i$$

where $g(\cdot)$ is a unknown continuous function. Although unknown it can be approximated sufficiently well by polynomials in D and its full interactions with the awards indicator T . In practice below, we determined that the fifth-order polynomial was the most parsimonious specification implied by the underlying data.²⁰

A minor issue is the fact that we have only considered the causal impact of the binary treatment. In fact, there are different levels of the award, or varying treatment intensities. This can be accommodated easily by recasting the problem as a simple instrumental variables estimator for the following equation system:

$$\begin{aligned}
 Y_i &= \alpha' + \theta A_i + g'(D_i) + \epsilon'_i \\
 A_i &= \alpha + \psi T_i + h(D_i) + \nu_i \\
 Y_i &= \alpha + \beta T_i + g(D_i) + \epsilon_i \\
 &= \alpha + \theta \psi T_i + g(D_i) + \epsilon_i
 \end{aligned}$$

In this set up, A_i is the endogenous regressor, the second equation is the “first stage” equation, and the third equation is the “reduced form” equation for the outcome. Our estimate of θ is merely the indirect least squares estimate $\frac{\widehat{\beta}}{\psi}$ – the ratio of the discontinuity in the outcome equation to the ratio of the discontinuity in the award equation. If we wish to consider the case where the treatment effect is random, then provided that monotonicity holds (the effect of higher growth scores is uniformly to increase the probability of receiving treatment) then this parameter identifies the local average treatment effect or the effect of the program on those schools induced to win the award by their score (Hahn, Todd, and van der Klauuw 2001). Finally, just as in an RCT,

²⁰To illustrate this, our graphs below superimpose the predicted values from the polynomial on top of the means for each discrete point of support. We explain these graphs and other specification tests below.

provided that the X 's are balanced (a restriction we test), we can also include exogenous covariates X for variance reduction purposes, although they are not required for consistent estimation of the parameter of interest.

3.1 Estimation Issues

In order to implement the RD, we need to first verify that the Governor's Performance Award Program rewarded API growth in both the 1999 and 2000 academic years according to the established rules. Figures 2a and 2b consider the probability of an award payment and Figures 3a and 3b the per pupil payment for schools at each distance from their eligibility threshold in 1999 and 2000. Across figures, the solid lines represent a parametric estimate of the conditional probability of an awards payment (2a and 2b) or of the per pupil payment amount (3a and 3b) to schools at each distance, D , from the eligibility threshold. Each school's distance to its eligibility threshold is just $((API_t - API_{t-1}) - Target_t)$ or the difference between its gain score and growth target as defined above. Consistent estimation of the treatment effect requires that our polynomial be sufficiently flexible to capture the true underlying continuous function.

Operationally, our parametric estimates are just the least squares fitted values from the following equation:

$$A = \delta T + P' \alpha_0 + TP' \alpha_1 + X' \beta + \varepsilon \quad (5)$$

where A is either the probability of awards reciprocity or the per pupil award payment, $T \equiv 1(D \geq 0)$ is an indicator for whether a school crossed the eligibility threshold, $P' = (D, D^2, D^3, D^4, D^5)$ is a fifth order polynomial of the distance, D , to the awards threshold and TP' is the interaction of our eligibility indicator with this fifth order polynomial. We include the interactions so as to allow the polynomial fit to differ on either side of the eligibility threshold.²¹

²¹For variance reduction purposes, we also include X , a set of control variables that include: a school's total enroll-

One potential problem, more apparent than real, is that our polynomial provides an inadequate parameterization of the true underlying continuous function. While in principle we could use nonparametric local linear regression or other techniques, as Card and Lee (2006) observe, such techniques confront the difficulty that the underlying data (changes in test scores) are not actually continuous but rather discrete. In this case, the “true” nonparametric estimator is just the set of mass points in the running variable.

Consequently, as we explain in more detail below, we assess the adequacy of our parametric representation by comparing our model to the fully saturated model that includes a separate indicator for every specific value of the running variable:

$$A = d \sum Z_d \gamma_d + X' \beta + \mu \tag{6}$$

where Z_d is a dummy variable that equals 1 if the school’s distance is equal to d , and 0 otherwise, γ_d are fixed effects for each distance to (and including) the awards eligibility threshold, and X is the set of covariates defined above. The coefficients γ_d , which are just regression adjusted average award payment probabilities or per pupil award payments, are represented as open circles in our figures below. Plotting these coefficients along with our parametric fit allows for a simple “eyeball” test of the extent to which our estimates are a spurious consequence of “noise” in the data. Below, we also describe and present the results of more formal tests.

Figures 2a and 2b show that, as per the award rules, there is a marked discontinuity in the probability of receiving an award at the eligibility threshold. Schools that failed to meet their

ment, the number of numerically significant subgroups in the school, the percent of tested students by race/ethnicity (white, black, Hispanic, Filipino, Asian, Pacific Islander, or American Indian), the share of students qualifying for free or reduced price meals, and dummies for whether the school is an elementary or high school (with middle school the omitted category). All covariates correspond to the academic year of the growth year score, i.e. t not $(t - 1)$. Standard errors are adjusted to allow for an arbitrary correlation in errors at the level of D , the distance to the awards threshold, as suggested by Card and Lee (2006) to account for potential misspecification.

target, and thus are to the left of the eligibility threshold, are not observed receiving an award payment.²² At the threshold, where a school's API equals its target, the probability of receiving an award jumps to almost one. For 1999, the estimated discontinuity is 0.91 with a t-stat of over 25 and for 2000 it is 0.84 with a t-stat of about 56. Both the regression adjusted averages and the polynomial fit past zero are also strictly below one because a small fraction of schools, about 8 percent in 1999 and almost 10 percent in 2000, made their API target but did not receive an awards payment. According to the California Department of Education these schools may have been disqualified because of "data irregularities," too high a share (over 15 percent) of parents obtaining exam waivers for their children, or student population changes that invalidated the school's API.

Figures 3a and 3b are analogous to 2a and 2b except that the dependent variable is the per pupil award payment rather than a simple dichotomous measure of reciprocity. As in these previous graphs, schools to the left of the eligibility threshold did not receive award apportionments and thus have per pupil award payments of \$0. At the discontinuity the per pupil award payments jump sharply. For 1999, the estimated discontinuity is about \$60 per pupil with a t-stat of almost 25 and for 2000 its \$50 per pupil with a t-stat of about 51. The difference in the per pupil award payments across these two years are likely attributable to two causes. First while the number of 2000 school recipients was just over 70 percent of the 1999 figure, the budget for the 2000 awards program was set at \$144.3 million less than 64 percent of the 1999 budget. Second, at least in our data, a higher share of schools in 2000 appear to have met their award targets but did not receive apportionments.

A visual comparison of the estimates from our parametric models and the regression adjusted

²²In actuality, in 1999, 5 schools or 0.3 percent of schools that according to our data failed to meet their target received awards payment from the State averaging \$60.5 per pupil. In 2000, none of the "failing" schools received GPAP payments.

averages of award reciprocity suggests that the fifth-order polynomial fits are reasonable. We confirm this using a more formal test. Following Lee and Card (2006), we calculate a goodness of fit statistic, $G \equiv \frac{(RSS_r - RSS_{ur})/(J-K)}{RSS_{ur}/(N-J)}$, where RSS_r is the residual sum of squares from the restricted (polynomial-fitted) model, RSS_{ur} is the residual sum of squares from the fully flexible or unrestricted model, J is the number of parameters in the unrestricted model, K the number of parameters in the restricted model and N the number of observations. Under normality G is distributed $F(J-K, N-K)$. With this F -statistic we can test the null hypothesis that the fit of the polynomial model is as good or has as much explanatory power as the fully flexible model. For both the 1999 and 2000 awards reciprocity models, G is less than one (or 0.773 and 0.743, respectively). For the corresponding per pupil award payments regressions, G is 0.765 and 0.637 for 1999 and 2000, respectively. In all cases, the F -statistics indicate that we cannot reject the null that the restricted and unrestricted models have similar goodness of fits.

