Title

Essays in Public Economics and School Finance

Author

Ugo, Iwunze Lawrence

Publication Date

2023

Peer reviewed|Thesis/dissertation

Essays in Public Economics and School Finance

By

IWUNZE LAWRENCE UGO
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Economics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

_____

Scott Carrell, Chair

_____

Heather Rose

_____

Marianne Bitler

Committee in Charge

2023

Table of Contents

Acknowledgements

Thank you to my wife, Omobhene, for all the support through the years it took putting this work together. Thank you to my family and friends for energizing me along the journey. Thank you to Jon Sonstelie for opening the door to all of this for me. Thank you to Laura Hill for showing me there was a path here. Thank you to Heather Rose for your incredible generosity with your time. Thank you to Scott Carrell, Marianne Bitler and Michal Kurlaender for your patience and advice. Lastly, thank you to all of my teachers and classmates for bearing with me as I asked just one more follow-up question.

Dissertation Abstract

This dissertation consists of three papers on public economics and school finance. Each focuses on school districts in California, exploring how they manage their finances. Given the structure of the California's school funding system, districts in the state uniquely situated compared to school districts elsewhere and other types of local governments more broadly. The results of the analyses, however, are instructive and could be relevant in other contexts.

In the first chapter, I explore categorical funding programs—additional funds available to applying districts that are targeted to specific purposes or student groups. These are the most significant of the limited options for California school districts to raise revenues. The use of these funds is restricted, however, and districts often face a choice between altering their educational program and forgoing the new revenues. This paper analyzes the impacts of the availability of categorical funds on school district decision-making. Using detailed school district financial data and focusing on results for four major programs, I use a difference-in-difference research design by exploiting the timing of each program's introduction or elimination and find that district spending decisions are significantly influenced by these categorical programs, as participating districts react by shifting their spending patterns to accord with program goals. The influence of categorical funding on district spending patterns is accompanied by mixed results for student outcomes. Suggestive evidence shows that some programs have benefits for students in participating districts while other programs have minimal impacts on student achievement.

The second chapter examines school district savings. In any given year, California school districts hold a significant portion of their revenues in reserves. While some amount of saving is fiscally prudent, such high savings raise a question of whether this reserve behavior is appropriate. Revenues held in reserve are not spent on services for the students that generated them and research in other contexts has raised concerns about agency issues—the self-interested use of savings on

the part of managers potentially at the expense of the broader organization. In this paper, we explore the determinants of reserves in California school districts. With estimates of these determinants, we then identify districts that hold significant excess reserves and examine whether these districts exhibit agency issues. We find that revenue volatility, total enrollments, and declines in enrollment are key determinants of district reserves. District that hold substantial excess reserves also spend more on administrative expenses relative to other districts, indicating some potential for agency issues. However, further analysis finds that these high-saving districts also have higher average student achievement, suggesting that their reserves behavior and the associated administrative spending may actually reflect a premium paid for higher quality management.

The third chapter is on the volatility of revenues. California school districts have little control over their revenues—most rely heavily on state aid for the bulk of funding. This funding fluctuates from year to year, and in the case of recessions, can be subject to substantial reductions. Typically, income or revenue volatility suggests a need to smooth consumption in order to avoid disruptions when shortfalls occur. Using nearly 30 years of detailed financial data, we explore how districts manage this variation and the associated uncertainty. We find that, on average, California districts spend nearly all new revenues, responding myopically to annual changes. There is, however, substantial variation between districts in the management of revenue volatility. These differences in management are related to district characteristics like region or grades served, as well as student demographics and the level of available resources. We continue our examination and find that differences in management are related to student achievement and class sizes.

Chapter 1: The Influence of State Funding Restrictions on Local School Policy

1.1 Introduction

School district revenues in California are primarily composed of local property taxes, general state aid and restricted grants called categorical programs. Most states combine broad unrestricted grants and targeted categorical programs to fund local school districts (Smith 2013). Revenues for categorical programs come with a variety of restrictions. Some programs only provide revenue to districts according to certain characteristics or their enrollment of students in particular demographic groups, while others require funds be spent in prescribed ways (Timar 2010). School districts in the state have had access to several dozen programs, each with their own levels of available funding and sets of restrictions. While every program offers the obvious benefit of added revenues, districts must choose whether to apply for any given program and accept the requirements that come with them (Weston 2011a, 2011b). These choices offer insight into the way that school districts make decisions about the kind of educational program they offer. Districts may alter the programs and services they offer in order to conform to the requirements of a given program and secure the added funding. Alternatively, they may commit to a particular set of services and only pursue categorical programs that accord with their established priorities.

I analyze the ways in which California school districts respond to the incentives created by the availability of categorical funding. Over the last two decades, many programs have been introduced, eliminated, or both. I use a difference-in-difference research design to compare spending patterns over time in districts that participate in various programs to those that do not. By analyzing the different spending decisions that participating districts make in the years before, during and after a program is available, I can assess what effect, if any, categorical programs have on district spending priorities. The availability of a program is out the control of individual districts, so the establishment of new programs and the discontinuation of others represents an

exogenous change in the funding environment. Importantly, the change is different for districts that participate in a given program and those that do not. I then continue the analysis to consider the relationship between these effects and student outcomes. Given that there have been dozens of categorical programs over the last two decades, I focus on four categorical programs in my analysis that I selected using an approach described below. The programs I include are the Grades 9-12 Class Size Reduction Program, the College Readiness Block Grant, the Economic Impact Aid Program and the Teacher Credentialing Block Grant. Altogether, they offer examples of the broad types of categorical programs available.

In each of the programs, participating districts did spend more on the activities associated with their programs, thus showing that program participation did tend to align district spending patterns with the goals laid out in particular categorical programs. My analysis of spending trends over time provides evidence that it was the programs and their additional revenues driving school policy rather than districts themselves. In most cases, participating districts do not show an established preference for the activities associated with a program prior to it being available. And after a program is eliminated, so too is the spending differential in those associated activities between participating and non-participating districts. This finding indicates the leading role that categorical programs play in determining district policy.

Given the significant impact categorical programs have on school district policy, a natural question is whether the effort to shift local program offerings through state-level policy is worthwhile. I extend my analysis beyond the results for district revenues and expenditures to analyze how the changes in spending relate to student outcomes. I investigate whether these programs have meaningful benefits for students in participating districts. I find mixed results with

some programs showing significant benefits for affected students on related outcomes, while others are relatively muted.

In this paper, I make contributions to the school finance literature, and more broadly, to the literature on public finance—in particular, on questions concerning the administration of subnational governments. The relationship between school districts in California and the state government is unique compared to other states around the country and has shifted in many ways in recent decades. The example of California is particularly instructive for the broader literature because of the extent to which districts rely on the state government for much of their funding. This paper also contributes by using the state's district financial data in a robust way that goes beyond revenue and expenditure totals used in other work, and instead leverages the vast amount of detail in the data in order to glean insights about how school districts manage their funds in a constantly changing economic environment.

The next section, Section 1.2, provides an overview of school district revenues in California and describes where categorical programs are situated in the funding mix and their history in brief. Section 1.3 covers the detailed financial records and methods I use to conduct the analysis as well as the process that I use to select the categorical programs I focus on and their associated activities. Section 1.4 provides results on the effect of program participation on school district revenues and expenditures. Section 1.5 extends that analysis to student outcomes. In Section 1.6, I conclude.

1.2 California School District Revenues

California's school finance system has experienced a number of substantial changes in recent decades. Before the first set of major reforms in the late 1970s, schools were largely funded by collections from local property taxes. Though the state provided some foundation aid, the key

characteristic of the system was the significant level of inequality in per pupil funding across districts related to differences in fiscal capacity–that is, the ability of districts to collect tax revenues. This inequality was confronted directly by the California Supreme Court's decisions in the Serrano v. Priest cases. The plaintiffs in the case argued that the state's reliance on local property tax revenues to fund public schools produced differences in per pupil funding across districts that were large enough to effectively deny students an equal right to education. Siding with the plaintiffs, the Court required that the state's education funding system be reformed to equalize per pupil revenues across school districts. While equalization efforts carried on with mixed success over the next several years, the state's political context was also undergoing rapid change. Rising property values came along with rising property tax burdens and fueled a movement of fiscal retrenchment that led to the passage of the statewide initiative, Proposition 13. The ballot initiative addressed the issue of rising property taxes by imposing a 1% statewide cap on property tax rates and limiting the amount that assessed property values could rise each year.

Prop. 13 had the effect of dramatically reducing property taxes available to school districts and led to the formalization of school finance system called revenue limits. Most of the funding in the revenue limit system came from the mixture of local property taxes and general state aid that was meant to equalize funding differences across districts. Both of these sources, however, were largely out of a district's control. Property tax rates were restricted by Proposition 13, while the state aid was distributed according to an equalization formula.

Other sources of revenue included categorical programs, federal revenues and special local revenues like parcel taxes and sales. Categorical program revenues stood out as the largest among these other sources. They also have the notable feature of being discretionary. While district revenue limits are largely out of their control, categorical programs offer the most significant way

for districts to engage in fiscal policy and choose the level of revenues they desire. Of course though, this is mitigated by the number of programs a district is eligible for, the amounts of revenue available and their willingness to spend in ways that align with the goals of various programs. Given some of the history described above, categorical programs are a particularly important part of the funding mix for school districts in California relative to the rest of the nation. Districts—almost universally in need of funding—use many programs at high rates, and ambitious state legislators constantly add new programs in order to fund various initiatives.

Categorical programs have had a turbulent history throughout the decades since the establishment of the revenue limit system (Sonstelie 2000). The number of programs available as well as the amount of revenues they offered grew through the 1980s and 1990s. They were also subject of substantial political wrangling given their ability to partially reverse or amplify the equalization efforts in the state's school finance system (Timar 1994). At their peak, categorical programs played a major role in California school district budgets. In 2007, there were over 100 programs that altogether generated revenues of over 17 percent of the state total.

Two episodes then saw major reductions in categorical programs and program revenues. First, as a response to the Great Recession, the state legislature passed the Categorical Flexibility Act in 2009. The bill sought to make up for reduced state aid by allowing districts flexibility to spend revenues from categorical programs for any educational purpose. This effort essentially eliminated those programs by removing the restrictions that defined them. Then, in 2013-14, the state's school finance system was completely overhauled with the introduction of the Local Control Funding Formula (LCFF). The LCFF included a simplification of the system that eliminated nearly all of the remaining categorical programs in favor of a combination of base grants and additional funding for districts based on the number of high-need students they enrolled.

Figure 1.1 Distribution of School District Revenues (2003-2018)

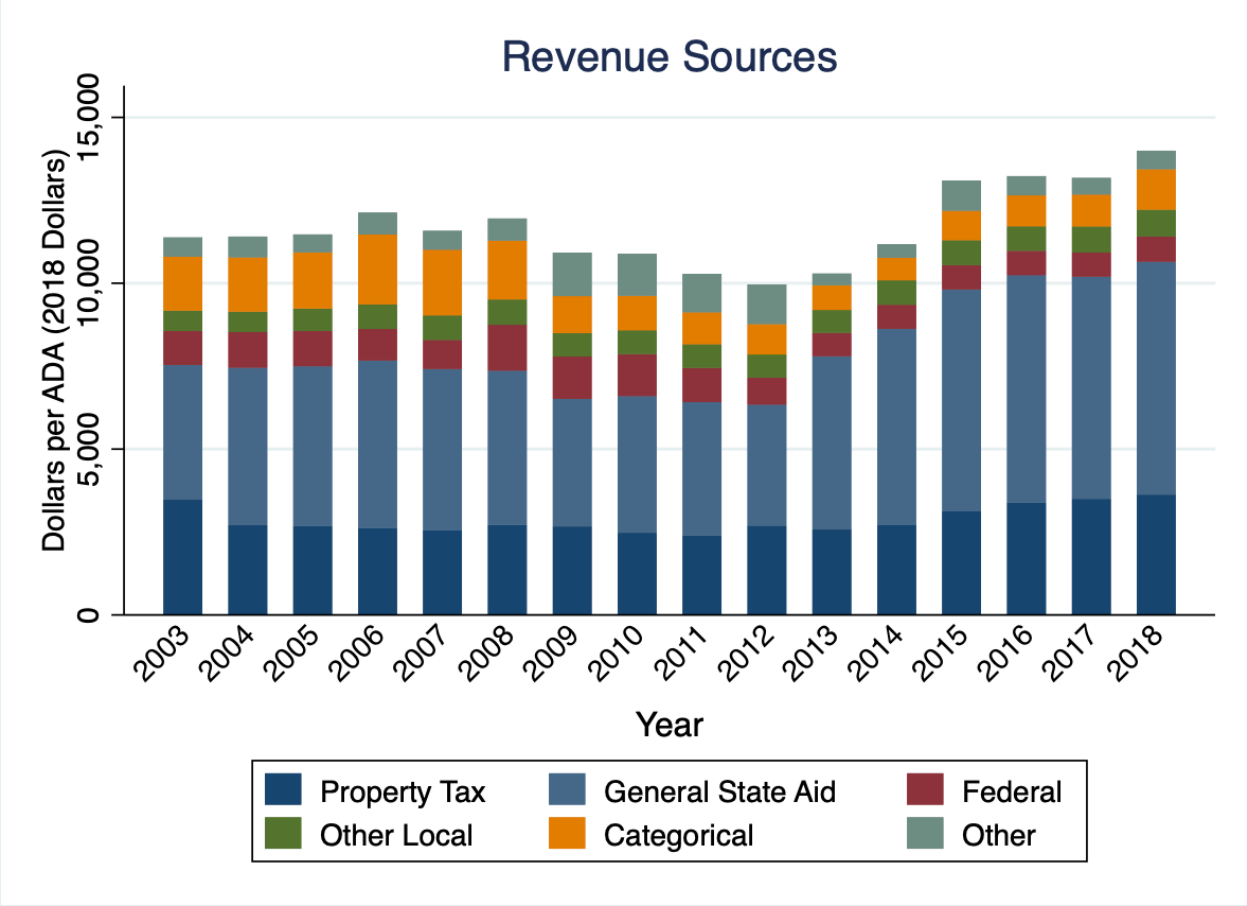

Figure 1 above shows district average per pupil revenues through much of the last two decades and the distribution of these revenues across the major sources. Starting in 2007, the graph shows both the decline in overall revenues during the Great Recession as well as the shift away from categorical programs in the revenue mix[1]. Overall spending rebounds quickly in 2013-14 with the introduction of the LCFF, increases in the Prop. 98 guarantee and the new revenues from the Education Protection Account established by passage of the Proposition 30 and its increased tax rates. However, the role of categorical programs remained diminished. Figure 2 shows this in greater detail, using box-and-whisker plots to track the distribution across districts over time of

[1] Property taxes and general state aid are the primary sources of revenue for most California school districts. Federal funds and categorical program revenues play a major role as well. "Other Local" in Figure 1 refers to other sources of locally-raised funds beyond the district's ad valorem property tax. These include parcel taxes and sales (i.e. rental income from leases of district property) as well as some, though not all, private donations. For more information, see Weston (2015).

the share of revenues coming from categorical programs. The edges of the box give the 25th and 75th percentiles of the distribution, while the center line gives the median. The whiskers give the extreme values, while some outliers are removed. The graph shows the range of the impact that categorical revenues had on individual school district budgets, with some districts deriving substantial shares of their revenues from them, while others relied on them very little. This variation in categorical program use is what allows my difference-in-difference research design.

Figure 1.2 Distribution of Share of Revenues from Categorical Programs

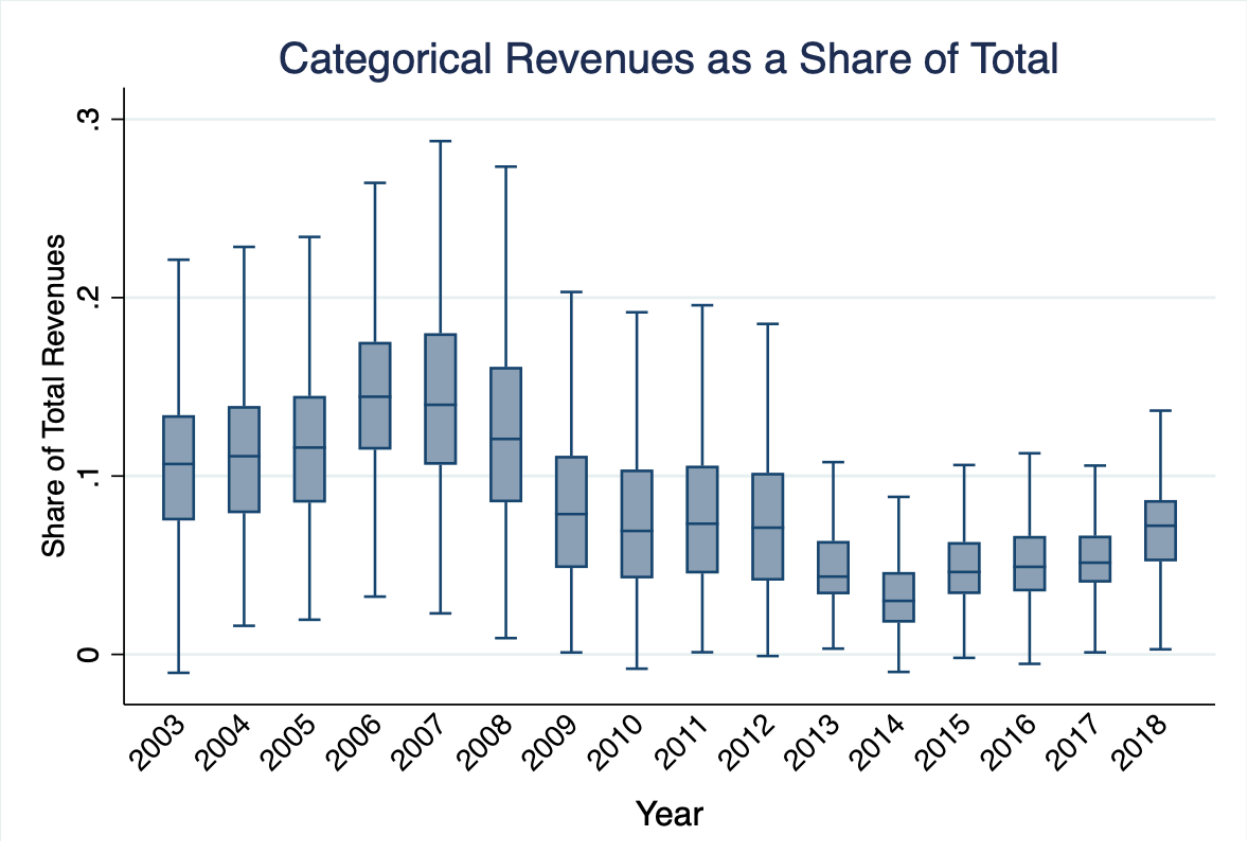

1.3 Data and Methodology

1.3.1 Data Sources

The main source of data for this paper is the annual financial data available from the California Department of Education. The Standardized Account Code Structure (SACS) data provide detailed financial information for the state's districts[2] including revenues, expenditures and remaining annual balances. The data, spanning from 2003-04 to 2018-19, include a variety of details like sources of revenue and spending on various activities, which facilitate my analysis of the impact of categorical funding on school district decision-making.