The framework used above to establish the discontinuity in the GPAP, is the same estimating equation we employ to determine the causal impact of the program. More specifically, we will use (4) to estimate the treatment effect on Y (instead of A), and we will use (5) to test the sensitivity of this estimate to our functional form assumptions. The ratio of the RD estimate of our outcome of interest, say math teachers or spending per pupil in year $t + 1$, to the RD estimate of award reciprocity or per pupil award apportionments in t , will form our causal estimate of the treatment on the treated.

One concern in interpreting this ratio as a causal estimate, however, is that other predetermined school characteristics may be changing at the same time. As discussed above, identification of δ requires that X is continuous at $d = 0$. For example, if schools near the discontinuity encourage certain types of students to transfer to other schools, we may see changes in the share of students

by race or socioeconomic status. Since these factors independently affect outcomes, they will, at least in principle, confound our estimates of the treatment effect of the awards program.

The fact that the State uses API *changes*, which should be considerably noisier and more difficult to manipulate than levels, to allocate awards should lend considerable credibility to our research design. But, while we cannot prove that all other predetermined characteristics are balanced, we can check to see whether observable characteristics are smooth through the discontinuity. To do this, we have plotted regression adjusted averages of and estimated polynomial fits to the 1999 values of all covariates described above as well as some other available characteristics, against the distance to the threshold for schools in either the 1999 or 2000 school years.

Appendix Table 1 reports the estimated discontinuities at the 1999 awards threshold for each of these predetermined characteristics. In 12 out of 14 cases, the estimated discontinuity is not statistically distinguishable from zero at even the 10 percent level. Thus, by in large, the estimated discontinuities reported in Appendix Table 1 provide support for the idea that schools close to the eligibility threshold are similar on predetermined characteristics.

Two cases merit some additional discussion. We estimate a small discontinuity in the percent of students qualifying for free or reduced price meals and the percent of tested students that are Asian American. Neither are significantly different from zero at conventional levels (i.e. 5 percent), but the p-values for each are only about 0.07. For free or reduced price meals, the point estimate implies a 4.7 percentage point or roughly 10 percent drop in the share of students qualifying in schools that just received awards relative to those schools that just missed receiving an awards. For the share of test-takers that are Asian American, the implied effect is a 1.8 percentage point or an almost 23 percent increase.

Appendix Table 2 reports the estimated discontinuities from these same (1999) characteristics

but relative to the 2000 awards threshold. Somewhat reassuringly, we no longer find evidence of a discontinuity in the share of students eligible for free or reduced price meals or in the share of test takers that are Asian American. In fact, relative to the 2000 awards threshold, we cannot reject zero in 13 of the 14 tested characteristics. The exception for the 2000 awards year is the share of 1999 test takers that were African American. Our estimates imply that schools that just barely qualified for an award in 2000 had almost 25 percent fewer (2 percentage points off of a base of 8.3 percent) African Americans than those schools that just qualifying.

The fact that those characteristics for which we cannot strictly rule out a discontinuity differ across the two award years, gives us some hope that they are generated by random variation. As is well understood, with a large number of comparisons a small fraction of these are expected to be “significant” by sheer random variation.

To the extent that these discontinuities are real, however, they suggest that our estimates may be slightly biased towards finding an impact of the awards program on academic achievement. For example, students qualifying for the school meals program, who are automatically characterized as socially disadvantaged, are more likely to perform poorly (see the API scores for socially disadvantaged subgroups in Table II). Similarly, African American students perform below average as a subgroup. Asian Americans, by contrast, perform well above the state average. Thus, at the discontinuity, a drop in the share of students qualifying for free or reduced price meals and a bump up in the share of Asian Americans might lead us to overstate any change in test scores (or other positive outcomes) associated with the 1999 awards program. By the same logic, a drop in the share of African American students at the discontinuity may overstate any change in outcomes associated with the 2000 awards program. Fortunately, our estimates also suggest that to the extent that such a bias exists, it will be small as long as higher scores do not induce some schools who –

in the absence of their higher score – would have received an award, to in fact be denied an award. This so-called “monotonicity” condition appears to be reasonable in our context.²³

4 Results

Having established that the data support the validity of our research design, we next discuss our estimates of the impact of the award program on achievement and resources. Where the pattern of outcomes appear similar relative to both the 1999 and 2000 awards threshold, we only provide figures of the former. We do this to minimize the number of redundant plots. In all tables, however, we present estimates of the magnitude of the discontinuity, its standard error and the F-test of the correspondence between our polynomial fit and the fully flexible model for both award years.

4.1 Evidence on Achievement

We first consider the impact of the awards program on the level of achievement. If schools that win awards are able to spend these resources in ways that positively impact achievement, then we should see a break in API scores at the discontinuity. In other words, schools that just barely won the awards should have higher scores in subsequent years than their counterparts that just barely missed winning an award.

Figures 4a, 4b, and 4c graphically represent our RD estimates of the impact of the 1999 awards program on achievement levels in 2000, 2001, and 2002. Figures 5a, 5b, and 5c present corresponding plots for the socially disadvantaged subgroup, the most common of all the subgroups. Recall that the t school year spans the t and $t+1$ calendar years and testing for the year t API is conducted in May $t+1$ of each school year. Because the first apportionment for the awards program was made in January 2001, the middle of the 2000/2001 academic year and the second and final payment in

²³For a discussion, see Lee (2005.)

March 2002, the following school year, we do not anticipate finding any impact on achievement in 2000 (as measured by test scores in May 2001).²⁴ Thus, it is reassuring that, as shown in Figure 4a, 2000/2001 school-wide API scores are smooth across the 1999 awards eligibility threshold. The same basic pattern holds for the 2000/2001 API scores for socially disadvantaged subgroups.

To the extent that the additional resources were put towards instruction, we might expect to find a discontinuity in the 2001/2002 school years. But, Figures 4b and 5b also indicate that scores were smooth across the award threshold. Moreover, the close correspondence between the polynomial fits to the API (the solid lines on either side of the awards target) and the regression adjusted average API scores at each distance from the threshold (the open circles) suggest that our estimates of the discontinuity in API scores (or lack thereof) is not an artifact of our modeling choices. This is confirmed by the F-statistics reported in Table 3. The estimates from 2002 are considerably noisier. But, the graphical analysis and the estimates provided in Table 3 broadly confirm a finding of no effect of the awards program on API scores. Similarly, as shown in Panel B of Table 3, we find no evidence of an impact of the 2000 awards program on schoolwide or subgroup achievement. In no case are the estimates of the impact of the 2000 award program on API scores significantly different from zero.

We also examine an alternative measures of academic achievement to test for an impact of the awards program. Specifically, we use measures of the percent of students in a school that tested proficient in English and language arts and in mathematics. These data, which are first available to us in 2001, are based on the California Standards Tests (CSTs) for grades 2-8 and the California High School Exit Examination (CAHSEE) for secondary school students. While these scores have been incorporated into the API over time – making them imperfect complements to the API –

²⁴See <http://www.cde.ca.gov/ta/sr/gp/history.asp> for details on the history of awards apportionments. Note that there were other awards programs for the 1999/2000 but these were paid to teachers and not schools.

they have the advantage of being reported in a very transparent form.²⁵ Yet, for neither the 1999 nor 2000 awards program can we detect any evidence of improvements in either the percent of students testing proficient in ELA or math for schools that just qualified relative to schools that just missed qualifying for an award. Analyses using API or proficiency gain scores, noisier measures of achievement, yield similar conclusions.

Finally, we also focus on the API scores of the subgroup that determined awards eligibility within the school. The idea is that since this subgroup is the worst performing group in the school, award resources in the following periods may be targeted to them. However, similar to previous analyses, we find no significant impact of the awards on these subgroup-specific outcomes.

4.2 Evidence on School Resources

One reason for these null findings, other than the simple possibility that resources do not translate directly or easily into academic improvements, may be that GPAP funds were not used for instruction. There are several ways this could happen. First, it is possible that districts, which have fiscal authority over schools, or the state, which provides much of the funds to districts, offset the awards paid out by the GPAP through corresponding reductions in other funds.²⁶ Under this scenario, award schools might have no additional financial resources to invest in achievement. Second, even if neither districts nor the state undo the monetary transfer required by the GPAP, schools are not constrained to spend these funds on instruction. In fact, they have almost complete discretion over the use of GPAP funds, needing only the approval of the local school board. Thus, to the extent that awards are spent on non-instructional staff, capital outlays, and so on, we may not expect to

²⁵The CST in English and language arts was added to the API measure in 2001. And, the CST in math and the CAHSEE were added to the API in 2002.

²⁶Baicker and Jacobson (2006) provide evidence that counties engage in this type of budgetary offsetting by reducing allocations to police agencies that receive financial rewards from state or federal government for drug enforcement activities. Similarly, Gordon (2004) finds that increases in federal spending on low income school districts are offset by reductions in local spending.

see any improvements in academic achievement.

Unfortunately, fiscal data on revenues and expenditures are not available at the school level. Rather, they are reported at the district level. This limits our ability to determine whether an individual school receives its award money from the district and, if it does, how it gets spent. We can, however, observe school-level inputs such as the number of teachers per pupil overall, the share allocated to each subject (math, English, science, physical education, and special education) as well as the number of instructional computers per pupil and internet-connected classrooms per 100 students.

We have estimated the impact of both the 1999 and 2000 awards program on each of these inputs in the 2000-2002 school years. Table 4 presents a subset of these estimates. Panel A presents our estimates of the impact of the 1999 award program on teachers per pupil, the share devoted to math instruction, the share devoted to English instruction, computers per pupil and internet connections per 100 students in the 2000/2001 academic year. Panel B presents estimates of the 2000 award program on the same category of outcomes but for the 2002/3 academic year. We consider this year because the 2000 award disbursements were not made until the 2002-3 fiscal year (specifically, July and October 2002).

We find little evidence of any changes in these inputs. Given that the award payments were a one shot deal and hiring requires a long term fiscal commitment, it may not be surprising that we find little impact of either the 1999 or 2000 award program on the number of teachers per student. On the other hand, we might expect to see a change in the allocation of teachers across subjects, if, for example, increased funds could be used to encourage some instructors to switch from their normal subject to one that is more valuable in an accountability system. But, our estimates of the impact of the award programs on teachers per pupil and the share of teachers in math or english

are neither statistically nor economically significant.

A more likely use of the award funds might have been for instructional resources. Indeed, the California Department of Education anticipated that these funds should be used “for the purchase of computers, instructional materials, or playground improvements.”²⁷ But, like the CDE, we find no evidence that award payments were used to increase the number of computers or internet connected classrooms within a school. If anything, the results for the 2000 award program suggest that computers per pupil increased less among those schools that just qualified for awards than those that just missed qualifying. But, this interpretation of the point estimate for computers per pupil (-.020 with a standard error of .006) should be viewed with considerable skepticism. As our goodness of fit statistic and the corresponding p-value (.0002) suggest, the polynomial fit for this model is poor relative to the fully flexible model. Thus, the estimated reduction in computers per pupil for awardees relative to nonawardees is likely driven by functional form. The more important lesson from Table 4 is that we find little evidence of increases in computers per pupil or any other measure of resource allocations among schools that received the GPAP.

4.3 Evidence on Fiscal Outcomes

We next turn to the fiscal data. Although we cannot observe changes in an individual school’s revenue or expenditures, we can determine whether the State offsets or alternatively disproportionately increases funds to districts that just barely qualified for the GPAP relative to those that just missed qualifying. And, if we find no evidence of offsetting, we should be able to trace out where (district-wide) any additional funds get allocated.

In order to characterize where each district lies relative to the school-level award eligibility

²⁷See page 4 of the statements from Patrick J. Chladek at the November 2002 National Conference on Teacher Compensation and Evaluation, <http://www.wcer.wisc.edu/CPRE/conference/nov02/chladek.pdf>.

threshold, we sort all of its schools by their distance from this threshold. We then assign to each district the *maximum* of the school-level “award gap.” Recall that because of the subgroup rules we characterize each school’s “award gap” as the minimum difference between the gain scores and targets for the school overall and each of its numerically significant subgroups. Thus, each district “award gap” is the across-school maximum of the within-school minimums. This definition of the district “award gap” then picks up whether any school in the district received treatment. It further characterizes the district according to the distance to the award eligibility threshold of the school that performed best relative to this target. If any treated school in a district was far from the award eligibility threshold, then the district as a whole will be characterized as far from the cutoff. If no schools were treated but at least one was close to its target, then the district as a whole will be characterized as just barely missing award eligibility.

Figures 6a and 6b are the district level analogue to figures 3a and 3b. They show the mean district-level apportionment per pupil in 1999 and 2000 by proximity to the awards threshold (open circles) as well as the polynomial fits to these data.²⁸ These figures (and the results reported in the first column of both panels in Table V) provide confirmation that we can still uncover the treatment at the district level. Because any given district may contain schools that won awards of varying per pupil amounts as well as schools that did not win awards, the estimated discontinuity in apportionments is below that for the school-level. For both 1999 and 2000, districts to the left of the eligibility threshold did not receive awards apportionments and thus have per pupil award payments of (roughly) \$0. At the discontinuity the per pupil award payments jump to about \$43 per pupil (with a t-statistic of about 12) for 1999 and about \$28 per pupil (with a t-statistic of 9)

²⁸To save degrees of freedom, we estimate polynomials with equal slopes on either side of the polynomial. We do not include covariates as these tend to decrease rather than increase the precision of the district-level estimates. Our F-tests suggest that these fits are still quite close to the fully flexible model.

pupil for 2000 test performance.

We next turn to district revenues and expenditure data. According to the School Fiscal Services Division of the CDE, GPAP apportionments are classified as unrestricted revenues. Figure 7a plots unrestricted revenue per pupil in 2000 relative to the 1999 award threshold. The jump in per pupil revenue at the awards threshold is, in fact, much larger than the \$43 apportionment shown in Figure 6a. The RD estimate (also reported in Table V) suggests a jump of \$102 per pupil (with a t-statistic of 21) at the discontinuity. Why should this be so much larger than the apportionment estimate? Other funds are also included in the unrestricted revenue category.²⁹ But to the extent that our RD provides quasi-random assignment, these funds should not differ systematically at the awards eligibility threshold except through the (direct and indirect) effects of the award program. Thus, Figure 7a implies that for every per pupil dollar a district was supposed to get through the 1999 awards program, the district, in fact, received closer to \$2.50 per pupil (with a t-statistic of over 4). Some of these additional funds may be attributable to the \$350 million Schoolsite Employee Performance Bonus program. Bonuses, which were only available for 1999 performance, were allocated based on the number of FTEs in schools winning the GPAP and were divided equally between a school and its employees. We were not able to locate any documentation of disbursements under this program. Any additional funds not attributable to this program may simply be evidence of crowd-in. In other words, the state may be further padding the budget of districts with award winning schools. We find no evidence, based on this revenue category, that such activity continued in 2001 or 2002.

For the 2000 awards program, we do not detect any increase in unrestricted revenues until 2002. This makes good sense since apportionments for the 2000 awards program were not made

²⁹Specific sources of revenue within this category are not available from district-level fiscal data.

until the 2002-3 fiscal year. More importantly, as shown in Figure 7b, the RD estimate suggests that the unrestricted revenues only increase by about \$20 per pupil (with a t-statistic of 3) at the discontinuity. This implies that for every per pupil dollar a district was supposed to get through the 2000 awards program, they, in fact, received only about \$0.72 per pupil (with a t-statistic of almost 3). We cannot reject that this point estimate is significantly different from one. Thus, the results for the 2000 program, when the only monetary awards at stake were from the GPAP, are also inconsistent with crowd-out.

To see if the apparent increase in unrestricted per pupil revenues associated with the 1999 award program and the decrease associated with the 2000 award program were in fact real, we next consider total revenues per pupil. This will help us get around any reporting problems as well as allow us to estimate the net effect of the awards program on revenues. For the 1999 award program (shown in Figure 8a), we estimate a jump in total per pupil revenues in 2000 of about \$340. This estimate is only statistically distinguishable from zero at the 10 percent level. Taken literally, it suggests that for every \$1 per pupil increase in funds due to the 1999 awards program a district actually received closer to \$8 per pupil in revenues. This point estimate is only statistically distinguishable from zero at about the 11 percent level. Moreover, some of the apparent crowd-in may be driven by the additional awards programs that were in place in the 1999 school year. For example, \$100 million was paid to teachers in the form of bonuses as part of the 1999 Certificated Staff Performance Incentive Awards Program. Because the eligibility target for this program was twice that for the GPAP, only schools that qualified for the GPAP could have received it. It is important to point out, however, that these awards were not classified as unrestricted revenue (but rather had a category of its own). Thus, this cannot explain the more than dollar for dollar increase in unrestricted per pupil revenues.