The SACS data describes revenues by their source, with categories including property taxes, state aid and federal grants. This detail extends even to the specific categorical programs generating revenues. Details on specific program revenues allow me to identify which programs districts apply for and take advantage of, while also giving the amounts of revenues they receive from various programs.

On expenditures, the SACS data include codes that identify expenditures on specific goods and services like staff salaries, books, materials, equipment and more, with added detail on the intended functions like instruction, administration, maintenance or security. I use this level of detail to identify the activities that are the main uses of the funds generated by various categorical programs. The structure of the data allows a key flexibility on two counts. I can examine spending on these main activities of a program in districts that participate in a program as well as those that do not, and I can also track this spending both in the years that a program is available as well as those either before the program was introduced or after it was eliminated.

[2] The analysis here is limited to school districts, and thus not inclusive of charter schools. These independent public schools largely report their finances in a different system. Further, charter schools are not eligible for most categorical programs. Instead, they receive a general purpose block grant meant to equalize total funding given their lack of access to most categorical programs.

The finance data are supplemented by data on a variety of metrics also publicly available from the CDE. We use student demographic information at the district level on race and poverty, which is measured by participation in the free- or reduced-price meals program. Participation in the school meals program only provides an approximation of student family income. Rates are sensitive to differences in take-up of the program, and the binary nature of the measure does not capture gradients in income both below and above the eligibility cutoff. We measure enrollment by average daily attendance (ADA), which gives the number of students in school on average rather than the slightly higher number of students counted on census day in October. ADA is used as the pupil count for the basis of funding for most programs. Despite that however, on the questions of district financial management that we explore here there is valid concern that total enrollment rather than average attendance ought to drive our analysis. After all, districts ought to have resources available for all of the students that may come on any given day, independent of the usual pattern of absences. Unfortunately, data on enrollment that matches the districts as they are formulated in the financial data are not readily available.[3] Checking the relationship between ADA and the enrollment data we have, we find that the ratio of ADA to enrollment is somewhat relatively constant at about 0.92. Student achievement is measured by performance on the state standards test. The California Standards Test (CST) was used from the 2002-03 to 2012-13. The LCFF reforms included a transition to a new test, the Smarter Balanced Assessments—often referred to as the SBAC tests. The new exams differed from the previous system in a number of

[3] Enrollment data is available at the school level while the SACS data are collected at the district. The school level data include counts for charter schools within a district, but it is difficult to accurately account for charter schools in the financial data given that the way in which charter revenues and expenditures are reported varies based on individual school and district practice.

ways. They were aligned with the state's new Common Core curriculum, and they were computer adaptive.[4]

1.3.2 Methods

I conduct the analysis in this paper using a difference-in-difference research design. With the detailed financial data described above, I identify which districts participate in any given program and compare their revenues and expenditures to districts that do not participate in that program. These non-participating districts serve as a control group in this design. Observing district spending in years where a program is available as well as in years where a program is not is key to the quasi-experimental interpretation of our setup. Given that individual districts do not control when new programs are introduced, and certainly not when they are eliminated, they must take the availability of these programs at one time or another as exogenous to the decisions they make in a given year. These revenue and spending patterns over time stand in for the counterfactual trend in participating districts had they not taken up a program. To better facilitate this comparison, I control for several district characteristics that could plausibly relate to a district's decision to participate in a program. By adjusting for factors that could confound the program participation decision, the remaining comparison is closer to the causal impact of a district choosing to take up a program. The difference-in-difference design requires that the trends in our outcome variable be similar across treatment and control districts in advance of the policy change. Similarly, other factors that could influence the pre-trend must be investigated as well. These include evidence of anticipation for treatment or control districts and other kinds of heterogeneity separating participant districts from the others.

[4] "Computer adaptive" refers to the manner in which questions are presented to students taking the SBAC. Students all answer different sets of questions. The difficulty of questions is adjusted dynamically based on whether students answered the previous questions correctly. This is meant to reduce the time required for the examination and produce more accurate assessments of student learning.

Throughout the analysis, I rely on a regression specification of the form,

$$Y_{it} = \beta_0 + \beta_1 Post_t + \beta_2 (Post_t X Part_i) + \beta^k X_{it}^k + \lambda_i + \sigma_t$$

where Y_{it}, the outcome, represents in a given district, i, and year, t, revenues, expenditures and other outcomes I relate to district program participation. As in a typical difference-in-difference model, the causal effect of program participation is identified by the interaction between a binary variable indicating years before and after a change in program availability—here given by $Post_t$–and another binary variable, $Part_i$, indicating districts that participate in a program when it is available. The model extends the typical specification to a generalized difference-in-difference model, including multiple years in the time variable, which allows me to assess how the impacts of program participation evolve over time. X_{it}^k represents a vector of controls for district characteristics and student demographics. These include total enrollment, recent enrollment growth, the presence of charter schools in a district, and the share of low-income students in a district as well as distributions of student across race, gender and grade.

Although the research design outlined above is meant to produce causal estimates of the impact of program participation, I consider here potential issues. Clearly, a district's characteristics or the characteristics of its students likely influence whether that district participates in a particular categorical program. Most explicitly, whether a district is eligible for a program at all often relates to those characteristics—an elementary school district would not be eligible for the Grades 9-12 Class Size Reduction Program, for example. Beyond that though, district characteristics could translate to a higher likelihood of participation rather than an explicit requirement. For example, a growing district would be more likely to apply for a program to train new teachers than one with stagnant enrollment counts. I control for a variety of district characteristics in the specification detailed above, but the list is not likely to be exhaustive of every potential factor that could

influence program participation. Observed differences in outcomes may be due to differences in these unobserved factors that just happen to be confounded with program participation.

Another potential issue relates to the main results in my analysis, the effects of program participation on spending. In the results section below, I compare spending patterns between participating and non-participating districts in the particular activities that are associated with a given program. The class size reduction program mainly funds spending on teacher salaries, for example. Spending in this activity may come from any of the several revenue sources available to school districts like unrestricted revenues, the program of interest or any other categorical program. The spending differences between participant and non-participant districts in an activity would ideally be driven by the program of interest. However, these differences could instead be related to differences in spending from other categorical programs or spending out of unrestricted revenues that is simply unrelated to the programs I focus on. To address this, I consider spending differences coming out of other categorical programs separately. Also, since the regression identifies mean differences between participant districts taken as a whole and those that do not use a program, it is not likely that the set of districts participating in a program would systematically spend different amounts on the activities associated with a given program out of their unrestricted funds for reasons totally unrelated to the program of interest.

1.3.3 Selection of Categorical Programs

With so many programs, each having a variety of different characteristics, I narrow my focus to a limited number for clarity's sake. I choose programs that have significant shares of both participating and non-participating districts, as well as those that are introduced, eliminated or both within the years of available data. Within these criteria, I select programs that are large in terms of

the revenue they generate for participating districts or those that are diverse in the policy areas they are intended to address.

Participating districts are defined as those that receive positive revenues from a particular program. In order to ensure reasonably sized treatment and control groups–districts that participate in a program and those that do not–I select programs in which at least 25 percent of districts receive revenues from a program and 25 percent of them do not. The actual rates of participation and non-participation may vary from 25 to 75 percent, but the research design requires that neither of the two groups be too small for the sake of comparison.

This requirement is not overwhelmingly restrictive. It excludes many rather niche programs that only serve a few districts and generate little revenues as well as other programs that are nearly universal in take-up and thus not discretionary–at least in practice.

I select programs that are not available in every year across the time span of the data available. Given the history of categorical programs over the past two decades, though, this restriction is not extremely limiting either. Categorical programs saw rapid growth in the early 2000s with new programs added each year (Weston 2009). Though individual programs lapsed for various reasons over the years, the two major episodes described above—the Categorical Flexibility Act and the Local Control Funding Formula—saw many programs eliminated. These two waves of program elimination did not entail reductions in revenue for participating districts, but rather just removed the explicit connection of particular buckets of revenue to the purposes laid out in particular programs.

Table 1 below shows some examples of programs that are included in the analysis here and others that are not. I look at programs in their peak years–that is, when the share of participating districts and, often, the total amount of revenues generated by the program were the highest. Some

programs are nearly universal, while others are more niche even though they may generate substantial revenues for the few districts that participate in them. There are also some major programs ideally suited for this analysis. These programs have substantial shares of both participating and non-participating districts, as well as useful spans of years in which they are available and not. On top of that, these are major programs, generating several millions of dollars of revenues and substantial shares of revenues for participating districts relative to other, smaller programs.

Table 1.1 Selected Categorical Programs

Program	Year Introduced	Year Eliminated	Share of Participating Districts	Total Program Revenues (millions)	Average Program Revenue Share
Special Education	–	–	93%	9,138	10.3%
Targeted Instructional Improvement	–	2010	7%	1,132	2.5%
Grades 9-12 Class Size Reduction	–	2009	25%	158	0.6 %
Economic Impact Aid	–	2013	69%	521	1.5%
Teacher Credentialing	2005	2009	30%	119	0.2%
College Readiness	2015	–	44%	166	0.5%

After searching through the available programs based on the criteria above, I settled on four programs: the Grade 9-12 Class Size Reduction Program, the College Readiness Block Grant, Economic Impact Aid and the Teacher Credentialing Block Grant. The class size reduction program is one of the largest categorical programs. As its name suggests, the program provides funding to districts for them to reduce class sizes in the high school grades. It is broadly available to districts given that a district only needs to serve high school students to be eligible for funds (a similar program is available for districts that serve students in grades K-3). The College Readiness Block Grant is newer than the other programs considered here, having been introduced in 2015 after the transition to the LCFF. It provides funds for districts—with higher funding levels for districts with higher shares of disadvantaged students—to improve the rates at which their students are ready and go on to college. The Economic Impact Aid program was a precursor to the LCFF, providing funds to districts in proportion to their enrollment of high need students that the districts could use for a variety of purposes aimed at improving outcomes for these disadvantaged students. Finally, the Teacher Credentialing Block Grant provided funds to districts to support mentorship opportunities for beginning teachers. The program is somewhat niche compared to the other programs mentioned here, but it provides an example of a program that is tied very explicitly to a specific purpose.

1.3.4 Selecting Activities Associated with Particular Programs

Much like how I use information about the characteristics of various programs to choose which ones to focus on for this paper, I employ a similar process to pick among the variety of activities that districts fund with revenues from any given program. I use two measures to drive my choice of activities to focus on. First the activities funded by a given program are ranked by the share of total spending in the program that comes from a particular activity. For some programs,

this is incredibly straightforward. The class size reduction program for example sees most of its funding spent on teacher salaries with nearly all of the remaining funds spent on benefits for those teachers. Other programs, though, have much more varied spending. The funding from the economic impact aid program for example is used on everything from teachers and administrators to materials and consultants.

The second measure I use looks instead at how the use of activities is distributed across districts. Rather than looking at the share of program funds going to a given activity, I use the share of participating districts funding a particular activity. This second measure has a number of advantages. Focusing on district use allows me to avoid the popularity of any given activity being heavily influenced by a few large districts that may not spend the funds from a program in a way that is broadly similar to most other districts. Also, given that different activities have different costs—teacher salaries cost more than textbooks for example—shifting my focus away from dollars spent prevents the most expensive activities standing out above other activities that would better describe the typical use of revenues from a given program. Details on the programs chosen and the main activities I consider are shown in the table below.

Table 1.2 Categorical Programs and Main Activities

Program Name	Start/End Year	Mean Program Revenues Per District	Main Activity	Share of Program Revenues in Activity	Share of Districts Using Activity
Grade 9-12 Class Size Reduction	2009 End	$672,427	Teacher Salaries	79.9%	95.6%
College Readiness Block Grant	2015 Start	$934,557	Instructional Materials	32.6%	88.2%
Economic Impact Aid	2013 End	$1,120,542	Teacher Aides	12.2%	83.7%
Teacher Credentialing Block Grant	2005 Start / 2009 End	$491,052	Instructional Support	18.1%	23.9%

1.4 Results

Using a difference-in-differences design, I compare revenues and expenditures between districts that participate in a program and those that do not. These comparisons are over several years spanning both when a program is available and when it is not. For each of the four programs described above, I begin with a comparison of revenues between participant and non-participant districts. I consider differences in total revenues and differences in revenues from the program of interest. These models are meant to show that there is little difference between participating and non-participating districts in revenues from most sources, though only participating districts receive funding from the program of interest. Next, I show results comparing spending in participating districts to non-participating districts. In the data I use, spending in an activity can be linked to the revenue source that it came from. Most spending comes from unrestricted sources,

but a substantial amount of spending is linked to the categorical programs that generate their revenues. Districts participate in any number of programs, some of which may have overlapping objectives with the particular programs that are the focus of my analysis. I show two sets of results. The first compares spending from only unrestricted sources and the program of interest, and the second compares spending from all other categorical programs. The first graph presents the main result: how the availability of categorical programs affects spending out of what would otherwise be unrestricted funds in participating districts. The second graph provides more context, showing the extent to which spending patterns out of other programs increase or reverse the differences between districts that participate and those that do not.

1.4.1 Grade 9-12 Class Size Reduction

The Grade 9-12 Class Size Reduction Program was introduced decades ago before the start of my time series of data. The program continued in the same form until it was eliminated in the 2009-10 school year with the Categorical Flexibility Act. Though districts still received the same levels of revenues they had, they were no longer required to spend that money according to the rules of the now-defunct program. Figure 3 shows the evolution over time of the difference in revenues of various types in participating and non-participating districts. Note the absence of major trends in differences in most other forms of revenues. Property taxes and total categorical program revenues are broadly similar for participant districts during and after the class size reduction program was available, while the difference in general state aid declines over the years–but not discontinuously. General state aid does rise, however, with the introduction of the LCFF in 2013. Program revenues, though, see a steep drop as the difference for participating districts goes to zero once the program is no longer available in 2009. This stable context is key to the identification of the spending effects shown in Figure 4.

Figure 1.3 District Revenue Comparison by Participation in Grades 9-12 Class Size Reduction

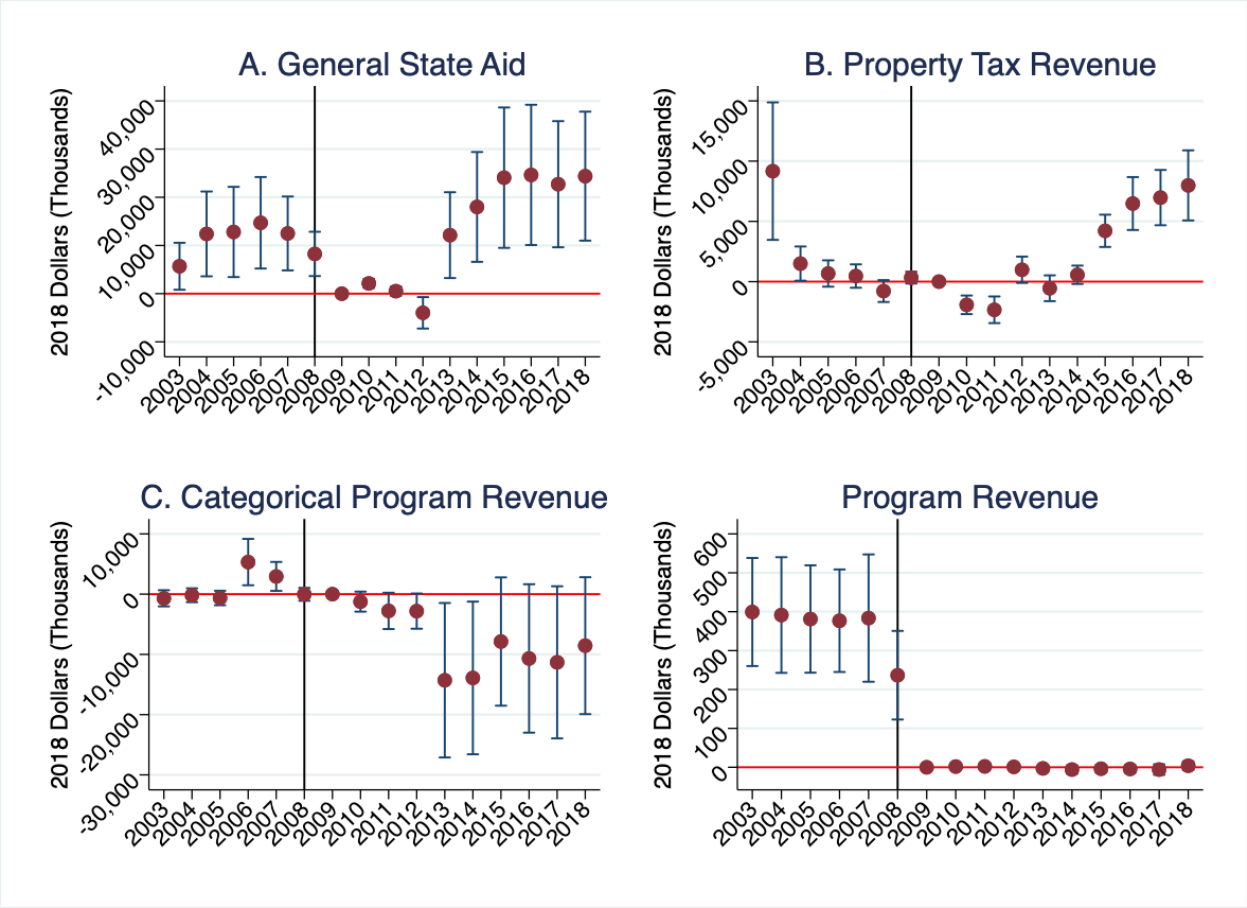

The graphs in Figure 4 tell a somewhat similar story as those in the previous figure. Panel A shows the difference in total spending on teachers, while Panel B shows the difference specifically from the combination of unrestricted revenues and the class size reduction program. In Panel C, the differences in spending on teachers from all other categorical programs are shown for context. There is a stark decline in the difference in spending for participating districts when we consider teacher spending from unrestricted and program funds. An average difference of nearly $3 million in spending on teachers was reduced to essentially zero immediately after the class size reduction program was phased out. Notably, this difference consistently hovered around $5 million throughout the years the program was available, before dropping to zero and remaining there for years afterward. There was little difference in teacher spending from other categorical

programs during the years the class size reduction was available. The spending difference here was also stable after the elimination of the class size reduction program before a consistent declining trend began in 2010-11.