The estimate for the 2000 award program (shown in Figure 8b), while quite imprecise, implies an increase in total per pupil revenue in 2002 of about \$123 per pupil. In other words, contrary to what we might conclude based on the unrestricted revenue alone, this estimate suggests that total per pupil revenues in 2002 may have also increased more than dollar for dollar as a result of the 2000 award program.

Total 2000 expenditures per pupil (shown in Figure 9a) also appear to jump by about \$343 dollars in response to the 1999 awards program. Similarly, 2002 expenditures per pupil (shown in Figure 9b) increase by about \$200 in response to the 2000 awards. This estimate, however, is not statistically significantly different from zero at any reasonable level of significance. Together, however, these estimates rule out the possibility of significant district-level crowd-out, and raise the possibility of significant crowd-in, as a result of California's Governors Performance Award Program.

5 Summary and Conclusion

In this study, we focus on a relatively understudied feature of accountability systems - the productivity of financial rewards for schools making "adequate" progress on state achievement exams. We analyze the the case of California, where for the 1999/2000 and 2000/2001 academic years, schools that met or exceeded their accountability targets were eligible to receive monetary awards through the Governor's Performance Award Program (GPAP). Because these awards were allocated based on a discontinuous function of school (and subgroup) test scores, we employ a regression-discontinuity design to evaluate how they affected achievement and resources. This design allows us to take advantage of the possibility (verified in our data) that schools close to the eligibility threshold are similar but for award receipt. In this framework, award receipt close to the eligibility

threshold is “as good as randomly assigned,” much like in an actual randomized controlled trial.

We find that the program did have a significant impact on the financial resources allocated to some schools. In 1999, the average value of the award was 60 dollars per student and in 2000 about 50 dollars per pupil. More importantly, the financial rewards from the GPAP appear to have been supplemented by payments from other award programs (in 1999) and state budgetary discretion (in 1999 and 2000). At its peak in 1999, districts with schools that qualified for GPAP awards received budgetary increases totaling about 5 percent of per pupil spending.

Despite the increase, we find little measurable improvement in standard metrics of achievement, such as exam performance, for those schools that received the award compared to those schools that did not receive the award. This is perhaps not surprising, as Project STAR, which increased resources by about 50%, yielded improvements in exam performance of less than a quarter of a standard deviation (Schanzenbach 2006). Moreover, because the resources from California’s program are more akin to a random shock than a guaranteed income stream, schools may not have been able to incorporate them into projects that determine educational achievement. However, we also find no increase in “capital expenditures,” such as computers or internet connections, which should be more responsive to a one-time shock.

It is also possible that the instability in the funding of the program weakened its incentives and failed to act as a strong signal of reward for teacher and school administrator effort. Because the awards ended up being distributed mostly as bonuses to teachers (and possibly support staff as well) in an effective school regardless of their individual contribution, the group-based incentives in the GPAP can also have the free rider problem. The free rider problem could give rise to teachers no longer exerting as much effort after award receipt, in the period when we evaluate the schools.³⁰

³⁰This is in contrast to recent work on the impact of paying bonuses to teachers based on schoolwide performance, for instance Lavy(2002).

Our estimates show that accountability “on the cheap” had no significant impact across schools that won awards versus those that didn’t. On the other hand, we cannot assess if the program would have a significant impact if implemented in conjunction with other reforms, such as reduced class sizes or raising teacher salaries.³¹ This also leaves the question of whether the competition for the awards itself made all schools in California more productive. However, a comparison of 1996-2000 4th and 8th grade Math NAEP scores shows California declined or performed the same as the rest of the country during this period.³²

Finally, our findings also suggest that California’s subgroup rules have had the (unintended) consequence of making diverse schools as well as schools that serve disadvantaged populations, more likely to fail their accountability targets. Because meeting these targets are tied to financial rewards, subgroup rules have had the unintended consequence of putting these schools at greater risk of receiving relatively fewer resources.

References

- Baicker, Katherine and Mireille Jacobson, 2007. “Finders Keeper: Forfeiture Laws, Policing Incentives and Local Budgets,” *Journal of Public Economics*, forthcoming.
- Carroll, Stephen J., Cathy Krop, Jeremy Arkes, et al., 2005. “California’s K-12 Public Schools: how are they doing?” RAND Education Report.
- Carnoy, Martin and Susanna Loeb, 2002. “Does External Accountability Affect Student Outcomes? A Cross-State Analysis,” *Education Evaluation and Policy Analysis*, 24(4): 305-331.
- Chay, Kenneth Y., Patrick J. McEwan, Patrick and Urquiola, Miguel, 2005. “The Central Role of Noise in Evaluating Interventions That Use Test Scores to Rank Schools,” *The American Economic Review*, 95(4): 1237-1258.
- CSR Research Consortium, 2002. “What Have We Learned About Class Size Reduction in California?” George W. Bohrnstedt and Brian M. Stecher, eds. September 2002.

³¹California implemented class size reduction beginning in 1996-97 at a cost of up to \$850 per student. While class size reduction is several times more costly than the GPAP, there is no evidence that the program had a significant impact on student achievement. However, evaluating the program is fraught with empirical difficulties. See CSR Research Consortium (2002).

³²Source: NCES, US Dept of Education.

- Figlio, David N. and Cecilia E. Rouse, 2006. "Do Accountability and Voucher Threats Improve Low-Performing Schools," *Journal of Public Economics*, 90(1-2): 239-255.
- Gordon, Nora, 2004. "Do Federal Grants Boost School Spending: Evidence from Title I," *Journal of Public Economics*, 88(9/10): 1771-1792.
- Hahn, Jinyong, Petra Todd and Wilbert Van der Klaauw, 2001. "Identification and Estimation of Treatment Effects with a Regression–Discontinuity Design" *Econometrica*, 69(1), January, pp. 201-209
- Hanushek, Eric A., 2003. "The Failure of Input Based Schooling Policies," *The Economic Journal*, 113 (February): F64–F98.
- Hanushek, Eric A. and Margaret E. Raymond, 2005. "Does School Accountability Lead to Improved Student Performance," *Journal of Policy Analysis and Management* , 24(2): 297-327.
- Hanushek, Eric A. and Margaret E. Raymond, 2003. "Improving Educational Quality: How Best to Evaluate Our Schools," in *Education in the 21st Century: Meeting the Challenges of a Changing World*.
- Hoxby, Caroline, 2002. "The Cost of Accountability." NBER Working Paper 8855.
- Jacob, Brian 2006. "Test-Based Accountability and Student Achievement: An Investigation of Differential Performance on NAEP and State Assessments." NBER Working Paper 12817.
- Jacob, Brian and Levitt, Steven 2003. "Rotten Apples: An Investigation of the Prevalence and Predictors of Teacher Cheating." *Quarterly Journal of Economics*, 118(3): 843-877.
- Jacob, Brian and Levitt, Steven 2003. "Catching Cheating Teachers: The Results of an Unusual Experiment in Implementing Theory." *Brookings-Wharton Papers on Urban Affairs*, pp. 185-209.
- Kane, Thomas J. and Douglas O. Staiger, 2003. "The Promise and Pitfalls of Using Imprecise School Accountability Measures," *Journal of Economic Perspectives*, Fall, pp 91-114.
- Kane, Thomas J. and Douglas O. Staiger, 2003. "Unintended Consequences of Racial Subgroup Rules," in Paul Peterson and Martin West, eds., *No Child Left Behind? The Politics and Practice of School Accountability* (Washington, DC: Brookings Institution Press, 2003), pp 152-176.
- Lavy, Victor 2002. "Evaluating the Effect of Teacher's Group Performance Incentives on Pupil Achievement," *Journal of Political Economy*, 110(6): 1286-1317.
- Lee, David S. and David Card, 2006. "Regression Discontinuity Inference With Specification Error," *Journal of Econometrics*, forthcoming.
- Lee, David S., 2005. "Randomized Experiments from Non-random Selection in U.S. House Elections," *Journal of Econometrics*, forthcoming.
- Rogosa, David, 2003. "Four-peat: Data Analysis Results from Uncharacteristic Continuity in California Student Testing Programs," Unpublished Paper, September 2003.
- Schanzenbach, Diane, 2006. "What Have Researchers Learned from Project STAR?," *Brookings*

Papers on Education Policy, pp 206-228.