Figure 1.4 District Spending Comparison by Participation in Grades 9-12 Class Size Reduction

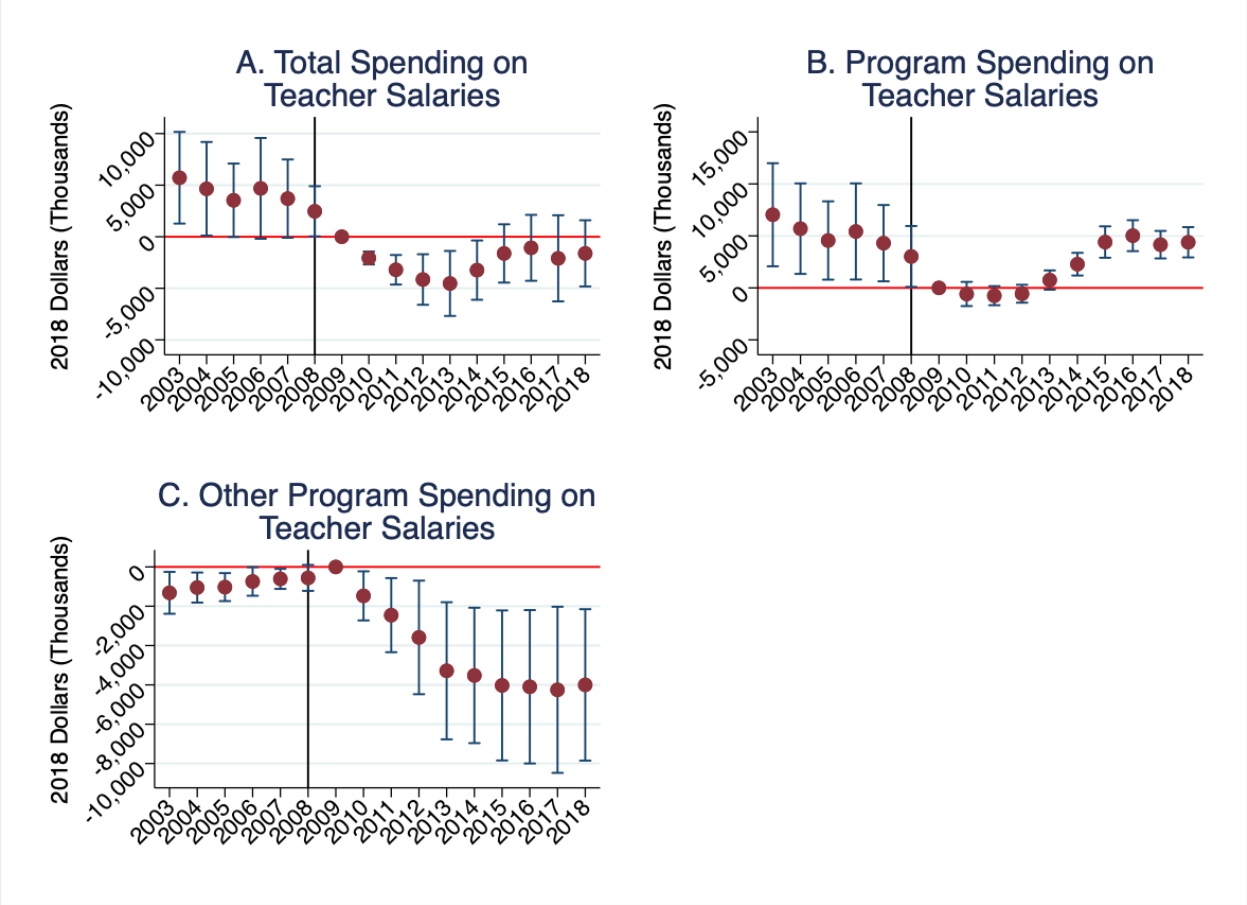

1.4.2 College Readiness Block Grant

The College Readiness Block Grant is the newest of the programs considered here. It provides funding to participating districts in order to increase the rates at which their students are prepared for and continue on into higher education. Larger grants are available for districts with high shares of disadvantaged students. In Figure 5 below, a comparison of revenues shows that participating districts typically had similar amounts of property tax revenue, lower levels of

general state aid and higher amount of revenue from categorical programs. Importantly, these differences are mostly constant over time, and in particular, stable around the time the program was introduced in 2015. Participating districts received about $500,000 in the second year of the program, with revenues rising to an average of $1 million in 2018-19.

Figure 1.5 District Revenue Comparison by Participation in College Readiness Block Grant

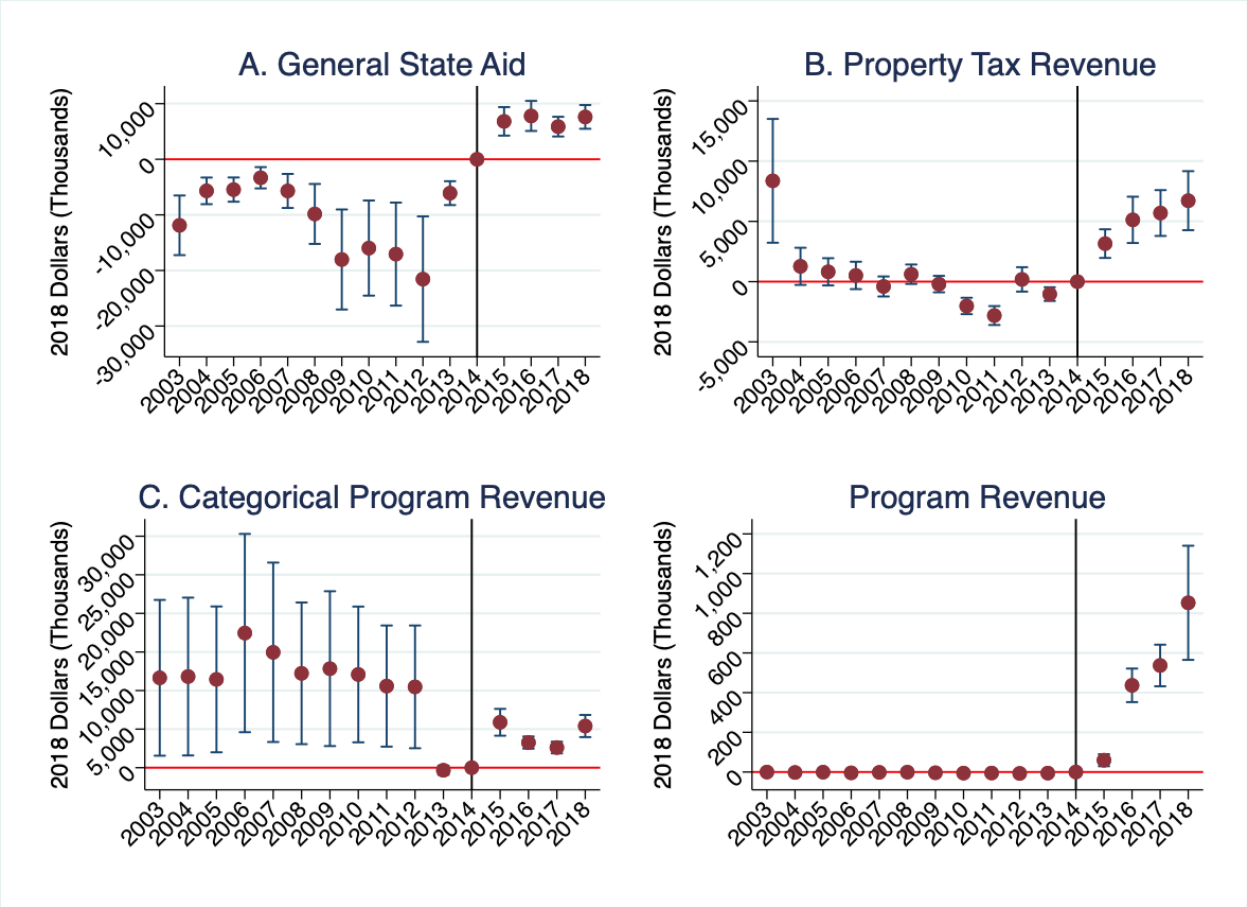

One of the primary uses of the funding in the College Readiness Block Grant was instructional materials. After the program was introduced in 2015, spending on instructional materials increases dramatically for participating districts as shown in Panel C of Figure 6. Prior to the program, these districts spent an average of $500,000 less, but they almost immediately began spending almost $1 million more once the grant was introduced. Differences in instructional

materials spending from other programs, however, were flat over several years and remained flat after 2015.

Figure 1.6 District Spending Comparison by Participation in College Readiness Block Grant

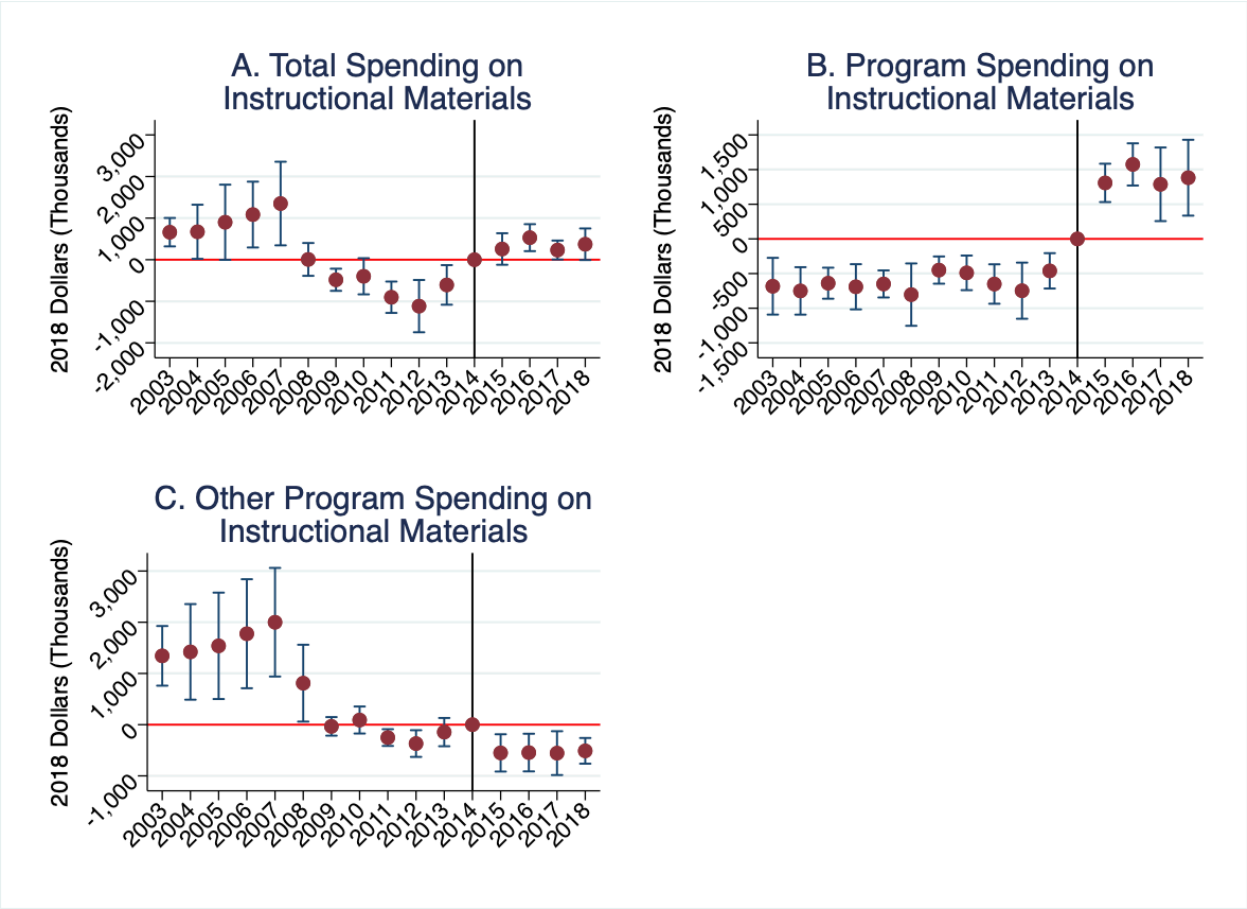

1.4.3 Economic Impact Aid

The Economic Impact Aid program is, like the Class Size Reduction Program, a longstanding grant that began before the start of my data in 2003. It was subject to major reforms in 2010, though, with its extension to juvenile court schools and its merging with a related program for English language development. The program was formally eliminated with the introduction of the LCFF, but this earlier expansion constitutes a sufficiently major change to stand in as the end of the original program. Figure 7 shows that participating districts tended to have lower revenues from general state aid, higher revenues from categorical programs and similar amounts of property

tax revenue. Despite the changes to the program in 2010, the revenues it generated continued going to participating districts until its formal elimination with the LCFF in 2013-14.

Figure 1.7 District Revenue Comparison by Participation in Economic Impact Aid Program

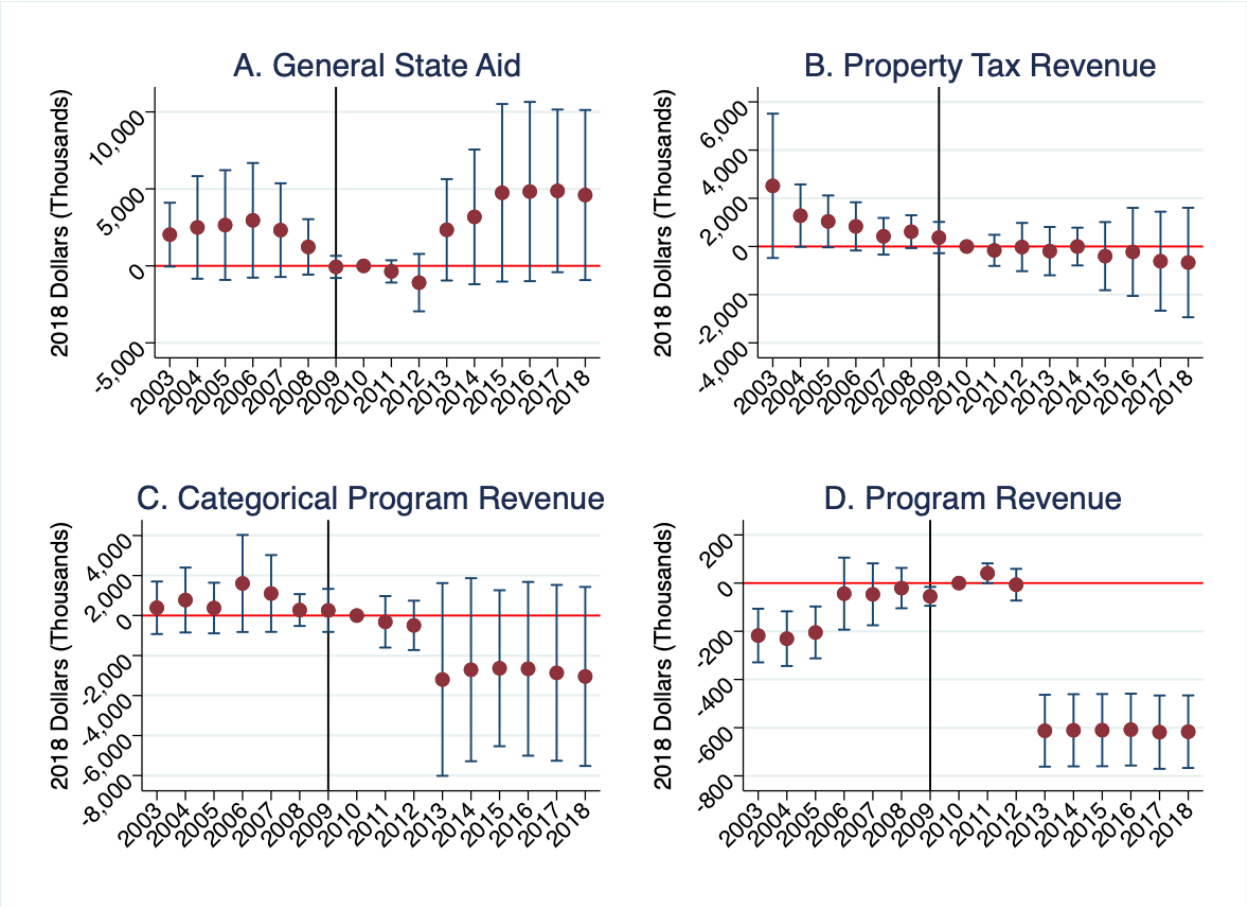

The spending results are shown in Figure 8. Teacher's aides emerged as one of the main uses of the revenues from the program. The graphs show that spending differences in this area were changed dramatically after the major changes in the program. In 2011-12, the first year with the loosened program requirements, spending on teachers' aides declined from an average difference of $50,000 to roughly zero. This, however, is the context of a broader move away from spending on teachers' aides in participating districts. Total spending in the activity also follows a downward trend, though the declines begin earlier in 2007-08 and are gradual unlike the discontinuous change seen for teachers' aide spending out of the program.

Figure 1.8 District Spending Comparison by Participation in Economic Impact Aid Program

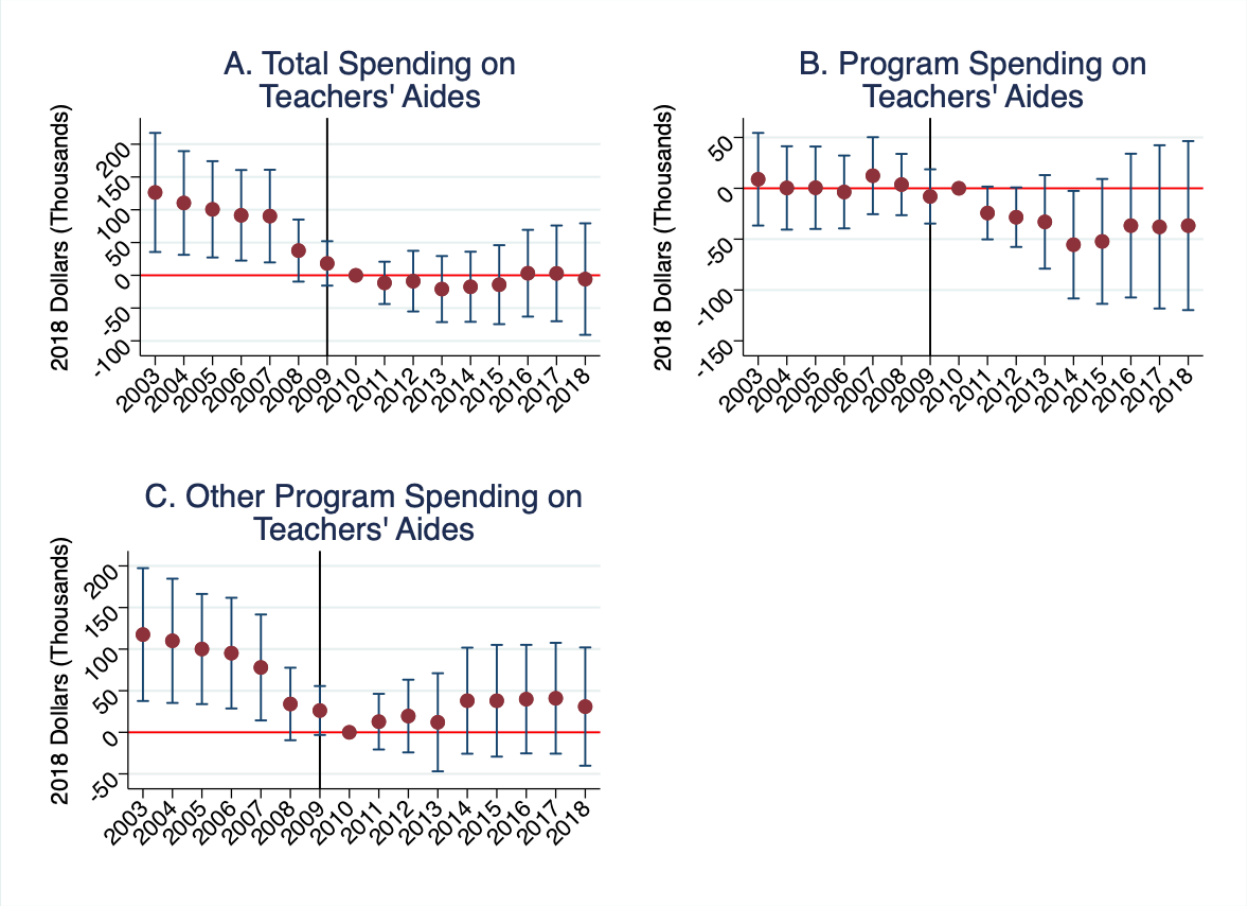

1.4.4 Teacher Credentialing Block Grant

The Teacher Credentialing Block Grant provides two changes to focus on: both the introduction of the program in 2005 and its elimination through categorical flexibility four years later in 2009. This program is also narrowly targeted at a particular use, namely mentorship support for beginning teachers. Starting with differences in revenues, Figure 9 shows that there were minimal differences between participating and non-participating districts in state aid, property taxes or total categorical program revenues. The program, though, generated about $300,000 per year for participating districts.

Figure 1.9 District Revenue Comparison by Participation in Teacher Credentialing Block Grant

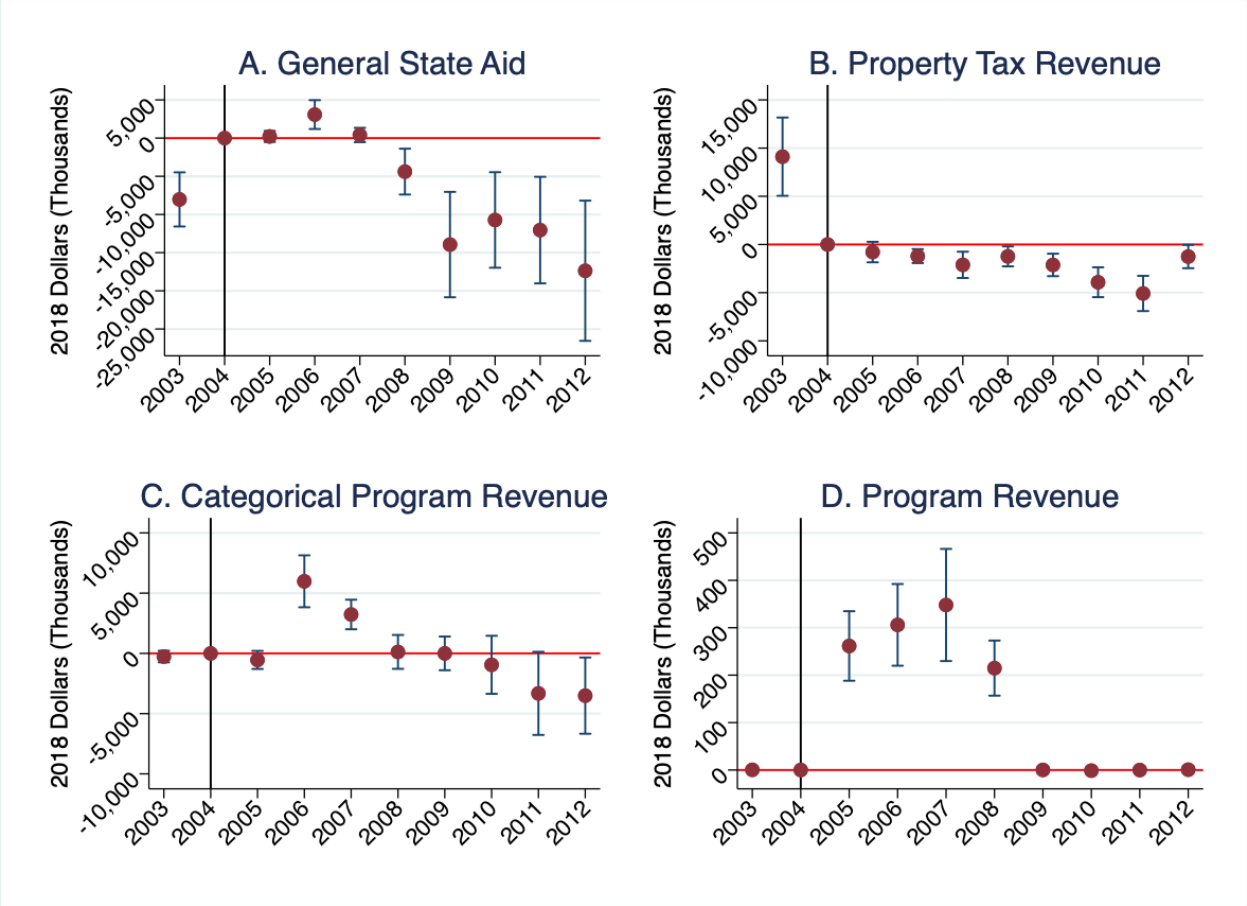

The differential in spending on mentorship for participating districts grows immediately in the first year of the program and continues throughout the years it was available as shown in Figure 10. After the elimination of the program, that spending differential falls back to zero. Spending differences from other categorical programs are minimal.

Figure 1.10 District Spending Comparison by Participation in Teacher Credentialing Block Grant

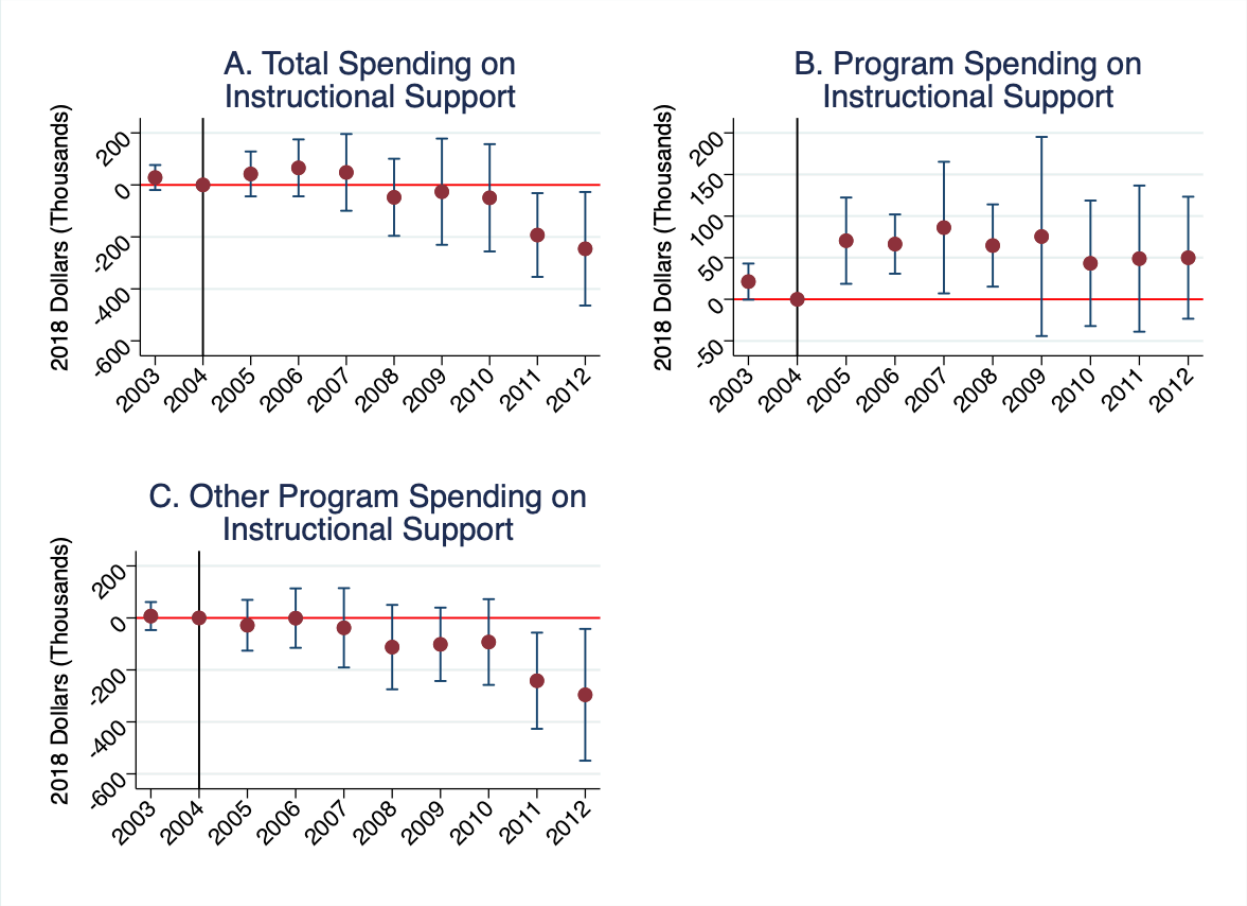

In the four cases above, a pattern begins to emerge. Each example shows a strong effect of the introduction or elimination of a program on the difference in spending on the main activities of the program between participating and non-participating districts. Program introductions are accompanied by spending differences that shoot up almost instantly, while the elimination of a program pushes that difference to zero just as quickly. Are districts policies shaped by the requirements in available categorical programs? Or, do districts instead develop their own priorities and then pursue only programs that serve those priorities? These results suggest an answer to that question. On the whole, districts do not appear to exert much control in their relationship with categorical programs. Instead, they take funds from these programs when they are available and spend them as intended so long as they are receiving revenues from it. Once the

program is removed, their interest or ability to funds those goals seem to go as well. Similarly, districts that ultimately choose to participate in a program demonstrate little investment in a program's goals until the program generates funds for them.

These results may point to the extent that districts are constrained in their annual budgets. Despite an interest in, say, heavily investing in mentorship for beginning teachers, districts may not be able to actually spend more in that area due to priorities elsewhere. It takes the introduction of a targeted program–which interested districts happily sign on to–in order to make the space in a district's budget for this niche activity and others like it. The restrictions themselves, not the additional revenue provided, appear to be the key features that define how districts experience categorical programs. The restrictions artificially elevate the priority of particular activities in the budget negotiation process. And, as the results show, this is the only way that districts that are interested in a particular set of investments can pursue them.

1.5 Student Outcomes

The results above showing the significant impact that categorical programs can have on how districts spend their revenues lead to an important subsequent question. Given that these targeted programs can so effectively steer the spending decisions of participating districts, are they being pushed in a positive direction? Students in a participating district are subject to a substantially different educational program compared to students in other districts. Does their experience of above average spending in these various areas lead to better outcomes? The spending changes induced by these programs could lead to more student success, or they could simply be ineffective at producing differential results—the connection between school spending levels and student achievement is a broad, contested literature of its own (Jackson 2016, Hanushek 1997). Worse still, the shift in spending toward particular activities brought on by these programs could

actually move districts away from a better mix of investments chosen by non-participating districts with less revenue and more flexibility. Below, I consider the impact of participation in the programs explored above on a variety of student outcomes.

Each of the programs explored here are broadly geared toward improving the quality of education that the students receive, whether that is through smaller class sizes, added help for disadvantaged students, better trained new teachers or a stronger emphasis on the pathway to college. Thus, looking at student achievement as measured by performance on the state's standardized tests offers a good general measure of the extent to which the programs are creating better outcomes.

I measure district-level student achievement by taking test scores standardized at the year-grade-test subject level and calculating weighted annual averages for each district based on the number of students for each grade and test. Naturally, the overall means of this measure are near zero, but this standardization allows a consistent comparison across several dimensions of variation. Students are in a number of different grades and take tests in a variety of subjects, each with different levels of average performance. In addition to that, districts serve different sets of students so that the grade span of students in one won't always match another. Finally, the state also shifted to a new set of tests with the transition to the LCFF, moving from the California Standards Test (CST) to the California Assessment of Student Performance and Progress (CAASPP), also known as the SBAC.

The graphs below show a comparison of district level achievement between districts that did and did not participate in the programs I have analyzed. The results are somewhat mixed. There seems to be little difference in achievement related to program participation for both the college readiness block grant as well as the economic impact aid program. On the other hand, student

achievement does seem clearly related to the timing of changes in the availability of both the class size reduction program and the teacher credentialing block grant.