Thistlethwaite, D., and D. Campbell, 1960. "Regression-Discontinuity Analysis: An alternative to the ex post facto experiment." *Journal of Educational Psychology*, 51: 309-17.

A Academic Performance Index (API) Database

Established as part of California’s Public Schools Accountability Act of 1999 (PSAA), the API is calculated by the California Department of Education annually for all public schools. These data can be downloaded from <http://api.cde.ca.gov/datafiles.asp>. Small schools, defined as having between 11 and 99 valid test scores, as well as “very small schools” (fewer than 11 valid scores) are evaluated under a separate “Alternative Accountability System.” Growth targets are not calculated by the Department of Education for these schools.

The API ranges from 200 to 1000 and reflects a school’s performance on various standardized tests. Students in grades 2 to 11 are tested. The tests and indicators used and weights accorded to each to calculate the API vary from year to year. Beginning with the 2004 API Base, the weights are applied at the individual student test level; prior to this, it was applied at the school level. The State Board of Education (SBE) “recognizes that the question of the appropriate test weights is a policy issue rather than a technical issue”; its members adopted test weights they believe “reflect the curriculum priorities in California public education.” In practice, each school’s content area weights are determined based on the test weights established by the SBE and also on the number of valid test scores in each content area and grade level at a school. Because of this, API calculations result in content area weights that may be slightly different for each school.

In 1999 and 2000, the API was based primarily on the Stanford 9, a nationally-normed test that is administered annually to California public school students in grades 2 through 11 as part of the Standardized Testing and Reporting (STAR) program. In 2001, the API was based on (1) the Stanford 9 and (2) the California Standards Test in English-Language Arts (CST ELA). In addition to (1) and (2), the 2002 API was also based on (3) the California Standards Test in Mathematics for grades 2-11, (4) the California Standards Test in History/Social Science for grades 10-11, and (5) the California High School Exit Exam (CAHSEE) for high schools. Beginning in 2003, instead of the Stanford 9, norm referenced assessment was based on the California Achievement Test, 6th Edition (CAT-6). Two new tests were also added: (6) the California Standards Tests in Science for grades 9-11 and (7) the California Alternate Performance Assessment (CAPA) for students with disabilities. With every addition/change in test components, the SBE changed the weights in calculating the API and also added a scale calibration factor to minimize the impact of shifts in tests and ensure that “the statewide average API does not fluctuate solely as the result of adding new API components.”³³

Although the components of the API have changed over time, the performance calculation has not. In each year, a school’s “API growth target” is five percent of the distance from the previous year’s API to the statewide target of 800 or a minimum of one point in 1999 or 5 points thereafter, as described more formally in the text.

Schools receive API scores as a whole as well as for each “numerically significant ethnic and

³³For additional details, see API information guides for each year at <http://www.cde.ca.gov/ta/ac/ap/>

socio-economically disadvantaged subgroup in the school.” These subgroups and what constitutes numerical significance are described more fully in the text. The subgroup target is calculated by first multiplying the school-wide target by 0.8 and then rounding the product to the nearest whole number.

B Awards Apportionment Data

Under state PSAA requirements, if a school meets or exceeds its growth target, it may be eligible to receive monetary or non-monetary awards through 3 programs: (1) the Governor’s Performance Award Program (GPAP) ³⁴; (2) the Certificated Staff Performance Incentive Act; and (3) the Schoolsite Employee Performance Bonus (authorized for 1999-00 only). Currently, no funding is appropriated in the budget for monetary awards.

Governor’s Performance Awards (GPAs) were paid for performance in the 1999/2000 and 2000/2001 school years. Awards were suspended thereafter because of budgetary problems. Existing site governance teams or the school-wide council decide how the funds from the GPAs are used, which then gets ratified by the local school board. The data on the apportionments received by each school for performance in the 1999/2000 and 2000/2001 school years can be downloaded from <http://www.cde.ca.gov/ta/sr/gp/ap/apport00a.asp> and <http://www.cde.ca.gov/ta/sr/gp/ap/apport99a.asp>.

The GPA program was funded up to \$150 per pupil, but there is substantial variation in the rewards paid out in our data. There are 4 observations that suggest payouts over \$150 per pupil. Excluding those, the average reward was \$66 per pupil (standard deviation \$8.25) and ranged from \$21 per pupil to \$111 per pupil for 7433 schools over the two school years.

The Certificated Staff Performance Awards (CSP) were paid only in 1999/2000 to schools that made “substantial improvement in their API,” which meant at least twice their API target. The distribution of awards was decided by the local district in negotiation with the teachers’ union. The payment setup was such that teachers from schools with the highest growth received the largest bonuses. In particular:

- 1000 certificated staff in schools with largest growth got \$25,000 each;
- 3750 certificated staff get \$10,000 each;
- 7500 certificated staff get \$5,000 each.

In our study, we evaluate the GPAs only. While we were able to obtain data on CSP payouts, those awards were distributed directly to teachers and would not have directly affected school resources. Moreover, the threshold for receiving those awards was twice that for the GPAs and thus any resource implications for schools from the CSP will be captured in our current estimates. While the Schoolsite Employee Performance Bonus was split between schools and employees, we were unable to locate any microdata on these awards. Since the eligibility was the same for the GPAP, the impact of this program (like the CSP) should also be captured in our current estimates.

³⁴To qualify, elementary and middle schools (high schools) must have a 95% (90%) test participation rate, in addition to the school and all its subgroups meeting or exceeding its API growth target.

C School Characteristics

Data on school characteristics come from the California Department of Education’s California Basic Educational Data System (CBEDS), an annual school-based census. From CBEDS, we collect data on student enrollment, the allocation of teachers across subjects, and a host of demographic characteristics such as racial breakdown, parental education, and percent of students receiving free lunch. These data can be downloaded from <http://www.cde.ca.gov/ds/sd/cb/filessethsch.asp>.

D Fiscal Data

Fiscal data on revenues and expenditures are available not at the school level but at the school district level. We also obtained this data from the California Department of Education. They can be downloaded from <http://www.cde.ca.gov/ds/fd/fd/>. From these unaudited fiscal data, we create variables such as district level total spending per pupil, per pupil expenditures on all staff’s salaries and benefits, teacher salaries, all classified staff salaries, all certificated staff salaries, books and instructional materials expenditures, and other line items.

E Data for Analyses

We created two datasets for the analyses in our paper. The first, which is at the school level, is used to analyze the impact of the award program on school outcomes such as achievement and school-level resources. It was created by merging the API, Award Apportionments, and CBEDS databases together using the unique school identifiers. Our analysis sample includes all schools with valid scores and for whom we could determine eligibility to receive an award. These include elementary and middle schools with at least a 95% test participation rate and high schools with at least a 90% test participation rate.

Since we were also interested in examining the impact of the program on fiscal outcomes, we also created a second district level dataset. To merge API and award eligibility information to the district-level fiscal data, we first collapsed our school-level dataset to the district level. For the purposes of the regression discontinuity design employed in the paper, we characterize each district by the *maximum* of the school-level “awards gap” or eligibility threshold. As described in the text, due to the subgroup rules we characterize each school’s “awards gap” as the minimum difference between the gain scores and targets for the school overall and each of its numerically significant subgroups. Thus, each district’s “award gap” is the across-school maximum of the within-school minimums. The district “award gap” as defined then picks up whether any school in the district received treatment. Finally, we merged this collapsed data to the fiscal data using unique district identifiers.