Figure 1.11 Comparison of District Average Student Achievement by Program Participation

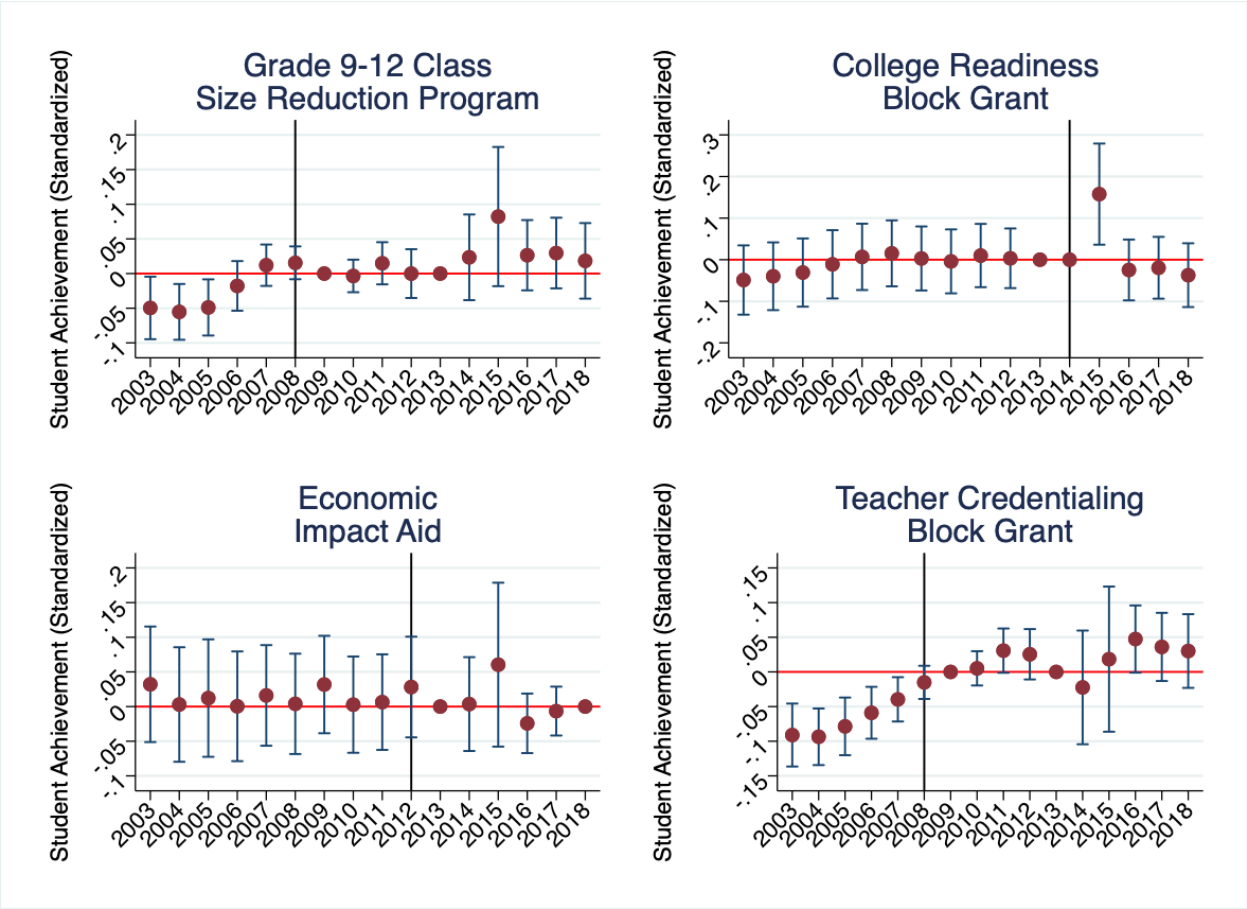

Prior to the introduction of the class size reduction program, average scores are consistently about .05 standard deviations lower in participating districts from 2003 to 2005, but scores improve quickly starting in 2006 and becoming level with non-participants by 2007. That pace of improvement is stalled though and does not continue after the program was eliminated in 2009.

The teacher credentialing block grant shows a similar pattern. Student performance in participating districts lags behind others before the program is introduced, but immediately begins catching up year after year while the program is in place. After the program is eliminated however, that pace of progress quickly levels off.

The results here point to some potential mechanisms for how categorical programs are related to student achievement. It may be that the additional money generated by the programs funded additional investments in specific areas that produced these positive results. Alternatively, program participation may be a broader indication of differences in district capacity. It could be the case that districts that are better equipped to apply for funding from the special programs and quickly shift their spending priorities in order to meet program requirements are also able to produce better results for students for reasons unrelated to the extra funding itself.

1.6 Conclusion

Categorical funding has long been one of the most popular tools available to California school districts looking to increase their revenues. Though this additional funding comes with restrictions on how the money must be spent, many districts choose to participate in these programs. My observation of several programs here showed that district spending policy was strongly influenced by participation in categorical programs as they shifted their spending toward activities related to program goals. Despite the influence that categorical funding could have on district spending, I also find that the shifts caused by these programs do not produce durable changes. In every case examined, spending differences between participating districts and those that did not participate in a given program only persisted for as long as the program was available. Additionally, in the cases where I examine spending differences between districts before a program is introduced, there is little difference between the districts that ultimately choose to participate and those that do not.

These results could point to a general lack of sincere interest in the goals of a given program on the part of the districts that participate in them. Districts may be happy to take on a program and even spend the added funds as dictated–perhaps shifting around funds elsewhere in their

budget according to their own prerogatives–as long as that extra revenue is on offer. However, these results could also point to the great extent to which district budgets are constrained. It could be the case that the vast majority of district budgets are committed to particular activities before the budget process even begins. Categorical programs, then, provide a means through which other activities that may not have a high priority compared to large and largely inflexible expenditures, like teacher salaries for example, can be funded when they would not be otherwise. These programs then open up space for a more varied educational program in both a district's coffers as well as its budget priorities.

I find that districts respond accordingly to the incentives available in these categorical programs and, in some cases, that their students are better for it. Average achievement rises for students in the districts participating in some of the programs I analyze. The timing of these gains also point to this effect being related to the added funding and the shift in educational investments brought about by these programs.

Although categorical programs provided added revenues for participating districts and, in some cases, led to district policy changes that produced better outcomes for students, they were also the subject of considerable criticism. The myriad programs available further complicated an already complex school funding system in California and there was wide concern that this stream of funding outside of the state's equalization formulas could reproduce the inequities that those formulas worked to eliminate. At first with the pressure of the economic downturn in 2008 and then again with the shift to the LCFF in 2013, categorical funding was nearly completely removed from the state's school finance system. However, much like how the previous system wasn't built up overnight, categorical funding has begun to reemerge. The years following the transition to the LCFF have seen the sporadic introduction of special programs that could, in time, add up to a

system similar to what existed before the Categorical Flexibility Act and the LCFF. As that system continues to grow, the results shown above highlight some of the implications of incentives these programs create for school districts, how they may respond, and what that could mean for the students they serve.

References

Hanushek, E. A. (1997). Assessing the effects of school resources on student performance: An update. Educational evaluation and policy analysis, 19(2), 141-164.

Jackson, C. K., Johnson, R. C., & Persico, C. (2016). The effects of school spending on educational and economic outcomes: Evidence from school finance reforms. Quarterly Journal of Economics, 131(1), 157-218.

Smith, J., Gasparian, H., Perry, N., & Capinpin, F. (2013). Categorical Funds: The Intersection of School Finance and Governance. Center for American Progress.

Sonstelie, J., Brunner, E., & Ardon, K. (2000). For Better or For Worse? School Finance Reform in California, Public Policy Institute of California. In University of California.

Timar, T. B. (1994). Politics, policy, and categorical aid: New inequities in California school finance. Educational Evaluation and Policy Analysis, 16(2), 143-160.

Timar, T. B., & Roza, M. (2010). “A false dilemma”: Should decisions about education resource use be made at the state or local level?. American Journal of Education, 116(3), 397-422.

Weston, M. (2011a). California's New School Funding Flexibility. Public Policy Institute of California.

Weston, M. (2011b). Rethinking the State-Local Relationship: K-12 Education. Public Policy Institute of California.

Weston, M., Cook, K., Murphy, P., & Ugo, I. (2015). Voluntary contributions to California’s public schools. Public Policy Institute of California.

Weston, M., Sonstelie, J., & Rose, H. (2009). California school finance revenue manual. Public Policy Instit. of CA.

Chapter 2: The Determinants and Implications of Savings Behavior

2.1 Introduction

School districts in California receive billions of dollars in revenue each year to provide educational services for their nearly 6 million students. In a typical year, districts save a portion of that revenue for future use. These savings help smooth cash flows and buffer lower levels of revenue in economic downturns. Although providing stable educational programs over time is a prudent fiscal goal and standard accounting practices provide guidance on specific levels of savings, excess savings present two potentially serious issues. First, excess savings by governments can be considered a failure to provide services to taxpayers who are currently paying for them. Second, excess savings may be related to agency issues. The term "agency issues" is a reference to the economic concept of the principle-agent problem, which describes the difficulties that can arise when the primary stakeholder in an organization and the day-to-day manager have different objectives. Limited oversight provides space for managers to act toward to their own purposes—potentially even against the mission of the organization as a whole. In the corporate world, such problems can manifest as self-interested managers engaging in value decreasing mergers (Opler 1999). In non-profits, Core et al. (2006) establish that larger endowments are related to higher manager salaries and lower program expenditures. Studying municipal governments, Gore (2009) demonstrates that excess reserves are related to higher administrator salaries and expenses. Although the robust literature examining cash holdings in the corporate environment has been extended to municipal governments like cities and counties, little research has explored these issues in the context of school districts. This paper aims to fill that void.

The research on reserves started with corporations. Original questions revolved around whether slack (various measures of fund balance) was beneficial in that it could even out revenue fluctuations and lead to positive outcomes for the firm, or whether it was harmful, because they

led to agency issues in which firm managers benefited at the expense of company performance. Opler (1999) finds that companies with larger cash reserves tended to engage more often in value-decreasing mergers—an indication of agency issues in this context. Tan and Peng (2003) add to this work by showing how fiscal slack—a slightly broader conception of reserves—can help during extreme downturns. Throughout this literature, reference is made to optimal level of savings, but the reality is that little is understood about what savings should be. Some accounting rules provide guidance, but even in the world of corporate finance it is difficult to determine what is optimal.

Despite the broad literature on corporations, little research extends this to the savings behaviors of government entities. Gore (2009) and Hand, Pierson and Thompson (2016) consider local governments such as cities and counties. Here, we replicate those papers with our contribution being the novel extension to the school district context. Our paper documents the level of reserves across California school districts from 2003-04 to 2021-22 and examines the drivers of these reserves based on the standard factors used in prior literature (Gore 2009 and Opler 1999). These factors include revenue volatility, district size and district growth, among others. We find similar results—with district size, and revenue volatility standing out as key determinants of district reserves. The extension of this type of analysis to school districts is important because school districts are unique even among local governments due to their interlocking relationships between the state and the federal government, students, parents and teachers. There is also relatively little prior research on school district saving and many questions remain to be answered.

After establishing the savings norms in the context of school districts, we extend the analysis to consider the potential for agency problems. Specifically, we replicate Gore's (2009) approach and measure the extent to which districts that hold high levels of reserves relative to their peers also have high spending on administration and higher administrative salaries. We find some

evidence that they do. Gore's interpretation of this positive association is that it is evidence of agency concerns. Hand, Pierson, and Thompson (2009), however, suggest that the positive association between higher than predicted reserves and higher administrator benefits may actually be evidence of better managerial competence. After establishing these results in the school district context, we make another key contribution in this literature, attempting to disentangle these two hypotheses. We do so by investigating whether districts with high relative savings see better productivity—here, in the form of higher student achievement. We find that higher than predicted reserves are partially related to better performance, running counter to concerns about agency issues.

The rest of the paper proceeds as follows. Section 2.2 describes the factors we expect to determine the level of school district reserves and model we use to estimate reserves. Section 2.3 describes the data sources we draw from for our study and the regression methods we use. Section 2.4 gives the results from these regressions and discusses their implications for variations in reserves holdings across California school districts. In Section 2.5, we shift to a consideration of districts that hold excess reserves and examine how district savings may relate to agency issues. In Section 2.6, we conclude.

2.2. Potential Drivers of School District Reserves

We construct a model of school district reserves drawing from the previous literature (Gore 2009), with some adjustments to account for the institutional differences between municipalities and school districts. Districts receive revenues from a variety of sources annually and spend them throughout the year. The remaining difference at the end of each fiscal year is held as reserves, adding to or decreasing the total amount of reserves from the year prior. In this paper, we focus on the unrestricted portion of a district's reported year-end balance in the General fund and a limited

number of other funds[5] that house monies that are similarly for general purposes.[6] Unrestricted reserves—fund balances that are available for any educational purpose, as opposed to restricted reserves[7] that are designated for particular uses—make up about 85 percent of total reserves.

We measure reserves as a school district's unrestricted ending balance as a share of its current annual spending. This scaling provides a more comparable measure across districts and doesn't depend on district enrollment. Formulated this way, we can describe a given level of reserves as the amount of funds that would allow districts to operate for some amount of time without new revenues. Reserves are typically about 25 percent—that is, the average district could operate for about a quarter of a school year on reserves alone.

School districts may hold reserves for several different reasons. Major considerations listed in a 2015 report from the Legislative Analyst's Office include managing cash flow, covering unexpected expenses, mitigating revenue volatility, saving for major purchases and improving the district's credit rating. Given those rationales, a number of factors emerge that we would expect to relate to larger reserves. We consider each in turn. Summary statistics for these important district fiscal variables are shown in Table 1 below.

Table 2.1 Summary statistics of determinants of reserves

Variable	Mean	Standard Deviation
Enrollment (ADA)	8,238	24,591
Average Change in Enrollment	-0.82%	6.96%
Growth	3.37%	4.29%

[5] Ending balances are taken in total from the General fund a limited number of other funds. These funds include the Cafeteria Special Revenue fund, the Deferred Maintenance fund, the Special Reserve fund, the Special Reserve for Postemployment Benefits fund, the Cafeteria Enterprise fund, the Pupil Transportation Equipment fund and the School Bus Emissions fund. Notably, we exclude Capital funds along with other funds that are tied to particular uses outside of typical classroom activities.

[6] The General fund makes up 80.3% of unrestricted reserves, while the Special Reserve fund accounts for 8% of reserves and the Deferred Maintenance makes up 6.1% of reserves. The remaining funds make up less than 6 percent of reserves together.

[7] Restricted reserves include money that is set aside, but not available to spend due to legal or temporal issues. For example, funds to pay staff salaries between state apportionments are accounted for as reserves, but they are not available to be spent elsewhere given a district's contracted obligations to pay employees.

Decline	-4.19%	-3.86%
District Type		
Elementary	45.41%	(19.3% of enrollment)
High	10.40%	(9.2% of enrollment)
Unified	44.19%	(71.5% of enrollment)
Revenue Volatility	7.76%	2.09%
Reserves	24.49%	11.94%
Restricted	3.35%	1.79%
Unrestricted	21.14%	11.72%
Revenue Sources		
General State Aid	39.99%	17.80%
Federal Revenue	9.61%	4.64%
Property Taxes	26.85%	19.04%
State Restricted Revenues	13.57%	3.35%
Other Local Revenues	5.86%	5.17%
Other Revenues	4.12%	3.15%
Number of Districts	654	

Note: Revenue volatility is measured by the coefficient of variation. This is the ratio of the standard deviation of revenues to average revenues. School district reserves are measured as the ratio of reserves to total spending in the same here. The resulting percentage gives the amount of years a school district could operate on reserves alone.

2.2.1 District size

Larger school districts potentially benefit from economies of scale and are therefore predicted to need fewer assets in reserve. The vast of majority of district revenues are generated by formulas tied directly to their student population, so larger districts tend to have greater revenues. District type also plays a role because after 2014 state school funding varies according to grade span, with high school students typically generating more revenues. About 44 percent of districts in our sample are unified, serving students in kindergarten through grade 12. However, they account for 72 percent of enrollment. Elementary districts and high school districts have both typically generated more revenue on a per-pupil basis. The connection between enrollments and expenditures is clear—additional students require additional educational resources, though fixed costs are spread more thinly in districts with larger enrollments. With greater total revenues, larger districts are more easily able to manage things like cash flow issues and unexpected costs through ordinary budget slack. Further, the sizes of the budget shortfalls that may arise due to these issues are much smaller as a share of the overall budget for larger districts than for small ones. Thus, the same reserves (as a percentage of revenues) are able to cover unforeseen expenses more easily in a larger district. Prior literature has found a negative relationship between size and reserves (Opler et al. 2009; Core et al. 2006; Gore 2009).

We use district average daily attendance (ADA) as our measure of enrollment. This measure counts the number of students that are present on a typical school day and is slightly lower than the enrollment data collected by school districts which measures the number of students enrolled in the school on the annual census day in October. Also, ADA is more directly tied to district funding as it is the basis on which per-student allocations are distributed. District sizes vary widely. The 220 largest districts in the state (about 1 in 4) enroll 4.1 million students, making up

80% of the state's total enrollment. On the other hand, there are 242 small districts with enrollments lower than 300 students. These make up under 1 percent of total state enrollment and we drop these smallest districts in our analysis. In the regression specification below, we take the log of ADA in order to account for non-linearities in district size.

2.2.2 Growth

Although larger districts may need to hold fewer reserves, prior literature suggests that districts with growing enrollment may opt to hold more reserves in anticipation of needing to expand educational programs. However, in California, student growth is met by additional revenue so there is less need to save. Alternatively, declining enrollment presents a major source of pressure on school district budgets. As enrollment declines, so does revenue, and fixed costs per student increase. Although general state aid revenues are often adjusted to temporarily hold districts harmless for falling revenue entitlements, this reprieve is only temporary. In the context of California school districts, we would predict a potentially nonlinear relationship between growth and reserves, where negative growth has a positive impact on reserves, and positive growth may not affect reserves at all.

A further complicating issue is that while the revenue increases associated with rising enrollments are somewhat constant, the expenditure requirements can be discontinuous—building a new school, for example. Considering the reverse, declining enrollment also presents a major challenge for school districts.

We measure enrollment growth by the percentage change in ADA between the current year and four years prior. Yearly swings in enrollment can be quite volatile, so long-term differences can better reflect durable changes in district student populations. In the regressions below, changes in enrollment are measured with two variables: one for

enrollment growth, and another for enrollment declines. The growth variable is coded as zero for districts with declining enrollment and vice versa for those that are growing. Separating the two variables allows for an estimate of the differential effects of rising and falling enrollments described above. Gore (2009) similarly measures population changes over a five-year gap, though no adjustment is made to allow for non-linear effects of growth relative to decline.

2.2.3 Revenue Volatility

Although there are a wide variety of considerations that may drive whether a district chooses to hold reserves, mitigating inconsistencies in regular revenues is perhaps the primary concern. We expect that districts that are typically exposed to more volatile revenues will hold larger amounts in reserve. With much of district budgets tied up in regular expenses like teacher salaries and routine maintenance of school facilities, interruptions in normal revenue generation create an immediate challenge. Reserves provide a valuable means to either completely address those challenges in the short term or dampen the pace of necessary fiscal changes that would otherwise be rapid and disruptive. To the extent that some districts are subject to more revenue volatility than others, this creates a need to hold larger reserves in order to respond to revenue shortfalls when they occur. Additionally, when volatile revenues produce budget surpluses, districts may choose to hold larger reserves in anticipation of a decline in revenues in future years.

Following Gore (2009), we use the coefficient of variation of revenues in the last five years. This measure takes the standard deviation of revenues in that time span and divides it by the mean value over that same time period in order to account for differences in scale across districts. Gore (2009) finds this measure of revenue volatility is positively related to the reserves of municipal governments. This scaling ensures that revenue volatility is measured in a way that is not sensitive to total district revenues.

2.2.4 Limited Revenue Sources

The composition of a district's revenues can be a key source of volatility and thus a determinant of its reserves. Districts with limited revenue sources—particularly ones that rely on relatively volatile sources—need to hold higher amounts in reserves as a hedge against potential shortfalls. School district revenues in California are largely derived from a combination of property taxes and general state aid, with much of the remainder coming from restricted state categorical program funds. Although districts exercise minimal control over the rates, property tax revenues are relatively stable compared to general state aid, which is largely funded by the much more volatile state income tax. State policy mostly equalizes per-pupil revenues across districts, though recent reforms provide additional funding to districts based on the number of disadvantaged students served.[8] The composition of revenues, however, varies according to the value of districts' local property tax bases. Some districts generate most of their revenues from property taxes, while others rely heavily on state aid to reach the amount set by state policy. Those districts where state aid comprises a larger portion of revenues are potentially exposed to more volatility, particularly in the case of potential recessions when furloughs, deferrals and other kinds of interruptions to funds from the state may be used. In our model, we control for the share of district revenues coming from general state aid, local property taxes, federal funds, categorical programs and other sources to capture the extent to which different districts rely on different kinds of funding—each with their own typical levels of volatility. Highlighting the differences between the school district setting and the municipal one, Gore

[8] The Local Control Funding Formula, which was implemented in the 2013-14 school year provides additional funding to school districts based on the percentage of students that are low-income or English Learners. These supplemental grants are further bolstered by concentration grants that provide even more funding to districts where more than 55 percent of students are in these targeted categories.

considers the shares of revenues from income, property and sales taxes finding that districts that rely more heavily on fewer sources of revenues tend to hold greater reserves.

2.2.5 Model of Reserves

The first phase of our analysis establishes a model of school district reserves that portrays how most districts operate. Given the large number of districts in California, and that they are each subject some level of public accountability, the typical savings levels of a district represent one measure of best practices. If not, these predicted reserves at least represent behavior that is broadly acceptable to voters in myriad school districts throughout the state. We model school district reserves using our 20-year panel of the state's hundreds of districts through a regression analysis that aims to mitigate confounding factors as much as possible, leaving only the relationship between the factors described above and the levels of reserves districts hold. Our empirical model is as follows,

$$\begin{aligned} Reserves_{\mathrm{it}} = {} & \beta_0 + \beta_1 Log(\mathrm{ADA})_{\mathrm{it}} + \beta_2 Growth_{\mathrm{it}} + \beta_3 Decline_{\mathrm{it}} + \beta_4 RevVolatility_{\mathrm{it}} \\ & + \beta_5 StateRev_{\mathrm{it}} + \beta_6 FedRev_{\mathrm{it}} + \beta_7 StateRestrictedRev_{\mathrm{it}} + \beta_8 LocalRev_{\mathrm{it}} \\ & + \beta_9 \mathrm{OtherRev}_{\mathrm{it}} + \beta_{10} \mathrm{DistType}_{\mathrm{i}} + \mathrm{YearFE} + \sigma_{\mathrm{it}} \end{aligned}$$

where per-pupil reserves in a given district, i, in year, t, are measured as function of a variety factors that we described in the previous section. Our regression specifications also include year fixed effects so that temporal statewide effects—an economic downturn, for example—do not confound our results. Standard errors are clustered at the district level to account for some of the correlation in district characteristics across time.

2.3 Data

This project relies on a variety of measures that are found in publicly available data sets from the California Department of Education (CDE). The data span 2003-2021, covering 896

districts and, on average, 5.4 million students annually. The CDE provides detailed financial information on California school districts through the Standardized Account Code Structure (SACS) data. This data set describes expenditures, revenues and reserves with added detail on the types and sources of each. Total district revenues in the state average nearly 68 billion dollars, with about 67 billion dollars of spending. Total state reserves average 12.4 billion dollars—though with significant fluctuations over time. Figure 1 shows trends of aggregate revenue, expenditures, and reserves over time. Per-pupil revenues and expenditures are relatively evenly distributed. Average per-pupil revenues in 2020 dollars range from $10,500 to $15,800 over the years our data cover, and typically 50% of districts have revenues within $2,000 dollars of the state average. The amounts of reserves that districts hold varies considerably. On average, they are 25 percent of expenditures with a standard deviation of 12 percent.

Figure 2.1 Total State Revenues, Expenditures and Reserves 2003-2021

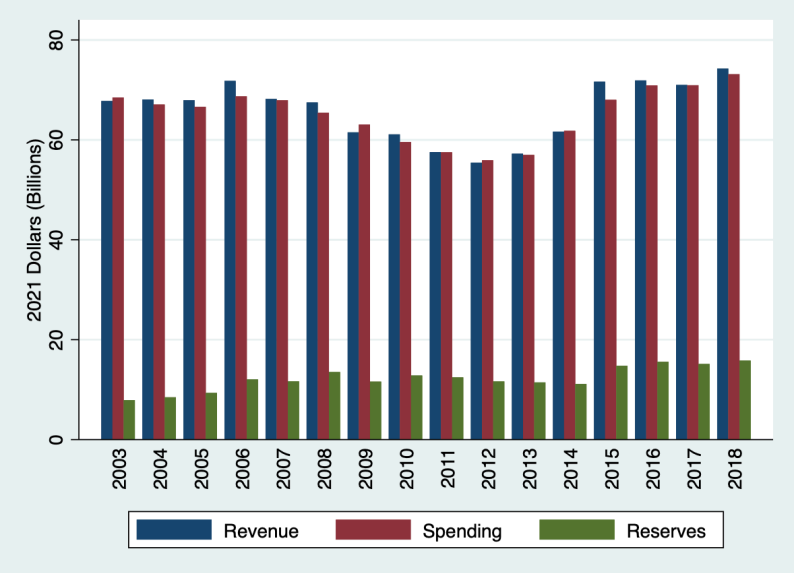

In the second part of this project, where we examine the implications of district savings behavior, we use data on the characteristics of district staff like the number of

administrators and salary levels. Data on administrator salaries comes from the Certificated Salaries and Benefits Schedule for the Certificated Bargaining Unit (Form J-90) data available from the CDE. These data contain the superintendent salaries of each school district in the state. Student achievement results from the state's standardized tests[9] are also used. These data are augmented with data on student characteristics like the share of students eligible for free or reduced-price meals (a common proxy[10] for family income) and the distribution of students across race.

2.4 Regression Estimates of Determinants of Reserves

We report the results from a regression of school district reserves on the determinants described above in Table 2. Generally, the reported coefficients are similar in direction to the results in Gore (2009).

[9] From 2003 to 2013, the state standards tests used was the California Standards Test (CST). The tests were replaced by the Smarter Balanced tests, often referred to as SBAC. The new tests are aligned with the Common Core State Standards and were a key part of the major reforms to the state's education system in 2014.

[10] Although participation in the free and reduced-price meals program is often used as a proxy for whether a student comes from a low-income family, it is not a perfect measure. Program eligibility does not always match program take-up, and low-income students are not counted if they don't sign up for the program. Further, the binary nature of the low-income status indicator can miss important gradients of family income both below and above the eligibility cutoff.

Table 2.2 Determinants of Reserves Regression Results

VARIABLES	Reserves
Log ADA	-0.0412***
	(0.00386)
Growth	-0.0170
	(0.0361)
Decline	-0.209***
	(0.0538)
Elementary District	0.0220**
	(0.00868)
High School District	0.0221*
	(0.0130)
Revenue Volatility	0.258**
	(0.121)
General State Revenue	-0.00878
	(0.0296)
Federal Revenue	0.102
	(0.108)
Other State Revenue	-0.165
	(0.172)
Local Revenue	-0.111
	(0.109)
Constant	0.552***
	(0.0515)
Observations	10,434
R-squared	0.221

Clustered standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

First, we find that district size is negatively correlated with reserves. On average, a 10% difference in enrollment across districts is associated with reserves holdings that are 0.4 percentage points lower. These results are consistent with the prior literature. District fixed costs and economies of scale appear to be a key driver of reserve holding behavior.

The results are more complicated to interpret when we measure the relationship between changes in district enrollments over time and reserves. We find that positive district growth is not significantly related to reserve holdings, whereas the significant -0.209 coefficient on declining enrollment implies that a 5 percent decrease in enrollment over the previous five years (the median change for districts that saw declines) is associated with a 1 percentage point decrease in reserves as a share of spending. This result suggests district savings are a luxury. Districts with declining enrollments anticipate lost revenues and prioritize current spending over saving for future needs. Additionally, even with declining enrollment, many fixed costs remain for districts. With a greater share of resources committed to those costs, there is less room for accumulating reserves.

Revenue volatility is strongly related to district reserves; districts with more unstable revenues tend to hold larger reserves. This result matches our expectation and the findings for reserves holding in other contexts. To the extent that incomes or revenues are uncertain, reserves are a popular means of transferring funds from good economic times over to when there are revenue shortfalls. The result here is strongly positive with a coefficient of 0.258 and a t-statistic of 2.1. Looking at the distribution of revenue volatility, districts in the 25th percentile have a coefficient of variation of 0.045—that is, the standard deviation of revenues in recent years is about 4.5 percent of revenue on average. The coefficient of variation at the 75th percentile is 0.1. Given the coefficient, moving from the first to third quartile implies reserves would typically be 2.6 percentage points higher in the districts with more volatility.

Finally, differences in sources of revenues across districts reveal some interesting correlations, but none that are statistically significant. Districts where general state aid makes up a larger portion of revenues appear to have smaller reserves, but the coefficient is small and not statistically significant. The opposite appears to be true for districts that receive a higher share of revenue from federal sources. Coefficients are negative, but insignificant for districts with greater shares of revenue from restricted state sources—for example categorical programs like special education and others—and other local sources (i.e. parcel taxes). The correlations are suggestive of the mechanical requirement to spend rather than save categorical funds and the relatively lower volatility in revenues from special local taxes like parcel taxes.

Altogether, the results here accord with those in Gore (2009) and Core (2006). Districts appear to save in a way that matches best practices. There is, however, considerable variation in individual district behavior as evidenced by the variation in residuals from our regression. Figure 2 below shows the distribution of residuals from our model of reserves. These residuals measure the excess reserves above (or below, if negative) what the model predicts. Although the median district holds reserves of 35.6 percent of expenditures, with a -0.02 residual, we find that districts at the 75th percentile hold reserves that are 5.4 percentage points higher than what our model would predict. At the other end of the distribution, we have that districts in the 25th percentile hold reserves that are 8.4 percentage points lower less than our model predicts. The higher than predicted reserves are what Gore has interpreted as leading to agency issues.

Figure 2.2 Distribution of Residuals

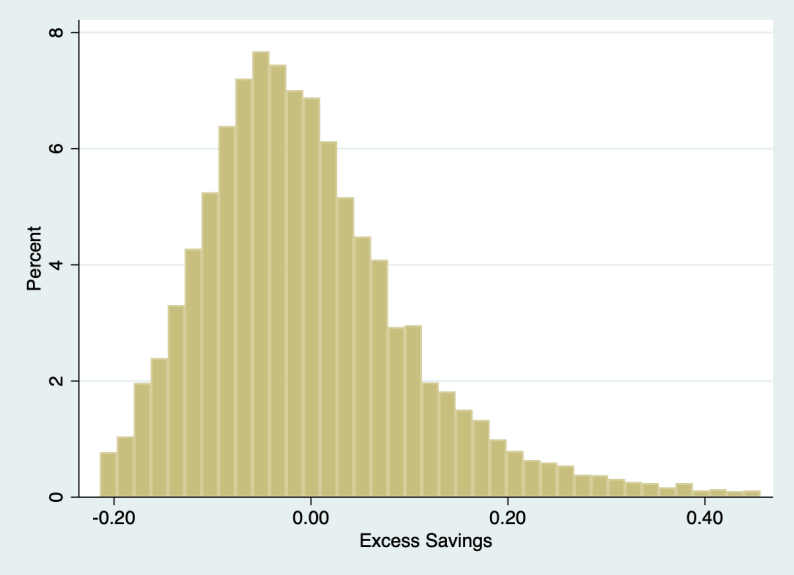

2.5 Agency Issues

As in Gore (2009), we consider agency issues to arise when administrators might spend money in a way that benefits them directly rather than the district. To test this, we consider two potential outcomes that may benefit administrators: the share of expenditures districts dedicated to administration and the salaries of superintendents in the district. We regress each of these outcomes on residuals from our first regression along with a variety of other covariates. Following Gore (2009), we lag the residuals in order to assess the effect of excess savings. The residual in the current year gives the remaining funds after the revenues and expenditures in that year. The one-year lagged residual, however, gives the amount of reserves that are available to spend in the current year. We consider two specifications with administrative spending as the dependent variable. The first controls for district type and district size. The second adds district demographics including the share of low-income students—measured by participation in the free or reduced price meals program—and the distribution of enrollment by race as well. In the latest year of our data, about 55 percent of students are low-income by our metric. A plurality of the state's students are

Hispanic (46 percent), while White (36%), Black (2.5%) and Asian (7.4%) students make up most of the rest. Both regressions also include year and region[11] fixed effects. The year fixed effects control for common statewide shocks, while the region fixed effects are meant to control for variation in regional economies that could be related to wage levels and administrative spending. Standard errors are clustered at the district level.

Table 2.3 Excess Reserves and Administrative Spending Regression Analysis

	(1)	(2)
VARIABLES	Administrative Spending	Administrative Spending
Excess Savings	144.1**	155.4**
	(63.9)	(62.4)
Log ADA	-99.08***	-100.6***
	(10.2)	(10.8)
Elementary District	-112.3***	-114.9***
	(18.8)	(17.7)
High School District	-19.95	-1.553
	(21.7)	(21.5)
Percent FRPM		233.5***
		(46.8)
Percent Black		419.4
		(296.1)
Percent Asian		-192.8**
		(76.5)
Percent Latino		-125.0**
		(53.0)

[11] School districts are sorted into regions based on Metropolitan Statistical Area, or MSA. See https://labormarketinfo.edd.ca.gov/definitions/metropolitan-statistical-areas.html for more information.

Percent Other Race		257.0*
		(133.7)
Constant	1,461***	1,366***
	(90.9)	(90.1)
Observations	9,759	9,759
R-squared	0.367	0.411

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

In both regression specifications, we find that districts with excess savings tend to spend more per student on administration. The 25th percentile of residuals is -0.079, while the residual is 0.054 in the 75th percentile. Comparing districts in the bottom quartile of excess savings to those in the top, the results in Model 1 imply that the districts in the higher quartile spend roughly $20 per pupil more on administration. For scale, that is about 3.1 percent of the average spent on administration per pupil. When we control further for district demographics, we find that the results are largely similar, with high-savings districts spending about $20 more per student on administration. Given the average ADA is 8,200 students, this extra spending translates into $164,000, enough to fund an additional position or two. Reviewing the other covariates, we see in both regressions that both elementary school districts and larger school districts tend to have lower administrative spending per pupil. Higher rates of student poverty are associated with more administrative spending, and the association with race and ethnicity are mixed. Greater shares of Black students are associated with higher administrative spending, while the share of Hispanic and Asian students correlates negatively with administrative spending.

Considering agency issues further, we also investigate the relationship between excess reserves and superintendent salaries. Table 4 below provides the results from a regression of imputed annual superintendent salaries on our residuals and a variety of covariates.[12] As with the regressions for administrative spending, the first regression controls for district type and size; the second adds district demographics. A third specification adds controls for the experience, gender and race of the district superintendent. All three specifications include year and region fixed effects along with standard errors clustered at the district. Across all three specifications we find that superintendent salaries are positively related to excess savings. Districts in the 75th percentile of excess savings have annual superintendent salaries that are $2,774 higher than districts in the 25th percentile. When we control for district demographics, that salary difference is slightly smaller. However, when we control for the demographics of the superintendent as well, the difference in salary is again just over $2,700. Notably, the regional fixed effects should capture average salary differences across labor markets, minimizing concerns that the residual may reflect labor market differences.

[12] Imputation is required because superintendents vary in the number of days they work each year, with some superintendents working half-time or splitting time as a principal in the case of some small districts. Using data from the J-90, which provides information on staff salaries for each of the state's districts, we multiply each superintendent's daily salary by the average number of days worked in the sample to get a measure of annual salary that is comparable across the entire sample.

Table 2.4 Excess Reserves and Administrative Salaries Regression Analysis

VARIABLES	(1) Superintendent Salaries	(2) Superintendent Salaries	(3) Superintendent Salaries
Excess Savings	20,904***	18,959***	20,627***
	(5,919)	(5,663)	(6,365)
Log ADA	21,057***	21,265***	21,170***
	(824.0)	(808.8)	(852.4)
Elementary District	-5,178***	-4,707***	-5,038***
	(1,644)	(1,620)	(1,698)
High School District	4,976**	4,286*	3,528
	(2,418)	(2,397)	(2,413)
Percent FRPM		-11,445**	-10,652**
		(4,713)	(5,002)
Percent Black		-75.91	3,069
		(14,423)	(14,875)
Percent Asian		12,132	12,899
		(9,769)	(10,170)
Percent Latino		806.8	-1,312
		(5,002)	(5,196)
Percent Other Race		7,436	2,708
		(6,019)	(6,858)
Total Experience			74.26
			(58.31)
Experience in Current District			-177.0***
			(54.43)
Female Superintendent			-2,416**

			(1,186)
White Superintendent			-167.3
			(1,434)
Constant	-19,086**	-14,048*	-10,390
	(7,400)	(7,419)	(7,933)
Observations	9,179	9,179	7,506
R-squared	0.773	0.776	0.779

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Elementary school districts tend to have lower superintendent salaries, whereas larger districts have higher pay. Student poverty is associated with lower salaries, and the student race effects are statistically insignificant. Controlling for both total experience and experience in the district, we see that experience in a particular district is negatively related to salary when total experience is held constant. Also, we see a negative relationship with salaries for superintendents who are women and those who are White.