Table I
Awards Calculation for Two Elementary Schools in AY 2000/2001^a

Mission Elementary, ID No.: 7616486084941

	<i>Students Tested</i>	<i>API_t</i>	<i>API_{t-1}</i>	<i>Gain Score</i>	<i>Award Target</i>	<i>Award Gap</i>
Overall	450	692	682	10	6	4
Black	71	610	589	21	5	16
Amer Indian	4					
Asian	14		Not	Numerically	Significant	
Filipino	9					
Hispanic	122	644	653	-9	5	-14
Pacific Islander	4					
White	225	735	720	15	5	10
Disadvantaged	211	632	632	0	5	-5
Min Awards Gap						-14

Salida Union Elementary, ID No.: 50712666113823

	<i>Students Tested</i>	<i>API_t</i>	<i>API_{t-1}</i>	<i>Gain Score</i>	<i>Award Target</i>	<i>Award Gap</i>
Overall	453	696	664	32	7	25
Black	16					
Amer Indian	1					
Asian	2		Not	Numerically	Significant	
Filipino	7					
Hispanic	218	627	576	51	6	45
Pacific Islander	5					
White	201	763	745	18	6	12
Disadvantaged	212	625	593	32	6	26
Min Awards Gap						12

^aNotes:

1. Only students in grades 2-11 are tested. Parents can obtain waivers for their children exempting them from the exam.
2. A numerically significant subgroup is any of the groups listed above with (i) at least 30 students with valid Stanford 9 scores and at least 15 percent of school's tested enrollment or (ii) at least 100 students with valid Stanford 9 scores (regardless of percent of tested enrollment).
3. The disadvantaged subgroup is not mutually exclusive, i.e. it may contain students counted in other subgroups. A student is classified as socio-economically disadvantaged if (1) he or she qualifies for free or reduced price meals or (2) neither parents has received a high school diploma.
4. In 2000, the awards target = $\max(40 - .05 * API_{t-1}, X)$, where $X = 5$ for schools and 4 for subgroups. This calculation is rounded up and is always based on school's API_{t-1} even for subgroups.
5. The minimum awards gap is the minimum of the awards gaps for a school overall and each of its subgroups. The value must be non-negative for a school to receive an award. School B received about \$50 per pupil or a total of \$42847 for performance in 2000/2001.

Table II
Sample Characteristics by Award Receipt Status^a

<i>Panel A</i>					
<i>Basic Statistics</i>					
	<i>All</i>	<i>No Awards</i>	<i>Award in 1999</i>	<i>Award in 2000</i>	<i>Award Both Years</i>
Percent by Category	100	22.8	30.7	14.7	31.8
Award Per Pupil (\$)	63.1 (7.76)	–	66.5 (2.67)	58.9 (9.80)	62.4 (8.21)
Total Award (\$)	48554 (29487)	–	53742 (33739)	52566 (37973)	45019 (23890)
School Enrollment	856 (606)	1068 (845)	824 (541)	884 (618)	720 (366)
Elementary	69.7	49.3	70.5	65.7	85.3
Middle	17.1	21.5	17.4	21.7	11.5
High School	13.2	29.2	12.1	12.6	3.2
# of Subgroups	1.17 (0.74)	1.34 (0.84)	1.19 (0.71)	1.18 (0.72)	1.04 (0.67)
“Lost” Award	17.9	45.1	15.1	20.3	–

<i>Panel B</i>					
<i>API Scores</i>					
	<i>All</i>	<i>No Awards</i>	<i>Award in 1999</i>	<i>Award in 2000</i>	<i>Award Both Years</i>
School	652 (110)	632 (108)	669 (105)	636 (112)	658 (111)
African Americans	550 (88)	527 (81)	568 (88)	533 (93)	572 (83)
American Indians	579 (98)	530 (76)	653 (113)	667 (63)	583 (86)
Asians	749 (121)	708 (124)	782 (111)	737 (116)	761 (120)
Filipino	743 (70)	709 (67)	758 (65)	742 (63)	770 (65)
Hispanics	574 (86)	551 (84)	589 (84)	558 (85)	582 (85)
Pacific Islanders	585 (120)	475 (30)	696 (19)	–	–
Whites	746 (75)	725 (77)	755 (71)	733 (76)	759 (71)
Socially disadvantaged	581 (86)	555 (84)	595 (84)	570 (88)	592 (84)

<i>Panel C</i>					
<i>API Gain Scores, (API_t - API_{t-1})</i>					
	<i>All</i>	<i>No Awards</i>	<i>Award in 1999</i>	<i>Award in 2000</i>	<i>Award Both Years</i>
School	26 (30)	9.4 (27)	24 (32)	24 (26)	42 (23)
African Americans	27 (38)	9.4 (29)	28 (43)	30 (35)	51 (27)
American Indians	24 (39)	10 (31)	11 (29)	63 (55)	45 (34)
Asians	22 (28)	9.3 (24)	21 (28)	22 (27)	37 (25)
Filipino	20 (28)	7.1 (23)	22 (32)	18 (25)	37 (23)
Hispanics	31 (35)	13 (32)	28 (38)	30 (32)	49 (26)
Pacific Islanders	15 (57)	0 (29)	29 (78)	–	–
Whites	22 (31)	10 (30)	21 (33)	20 (29)	37 (24)
Socially disadvantaged	30 (37)	11 (34)	27 (41)	30 (32)	49 (28)

^aNotes:

1. Standard deviations are given in parenthesis.
2. Zeros are not counted in the award payment calculations.
3. American Indians are “numerically significant” in only 29 school-years and Pacific Islanders in only 8 school-years.

Table III
Impact of the Awards Program on API Scores^a

<i>Panel A: 1999 Awards Program</i>						
	<i>School Overall</i>			<i>Socially Disadvantaged Subgroups</i>		
	<i>2000 API Score</i>	<i>2001 API Score</i>	<i>2002 API Score</i>	<i>2000 API Score</i>	<i>2001 API Score</i>	<i>2002 API Score</i>
Mean	689	697	721	595	615	654
Treatment	-3.56 (5.75)	-3.38 (6.24)	-5.58 (4.99)	.690 (7.00)	-.536 (6.85)	-1.98 (5.37)
F-statistic	1.02	.949	.828	.979	.861	.778
p-value	.416	.679	.960	.567	.916	.989

<i>Panel B: 2000 Awards Program</i>						
	<i>School Overall</i>			<i>Socially Disadvantaged Subgroups</i>		
	<i>2000 API Score</i>	<i>2001 API Score</i>	<i>2002 API Score</i>	<i>2000 API Score</i>	<i>2001 API Score</i>	<i>2002 API Score</i>
Mean		699	723		616	656
Treatment		.502 (3.79)	1.97 (3.92)		-2.24 (4.74)	1.95 (4.78)
F-statistic		.974	1.04		.937	1.00
p-value		.591	.350		.725	.486

^aNotes:

1. Standard errors are in parenthesis and are clustered at the level of a school's distance to the awards threshold.
2. The p-value corresponds to the F-test of the explanatory power of the 5th order polynomial fits relative to the fully flexible model.
3. Differences in the mean API scores for a given year across award program samples occur because some schools evaluated for an award in 1999/2000 are disqualified in 2000/2001 and vice versa. Disqualifications are due to data irregularities, failure to meet required participation rates, and so on. See text for further details.

Table IV
Impact of the Award Program on School Resource Allocations^a

<i>Panel A</i>					
<i>2000 Allocations Relative to 1999 Award</i>					
	<i>FTE Per Pupil</i>	<i>Share FTE Math</i>	<i>Share FTE English</i>	<i>Computers Per Pupil</i>	<i>Internet Connections Per 100 Pupils</i>
Mean	.048	.040	.065	.152	3.32
Treatment	.0004	-.0003	.003	-.005	-.006
	(.003)	(.002)	(.003)	(.006)	(.306)
F-statistic	1.12	.762	.884	1.63	1.07
p-value	.121	.997	.875	.105	.252

<i>Panel B</i>					
<i>2002 Allocations Relative to 2000 Award</i>					
	<i>FTE Per Pupil</i>	<i>Share FTE Math</i>	<i>Share FTE English</i>	<i>Computers Per Pupil</i>	<i>Internet Connections Per 100 Pupils</i>
Mean	.049	.042	.064	.191	4.63
Treatment	-.0006	.0002	-.003	-.020	-.490
	(.0004)	(.002)	(.003)	(.006)	(.290)
F-statistic	2.22	1.10	.961	1.12	1.08
p-value	.0000	.167	.642	.0002	.205

^aNotes:

1. The first row gives the mean of the dependent variables.
2. FTE are full time equivalent teachers.
3. Standard errors are given in parenthesis.