Much like Gore (2009), we find results that are suggestive of agency issues with excess savings in the context of California school districts. Considering two measures of potentially self-interested behavior, the regressions described above find that high-saving districts—those with reserves that much higher than predicted by our model of the determinants of reserves—tend to spend more on administration in general and on superintendent salaries in particular. A variety of additional controls including year and region fixed effects, district size and type, student demographics and even the demographics of individual superintendents only affect this main result slightly.

The interpretation of similar results for municipalities throughout the U.S. in Gore (2009) was that the relationships between excess reserves and administrative spending implied the existence of agency issues—where managers are accumulating and distributing public savings for their own benefit. In response, Hand, Pierson and Thompson (2015) replicate the results, but offer a discussion of an alternative explanation: Perhaps excess reserves are improperly defined, and instead the level of reserves in high-saving municipalities—as well as the high administrative spending associated with them—reflect prudent fiscal management and a premium paid for high-quality administration.

2.5.1 Student Achievement

Given that test performance is likely a core goal of most districts, a positive relationship between excess reserves and student achievement could allay concerns about agency issues. If higher than expected reserves (and their related higher administrative spending and salaries) are able to provide a higher quality education (as measured by test scores), that would provide suggestive evidence that districts with higher achievement are striking the appropriate balance between using current resources for current students and maintaining adequate reserves.

Table 5 below reports the results of a set of regressions of student achievement[13] on excess reserves. We control for district type and size, as well as student demographics in our second specification. In a third regression, we further control for the square of the lagged residual. This allows a quadratic relationship between achievement and excess savings. The curvature in the regression line allows for estimation of the optimal level of reserves—or for evidence of increasing or diminishing returns to greater savings.

[13] Student achievement is measured by the share of students in a district achieving a score rated at or greater than Standard Met on the SBAC or Proficient on the CST. Both tests cover reading and math. The CST was administered in grades 2-11, while the SBAC is taken by students in grades 3-8 and 11.

Table 2.5 Excess Reserves and Student Achievement Regression Analysis

VARIABLES	(1) Student Achievement	(2) Student Achievement	(3) Student Achievement
Excess Savings	0.109***	0.0454***	0.0728***
	(0.0371)	(0.0170)	(0.0190)
(Excess Savings)^2			-0.0872**
			(0.0344)
Log ADA	-0.0118**	0.00520*	0.00435
	(0.00545)	(0.00265)	(0.00264)
Elementary District	0.0397***	0.0559***	0.0561***
	(0.0123)	(0.00529)	(0.00525)
High School District	-0.0496***	-0.0772***	-0.0769***
	(0.0142)	(0.00735)	(0.00730)
Percent FRPM		-0.364***	-0.363***
		(0.0185)	(0.0184)
Percent Black		-0.386***	-0.385***
		(0.0572)	(0.0565)
Percent Asian		0.125***	0.125***
		(0.0264)	(0.0263)
Percent Latino		-0.139***	-0.138***
		(0.0204)	(0.0202)
Percent Other Race		-0.0840***	-0.0810**
		(0.0310)	(0.0314)
Constant	0.442***	0.644***	0.651***
	(0.0516)	(0.0247)	(0.0246)
Observations	9,103	9,103	9,103

R-squared	0.392	0.825	0.825

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

In each specification, we find a positive relationship between student achievement and excess savings. When we control for district demographics, the size of the coefficient is smaller, but still positive and statistically significant. When we include the squared residual, we see that the linear coefficient is positive while the quadratic term is negative. This implies that there are positive but diminishing returns in student performance to greater levels of excess savings.

These results suggest that the typical district is actually saving less than is needed to produce peak performance. However, high-savings districts also tend to hold reserves that are larger than this optimum. Thus, while these districts do appear to be holding reserves that are excessive to some extent, their above-predicted level of savings is at least partially related to a higher level a performance that they may be paying a premium for in their administrative spending.

2.6. Conclusion

Our study extends the broad literature on the savings behavior of various private and public entities to the context of California school districts. Our findings accord with much of the prior literature. Size and revenue volatility are key determinants of district reserves. Following on work from Gore (2009) and others, we continue our analysis and find that districts with higher-than-average levels of savings tend to spend more on administration and offer higher superintendent salaries. While previous discussion of similar findings has linked this relationship to agency issues, we investigate this interpretation and find some evidence that the relationship may instead reflect a premium for management quality. Districts with higher savings than our models predict produce

better student performance than districts with more typical savings rates, implying that these districts strike a better balance between current spending and prudent foresight.

School funding in California, like most other public services in the state, is subject to considerable volatility given the state's reliance on the income tax. As the economic cycle turns, districts may need to rely on their savings to weather disruptions in normal revenues. This paper highlights the considerations that should factor into the level of reserves a district would ideally hold. While maintaining ongoing educational programs is imperative, this paper discusses the tradeoffs that factor into reducing current spending to support students in the future. However, it appears many California districts are not being prudent enough in their savings decisions.

References

Opler, T., Pinkowitz, L., Stulz, R., & Williamson, R. (1999). The determinants and implications of corporate cash holdings. Journal of financial economics, 52(1), 3-46.

Core, J. E., Guay, W. R., & Verdi, R. S. (2006). Agency problems of excess endowment holdings in not-for-profit firms. Journal of accounting and economics, 41(3), 307-333.

Gore, A. K. (2009). Why do cities hoard cash? Determinants and implications of municipal cash holdings. The Accounting Review, 84(1), 183-207.

Hand, M. L., Pierson, K., & Thompson, F. (2016). A Replication Study of "Why Do Cities Hoard Cash?"(The Accounting Review 2009). Public Finance Review, 44(5), 681-687.

Tan, J., & Peng, M. W. (2003). Organizational slack and firm performance during economic transitions: Two studies from an emerging economy. Strategic management journal, 24(13), 1249-1263.

Chapter 3: The Impact of Volatility in California School District Revenues

3.1. Introduction

School districts in California are largely unable to choose the level of their revenues. Most districts rely on state aid for a substantial portion of funding, and much of the remainder comes from property taxes collected at rates that—due to elements of the state's history and political structure—they are not able to set. Revenues in California are notably volatile relative to other states given its reliance on personal income and capital gains taxes, particularly those levied on high-income residents[14] (Hill 2005). School districts in the state must then constantly plan around an ever-changing level of revenue during the annual budget process, taking care to incorporate both general uncertainty and the threat of major shortfalls as was the case with the contractions after the "dot-com" bust in 2000-01 and financial crisis in 2007-09.

A key concern in the budget process is maintaining a consistent level of expenditures in order to reduce churn in the students' classroom environment and promote stability for school staff and personnel. We examine the extent to which districts are able to maintain expenditure stability over time—saving more when revenues rise in order to mitigate necessary spending reductions when revenues fall.

Using detailed district financial data from 1995-2022, we estimate the relationship between annual changes in revenues and expenditures. We develop a "revenue volatility index" that describes how well districts are able to smooth expenditures given the revenue changes in each year. We examine the distribution and determinants of this this index before turning to a further analysis that looks at the relationship between the index and two measures of district performance—standardized test scores and class sizes. We find that nearly 90 percent of annual

[14] Roughly 60 percent of general fund revenues in the state come from the income tax. About 45 percent of income tax revenue is collected from 1 percent of Californians.

changes in revenues are reflected in changes in spending. This rate varies substantially across districts—ranging from 25 percent at the low end to 100 percent or more at the top. The lower this percentage, the more a district smooths its expenditures from year to year. In further analysis of variation in expenditure smoothing across districts, we explore the relationship between the revenue volatility index and two outcomes, student achievement and class sizes. Both relationships are in the opposite direction of what we hypothesize, with districts that smooth less seeing better student achievement[15] and smaller class sizes. We explore these results in detail below and discuss potential reasons for the findings.

The rest of the paper proceeds as follows. We describe California's school finance system in greater detail in section 3.2. Then in section 3.3, we review the related literature. In section 3.4, we describe the methodology used in the analysis. Section 3.5 describes the data used and provides summary statistics. In section 3.6 we present the results. And in section 3.7, we discuss the findings and conclude.

3.2 California's school finance system

School district revenues in California—along with other major aspects of state policy—are strongly influenced by policies created through court decisions and the state's ballot initiative process. Ballot initiatives like Proposition 13 (1978) and 98 (1988) led to major structural changes to the state's school system. Proposition 13 restricted the ability of local school districts to set their property tax rates—giving teeth to the early efforts to equalize school funding after the 1971 Serrano v Priest decision from the California Supreme Court (Silva & Sonstelie 1995). Proposition 98 established a statutory commitment on the part of the state government to fund K-14 schools at

[15] Although the relationship between the revenue volatility index and student achievement is positive, the coefficient is not statistically significant at typical levels.

a consistent level (LAO 2017)[16]. Real per-pupil spending is required to be maintained from year to year, with increases mandated by a formula tied to growth in state per capita incomes and the General Fund.

Together, these two policies established much of the state's school finance system. Proposition 98's funding guarantee determined the level of school spending in the state and Proposition 13 essentially completed the construction of the revenue limit system that determined how state school funding was distributed to districts. The system established a prescribed level of per-pupil funding for school districts throughout the state that was meant to equalize per-pupil revenues across districts. These funding levels were high enough that most districts[17] were not able to reach their revenue limit through property tax receipts alone—the state provided the remainder, with general aid that pushing each district up to their equalized target level. The revenue limit system was reformed in 2014 with the Local Control Funding Formula (LCFF). The LCFF simplified the complex collection of historical adjustments and categorical programs of the previous system, replacing it with a combination of large, general purpose base grants and additional targeted funding for high-need students.

Despite the consistency afforded by Proposition 98, the funding environment for California school districts was not entirely stable, and districts constantly face revenue volatility. Most often, this volatility is only potential and comes through the form of uncertain revenues from sources districts do not control. Substantial volatility has been experienced, however, with the substantial increases in funding throughout the late 1990s and mid-2000s, both of which were followed by

[16] Proposition 98 establishes a "minimum funding guarantee" for annual state spending on K-12 schools and community colleges. The guarantee is typically about 40 percent of the state general fund but can sometimes be determined by formulas based on changes in student enrollments, state per capital personal income or the size of the state general fund. Occasionally, the requirements of Prop. 98 can be suspended by a vote of the state legislature as well.

[17] A subset of districts throughout the state collected enough property tax revenues to meet and even exceed their revenue limits. These districts were termed "basic aid districts" as they only received minimal foundation aid from the state that was provided to all districts. When property tax receipts for these districts exceeded state revenue limits, these districts were entitled to keep the excess.

major recessions and necessitated significant cutbacks. Figure 1 shows inflation-adjusted per-pupil school district revenues over time. The significant reduction in the pace of revenue increases during the recessions of 2000 and 2008—as well as furloughs and deferred payments—created substantial pressure on school districts to adjust their spending plans. Further, in the face of rising costs related to personnel and post-employment benefits, these changes in revenues resulted in significant reductions in the resources available to school districts (AIR 2011).

3.3 Related literature

The central question of this paper is whether school districts, acting as economic agents, smooth consumption over time in a manner similar to individual households or firms. Gruber (1997) provides an instructive example that considers the smoothing benefits of unemployment insurance. Given standard preferences—a utility-maximizing social welfare function—and discounting the value of future consumption, the typical model (Hall 1978) predicts that actors will seek to smooth consumption by spending a smaller proportion of income in periods when incomes are higher than average in order to accumulate savings that can then be spent in times when incomes are lower—for example, during an unemployment spell. Alternatively, actors may incur debt when incomes are low and pay back obligations when incomes rise again. This is functionally the same as saving, though with added concerns about the real costs of borrowing (as opposed to the opportunity costs of saving). The prediction of consumption smoothing in economics is closely related to the permanent income hypothesis or rational expectations. Despite the theory, research on whether the model reflects typical behavior is quite mixed. Hall (1978) rejects the hypothesis, while Sargent (1978) finds the opposite. Flavin (1981) reconciles the two, finding that the two approaches, when appropriate specified, are equivalent—and both reject the hypothesis. Flavin continues on with an alternative test of the hypothesis and convincingly rejects the hypothesis yet

again using a structural, as opposed to reduced form, model. Zeldes (1989) finds that liquidity constraints play a significant role in preventing the consumption smoothing that is implied by the permanent income hypothesis.

The school district context is different in a variety of ways, however, and the context of California district is even more different still. The two particular concerns are that California school districts have little control over their revenue generation and almost no ability to borrow significant sums in order to smooth consumption over time. Thus, in the California school district setting, the consumption smoothing objective is met almost primarily through the retention of revenue increases in reserve accounts.

The impacts of and approaches to revenue volatility are key questions in the public finance literature on sub-national state and local governments. School districts feature prominently, but state governments and municipalities have been extensively researched as well.

Poterba (1994) reviews state responses to unexpected fiscal shocks for state governments between 1988 and 1992 and finds that states typically respond quickly to adjust to shocks, both increasing revenues and reducing spending. Further, he finds that constitutional restrictions requiring balanced budgets and similar policies play a role in increasing the responsiveness of state spending to revenues. Partisan political consolidation also increases responsiveness. Chakrabarti (2014) and Hall and Koumpias (2018) both find that school districts respond to revenue shortfalls with tax increases, though the revenues raised are relatively small and reflect only part of the full fiscal response. In the case of California school districts, however, the limitations to debt accumulation and tax rate setting mentioned above, put even more pressure on spending responses to shortfalls and heightens the need to pre-emptively reduce exposure to revenue volatility through savings (Vanyolos 2005).

Lavertu and St. Clair (2018) examine the volatility associated with district errors in revenue forecasting and find that errors lead to significant declines in student learning largely caused by increased rates of teacher attrition. More broadly, Figlio (1997) finds in looking at tax revolts from 49 states that the reduced funding led to larger class sizes, lower salaries and worse student performance. Given that personnel—and teachers in particular—are one of the primary expenses of school districts, fiscal responses to revenue shortfalls often incorporate some adjustment to the teaching workforce. Negative student outcomes are associated with various adjustments such as hiring new teachers (Rockoff 2004), moving teachers between schools (Hanushek and Rivkin 2005) or shuffling teachers within schools (Atteberry 2017).

This paper contributes to the literature on the effects of revenue volatility in local governments—school districts in particular. Our results concern the unique context of California school districts. Districts are much more constrained in the available fiscal responses to revenue shortfalls compared to other local governments and even other school districts throughout the US. Further, with nearly 30 years of data, we can estimate funding relationships over a substantial period of time. This allows us to the investigate the impact of changing economic conditions in the state as well as major policy changes that have occurred over time.

3.4 Methods

In order to assess how California school districts manage their finances in the face of revenue volatility, we employ regression analysis to measure the relationship between revenues and expenditures. Using nearly 30 years of data for each of the state's hundreds of districts, we estimate a regression of the form,

$$Y_t = \beta_0 + \beta_1 X_t + \beta_2 Z_t + \epsilon_t$$

where Y_{it} represents the one-year percent change in expenditures and X_{it} the annual percentage change in revenues at time t. We run separate regressions for each district in order to focus on the idiosyncratic relationship between revenues and expenditures for each district without the restrictions of a pooled model.[18] The coefficient, β_1 , gives the proportion of annual revenue changes that are reflected in concurrent expenditure changes, which we describe as a revenue volatility index. This index then gives a summary measure of how a particular district manages the variability of its revenues. A value closer to 1 indicates that districts do less smoothing, instead spending new revenues as they are received or cutting expenditures to match revenue shortfalls. Conversely, a lower index shows that districts instead smooth consumption—reducing the variation in spending despite changes in revenues.

We continue our analysis with an exploration of the determinants of the individual district revenue volatility indices. Using multiple regression with standard errors clustered at the district, we estimate the relationship between the volatility index and a variety of district characteristics. These include measures of district scale like total enrollment and revenue, as well as student traits like the share that are low-income and the distribution of enrollment across racial groups. Broader district characteristics include district type—that is, whether a district elementary, high school or unified—and the region in the state where the district is located.

Our analysis concludes by relating the expenditure volatility index to two measures of district quality: class sizes and student achievement. The first measure reflects quality in terms of inputs. Given that smaller class sizes are often equated with a better, more productive classroom environment and learning experience, districts that better manage their funds may also lean on this

[18] A pooled model with year and district fixed effects would fix the constant and require joint estimation of variances across districts. Given our focus on the unique path of revenues and expenditures, we trade some of the precision we may have in the estimation of a general relationship for a design that instead allows us to focus on the variation across districts. Results from a statewide regression are given in Appendix XXX.

kind of productive investment. The second measure—achievement—aim to assess quality in terms of outputs. Despite the broad mission of the state's schools—better performance on the state standards tests is a common first approximation of the overall productivity of a district and its activities. We analyze whether the management advantages implied by a district's volatility index correspond to similar differences in test scores.

3.5 Data sources and summary statistics

The analysis described here relies on detailed financial data that is publicly available from the California Department of Education. The Standardized Account Code Structure (SACS) data contains detailed information on annual revenues, expenditures, and reserves. These data include not just totals, but also the composition of district revenues—the shares coming from various sources such as general state aid, property taxes and federal programs—as well accounts of what funds are spent on, from teachers to facilities maintenance and much more. Figures 1 and 2 show revenues by source over time and expenditures over time, respectively. Both have risen steadily over the years, though the aforementioned recessions in 2001 and 2008 show substantial slowdowns in that pace of growth. The SACS system was preceded by a similar[19] system called the J-200, which was phased out in 2003. Together, the finance data span the years 1995 to 2021[20]. Totals are adjusted for inflation using the implicit price deflator and are reported in 2021 dollars.[21]

[19] During the transition to the SACS system in 2003, the districts in the state used both the SACS and J-200 to report their finances. Having both systems in use in the same year, facilitates the creation of a direct crosswalk between the two system. The SACS system is more detailed than the J-200 in how it allows for more fine-grained accounting of the uses of funds, but we establish the still-substantial level of detail in the J-200 and carry it across the full span of time covered by the newer SACS system.