Table V
Impact of the Award Program on District Resources^a

<i>Panel A</i>		<i>2000 Per Pupil Allocations Relative to the 1999 Award Gap</i>			
	<i>Award Apportionment</i>	<i>Unrestricted Revenue</i>	<i>Total Revenue</i>	<i>Total Expenditures</i>	
Treatment	41.8 (3.36)	104 (20.6)	343 (206)	348 (186)	
F-statistic	.471	.721	.877	.716	
p-value	.999	.991	.829	.992	

<i>Panel B</i>		<i>2002 Per Pupil Allocations Relative to the 2000 Award Gap</i>			
	<i>Award Apportionment</i>	<i>Unrestricted Revenue</i>	<i>Total Revenue</i>	<i>Total Expenditures</i>	
Treatment	28 (3.10)	20.2 (6.92)	123 (318)	202 (302)	
F-statistic	.615	.786	1.12	1.08	
p-value	.999	.964	.174	.260	

^aNotes:

1. Standard errors are given in parenthesis.

Appendix Table I
Predetermined Characteristics Relative to the 1999 Awards Threshold^a

<i>Panel A</i>		<i>1999 Characteristics of California Schools</i>				
	<i>Enrollment</i>	<i>Share Free Meals</i>	<i>Share Elementary</i>	<i>Share High Schools</i>	<i>Number of Subgroups</i>	<i>Disadvantaged Subgroups</i>
Mean	842	.471	.706	.124	1.08	.715
Treatment	26 (54)	-.047 (.025)	-.033 (.060)	.008 (.043)	.039 (.059)	-.031 (.047)
p-value	.631	.066	.578	.853	.510	.506

<i>Panel B</i>		<i>Share of Tested Students by Race/Ethnicity</i>				
	<i>Hispanic</i>	<i>Black</i>	<i>Asian</i>	<i>American Indian</i>	<i>Filipino</i>	<i>White</i>
Mean	.380	.081	.080	.010	.024	.418
Treatment	-.015 (.024)	-.018 (.011)	.018 (.010)	-.004 (.003)	.008 (.005)	.005 (.022)
p-value	.543	.125	.068	.232	.116	.835

^aNotes:

1. The first row gives the mean of the dependent variable.
2. Standard errors are given in parenthesis and are clustered at the level of a school's distance to the awards threshold.

Appendix Table II
Predetermined Characteristics Relative to the 2000 Awards Threshold^a

<i>Panel A</i>		<i>1999 Characteristics of California Schools</i>				
	<i>Enrollment</i>	<i>Share Free Meals</i>	<i>Share Elementary</i>	<i>Share High Schools</i>	<i>Number of Subgroups</i>	<i>Disadvantaged Subgroups</i>
Mean	813	.438	.712	.121	1.01	.709
Treatment	-37 (54)	.004 (.019)	.042 (.039)	-.022 (.033)	-.066 (.045)	-.036 (.054)
p-value	.299	.843	.281	.519	.139	.505

<i>Panel B</i>		<i>Share of Tested Students by Race/Ethnicity</i>				
	<i>Hispanic</i>	<i>Black</i>	<i>Asian</i>	<i>American Indian</i>	<i>Filipino</i>	<i>White</i>
Mean	.375	.083	.083	.009	.025	.407
Treatment	-.008 (.016)	-.020 (.010)	.003 (.013)	.0002 (.003)	-.005 (.005)	.028 (.019)
p-value	.626	.044	.819	.990	.343	.127

^aNotes:

1. The first row gives the mean of the dependent variable.
2. Standard errors are given in parenthesis and are clustered at the level of a school's distance to the awards threshold.

Figure 1. API Growth Required to Qualify for Governor's Performance Award as a Function of School's Base API Score: 1999 and 2000 Rules




Figure 2a. Share of Schools Receiving an Award for 1999 Performance Relative to the Distance to the Awards Eligibility Threshold





Figure 2b. Share Receiving an Award for 2000 Performance Relative to the Distance to the Awards Eligibility Threshold



○ Regression Adjusted Probability — Fitted Value

Figure 3a. Per Pupil Award Payment for 1999 Performance Relative to the Distance to the Awards Eligibility Threshold

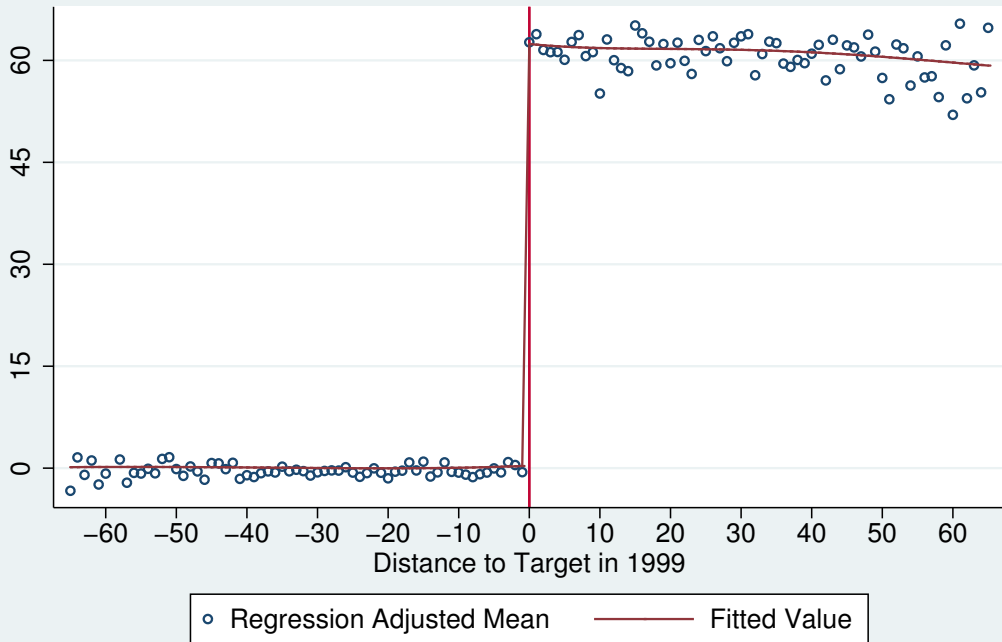


Figure 3b. Per Pupil Award Payment for 2000 Performance Relative to the Distance to the Awards Eligibility Threshold

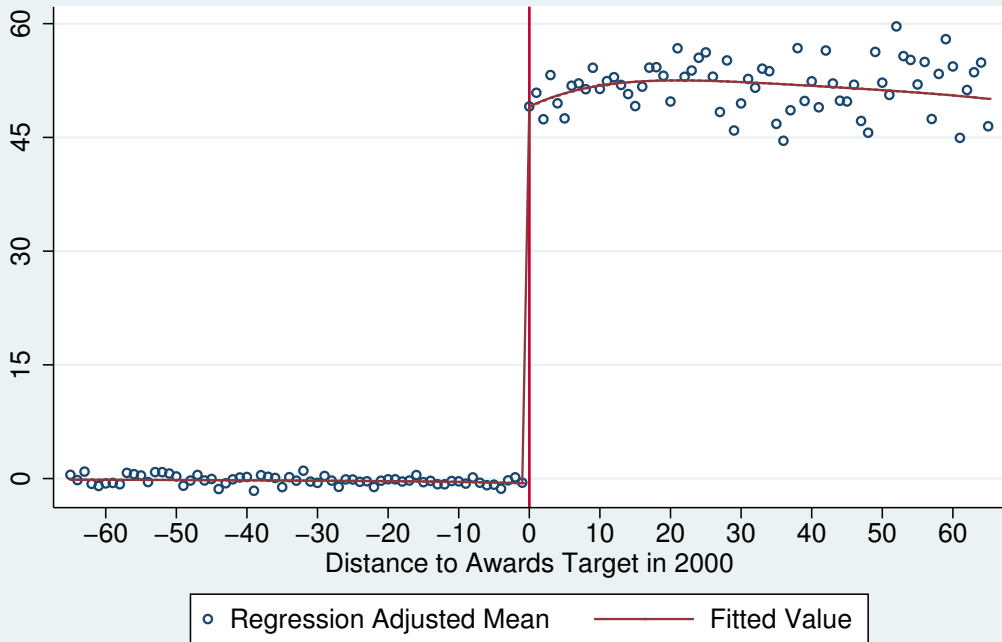


Figure 4a. 2000 API Score
Relative to the 1999 Awards Eligibility Threshold




Figure 4b. 2001 API Score
Relative to the 1999 Awards Eligibility Threshold




Figure 4c. 2002 API Score
Relative to the 1999 Awards Eligibility Threshold




Figure 5a. 2000 API Score for Disadvantaged Subgroups
Relative to the 1999 Awards Eligibility Threshold




Figure 5b. 2001 API Score for Disadvantaged Subgroups
Relative to the 1999 Awards Eligibility Threshold




Figure 5c. 2002 API Score for Disadvantaged Subgroups
Relative to the 1999 Awards Eligibility Threshold




Figure 6a. District-Level Per Pupil Award in 1999
Relative to 1999 Awards Eligibility Threshold




Figure 6b. District-Level Per Pupil Award in 2000
Relative to 2000 Awards Eligibility Threshold

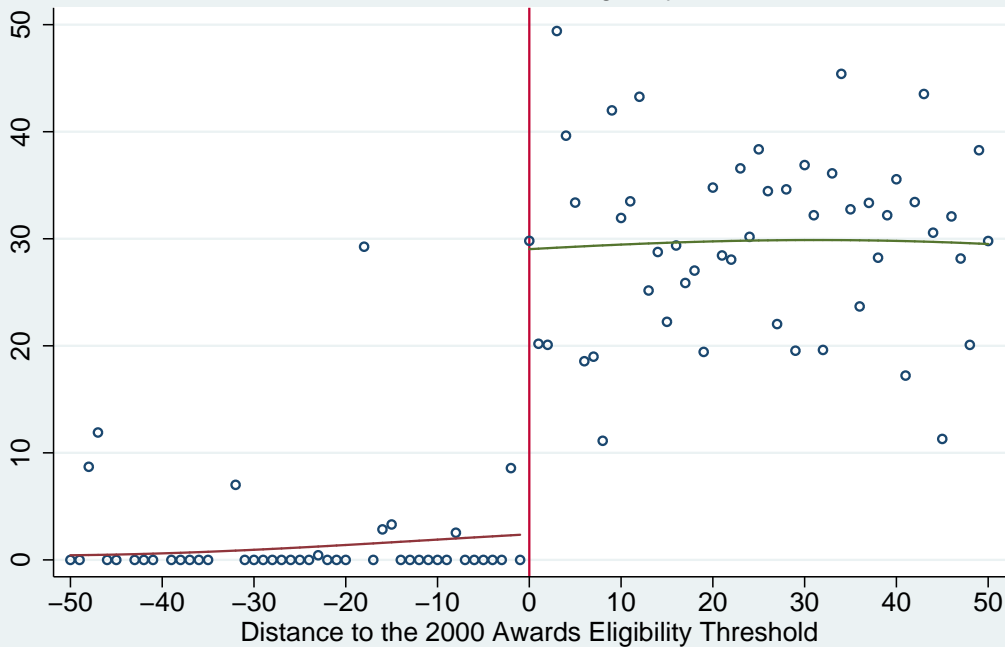


Figure 7a. District Awards Category Revenue Per Pupil in 2000
Relative to 1999 Awards Eligibility Threshold




Figure 7b. District Awards Category Revenue Per Pupil in 2002
Relative to 2000 Awards Eligibility Threshold




Figure 8a. Total Per Pupil Revenue in 2000
Relative to 1999 Awards Eligibility Threshold

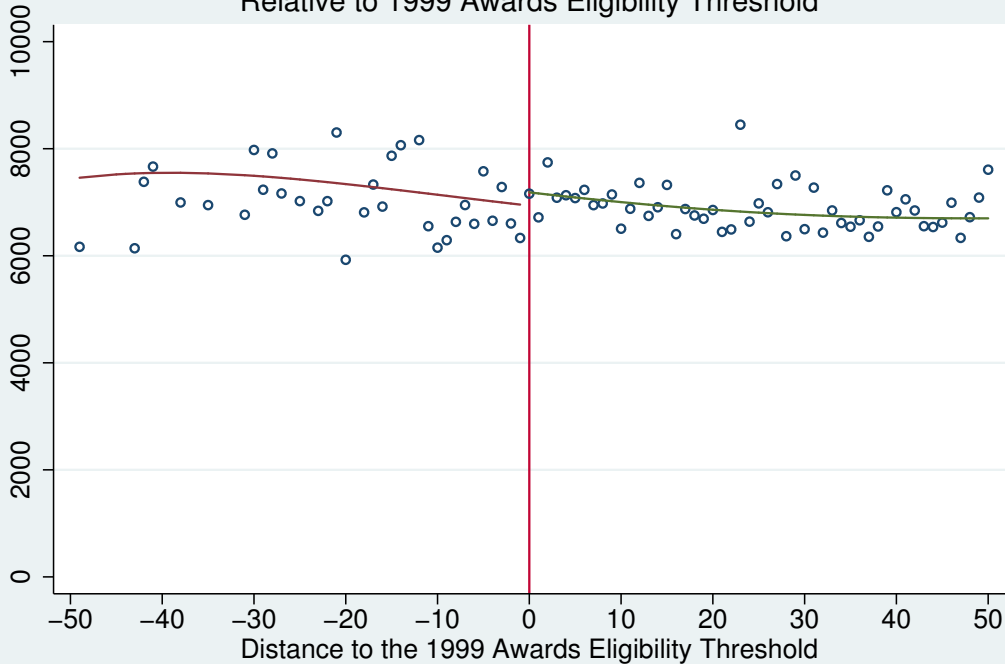


Figure 8b. Total Per Pupil Revenue in 2002
Relative to 2000 Awards Eligibility Threshold

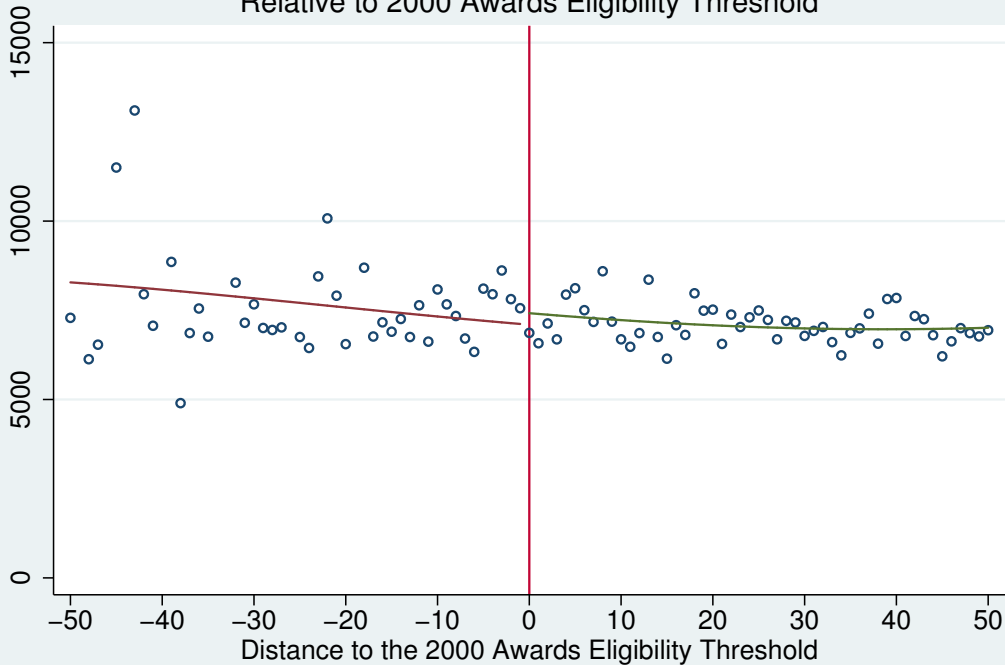


Figure 9a. Total Per Pupil Expenditures in 2000
Relative to 1999 Awards Eligibility Threshold




Figure 9b. Total Per Pupil Expenditures in 2002
Relative to 2000 Awards Eligibility Threshold