[20] The publicly available J-200 data are actually available going back to 1987-88, but district average daily attendance counts are not included prior to 1995-96. Given the difficulty establishing valid per-pupil finance totals with enrollment information from other sources, we opt to drop these earlier years and instead use the 27 more recent years instead.

[21] The implicit price deflator for state and local governments is used here in lieu of the consumer price index in order to reflect the bundle of goods that school districts typically buy. Further, we use quarterly estimates of the deflator in order to measure inflation over the school year, from Q3 to the following Q2, rather than the calendar year.

Figure 3.1 School district revenues over time by source

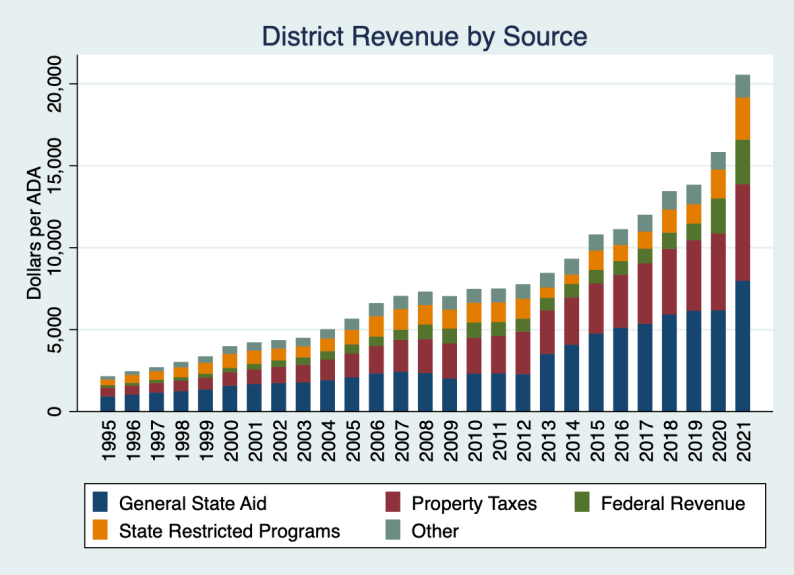

Figure 3.2 School district expenditures over time

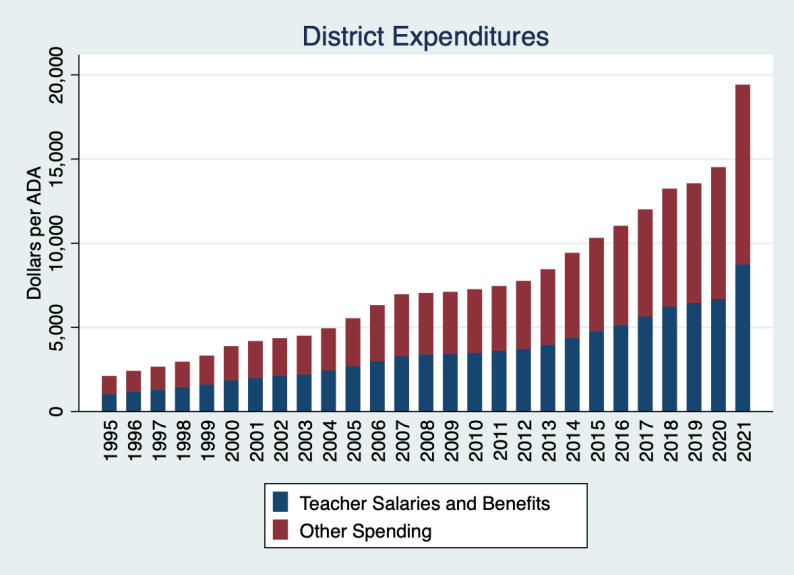

The finance data are supplemented by data on a variety of metrics also publicly available from the CDE. We use student demographic information at the district level on race and poverty, which is measured by participation in the free- or reduced-price meals program. Participation in the school meals program only provides an approximation of student family income. Rates are sensitive to differences in take-up of the program, and the binary nature of the measure does not

capture gradients in income both below and above the eligibility cutoff. We measure enrollment by average daily attendance (ADA), which gives the number of students in school on average rather than the slightly higher number of students counted on census day in October. ADA is used as the pupil count for the basis of funding for most programs. Despite that however, on the questions of district financial management that we explore here there is valid concern that total enrollment rather than average attendance ought to drive our analysis. After all, districts ought to have resources available for all of the students that may come on any given day, independent of the usual pattern of absences. Unfortunately, data on enrollment that matches the districts as they are formulated in the financial data are not readily available.[22] Checking the relationship between ADA and the enrollment data we have, we find that the ratio of ADA to enrollment is somewhat relatively constant at about 0.92.

We obtain teacher counts—in order to determine class sizes—from teacher salary and benefit data (Form J-90) available from the CDE. The data includes information on count of full-time equivalent, or FTE, teachers and other staff at each district. Student achievement is measured by performance on the state standards test. The California Standards Test (CST) was used from the 2002-03 to 2012-13. The LCFF reforms included a transition to a new test, the Smarter Balanced Assessments—often referred to as the SBAC tests. The new exams differed from the previous system in a number of ways. They were aligned with the state's new Common Core curriculum, and they were computer adaptive.[23] Although comparisons of performance across the

[22] Enrollment data is available at the school level while the SACS data are collected at the district. The school level data include counts for charter schools within a district, but it is difficult to accurately account for charter schools in the financial data given that the way in which charter revenues and expenditures are reported varies based on individual school and district practice.

[23] "Computer adaptive" refers to the manner in which questions are presented to students taking the SBAC. Students all answer different sets of questions. The difficulty of questions is adjusted dynamically based on whether students answered the previous questions correctly. This is meant to reduce the time required for the examination and produce more accurate assessments of student learning.

two tests is discouraged by the CDE, our focus is on the differences across districts within each given year.

California's school system includes nearly 1,000 districts, 10,000 schools and over 6 million students. There is a wide variety of districts, but to facilitate an analysis that is broadly representative of the experience of most students in the state, we drop a subset of districts that are substantially unlike the districts that serve most students. In particular, we limit our consideration to districts with at least 300 students. This removes 343 of the smallest districts in the state, but the remaining districts serve over 97 percent of the state's students. Further we remove districts that are not in the sample for all 27 years of our data (1995-2021). This is to ensure that the estimation of the revenue volatility index is made using the same set of years—and the broader macroeconomic context that that implies. The vast majority of districts are in the sample for every year, but this restriction removes 66 more districts. These districts again, however, only account for a small share of total enrollment (2.2%).

Figure 3.3 Average daily attendance

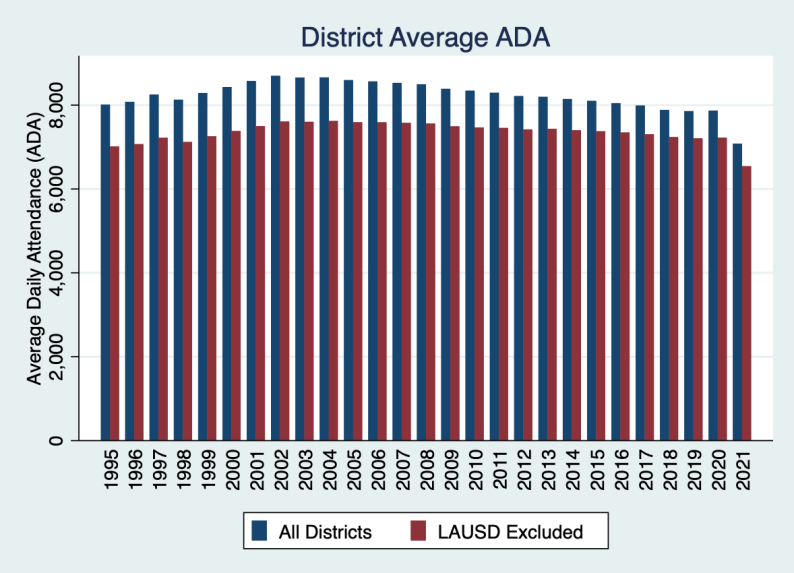

Our final sample of districts includes 628 districts. About 44 percent of districts are unified which serve grades K-12. Another 44 are elementary, with the remainder—11 percent of districts—being high school districts. However, elementary districts are typically much smaller than other district types. In fact, about 72 percent of students are enrolled in unified districts despite their lower share of the number of districts. Average district enrollment has hovered around about 8,000 for the years we cover. Los Angeles Unified School District, one of the largest in the entire country, is a massive outlier however with an average of 560,000 students over the years in our sample. As we see in Figure 3, when we exclude that district, the average ADA in our sample is closer to 7,000.

Figure 3.4 Student demographics – Student family low-income status

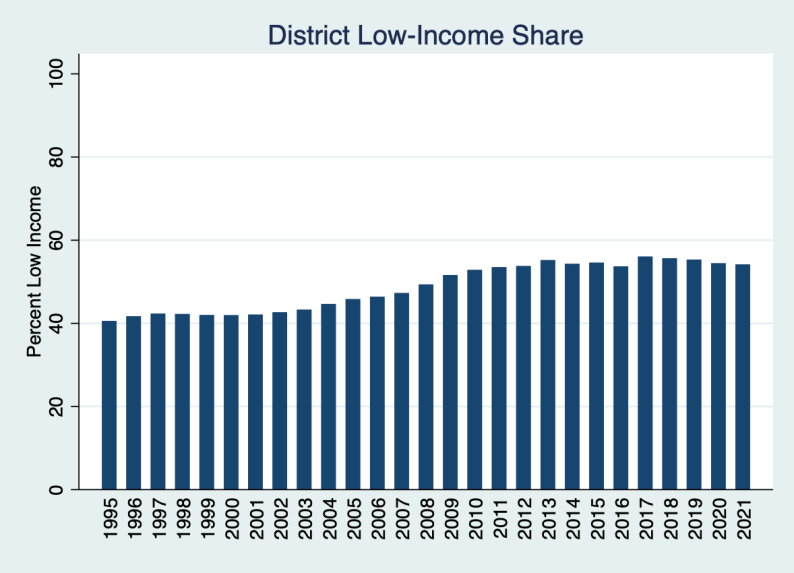

Looking at student demographics, we see in Figure 4 that the share of low-income students in a district—based on participation in the free or reduced-price lunch program—has grown from 40 percent to over 50 in recent years. And in Figure 5, the state's race/ethnicity distribution has seen the share of Latino and Asian students rise, while White and Black enrollments have fallen in relative terms.

Figure 3.5 Student demographics – Student race/ethnicity distribution

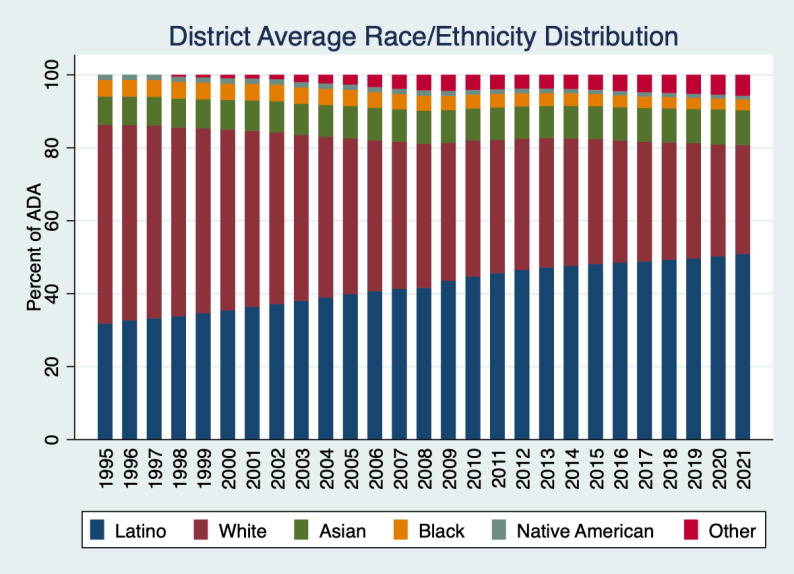

3.6 Results

Figure 3.6 Average annual changes in revenue and expenditures

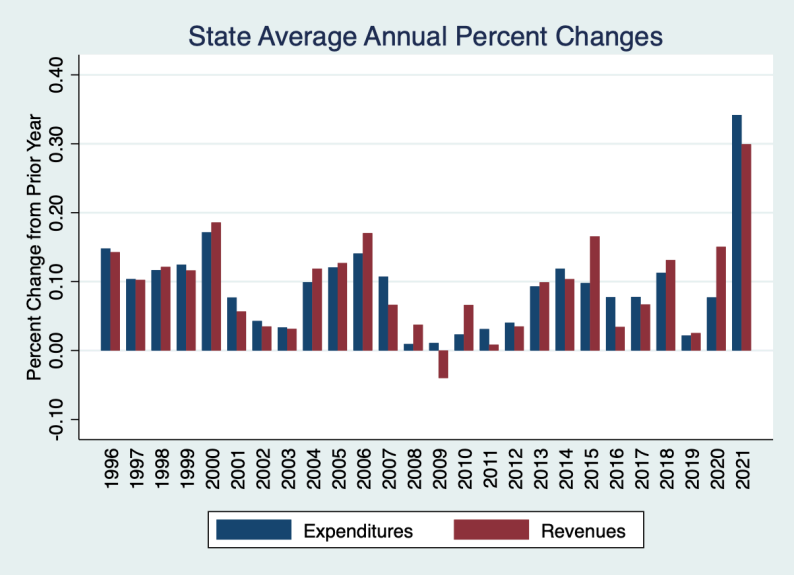

We begin by considering the relationship between revenue changes and expenditure changes statewide. Figure 6 shows the state averages of changes in both revenues and expenditures

over time. We see that the two series generally track each other closely, with larger gaps between the two when changes are relatively large or small. We begin our regression analysis by regressing the one-year percentage change in expenditures on the same change in revenues for each district. District revenue volatility indices—that is, the coefficient β_1 in the regression model described above—are estimated by regressing expenditures on revenues over the 27 years of data available for each district. The average revenue volatility index of 0.80, which indicates that over the years studied, the typical district matched responded to changes in revenues with corresponding changes in expenditures of about 80 percent of the magnitude. That is, we revenues rose by $100 per pupil, expenditures followed, rising by about $80 on average. There is considerable variation, however. Figure 7 shows the distribution of the revenue volatility index. A quarter of districts have indices lower than 0.71 while the same proportion have indices greater than or equal to 0.90. In fact, 55 of the 628 districts have an index greater than 1.00, which is possible due to reserves that existed prior to the start of our sample of data and allowed for districts to spend beyond their immediate means.

Figure 3.7 Distribution of revenue volatility index, β_1

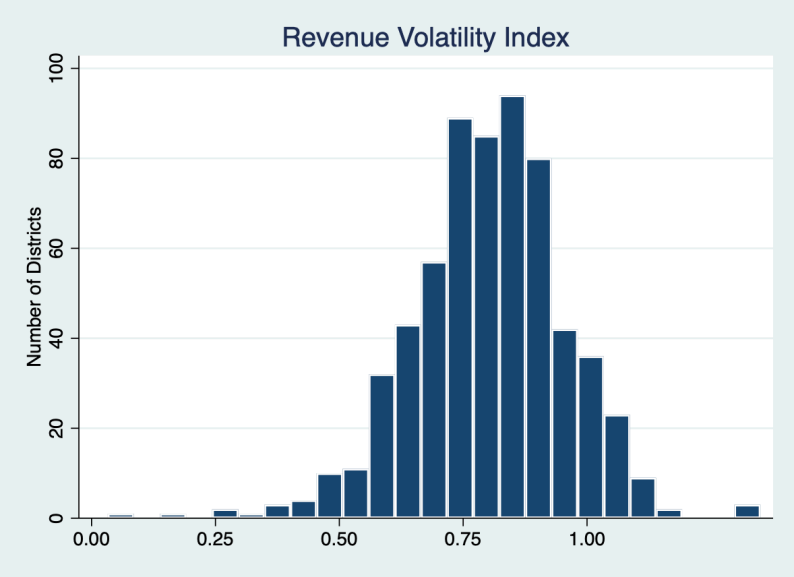

Given the considerable variation in the revenue volatility indices calculated for each district, we continue by investigating the correlates of this measure. If there are common, strong determinants of how well districts smooth expenditures or not, then an understanding of them would be useful. We consider a variety of measures that we suspect could relate to how districts manage revenue volatility. These include district type and region, as well as student demographics and both total enrollment and revenues. Each of the factors could play a role in determining how districts are able to respond to volatility in their revenues. Larger districts, both in terms of enrollment and revenue, may be able to mitigate volatility more easily than smaller ones due to economies of scale and a greater capacity for dedicated administrative staff. Districts in communities with greater shares of disadvantaged students typically receive more targeted funding, which could help manage some of the regular volatility most districts face—though those sources can be volatile themselves, which could make smoothing more difficult. While students in all grade spans need teachers, the mix of personnel and other resources needed can vary between them. These differences could relate to how districts respond to revenue changes as well. Also, districts in different regions face labor markets that differ—sometimes dramatically so. Regional economic conditions could relate to the volatility districts face and the options available to them to respond.

We test four different specifications, with results given in Figure 8. The first includes only region and district type. Next, we add student demographics, race and income. Then we add enrollment, and finally total revenues. We find that there are some significant differences in the revenue volatility index across regions. In particular, the North Coast, Bay Area and Central Coast regions saw statistically significant larger volatility indices. The share of low-income students had a consistently significant positive effect across our specifications. District shares of Asian, Latino

and Black students were statistically significant and negative in some specifications, which indicates that districts with more students from these race and ethnicity groups experienced more smoothing than other districts. District type, enrollment and total revenues were all not statistically significant in any specification.

Table 3.1 Determinants of the revenue volatility index, β_1

VARIABLES	β_1	β_1	β_1	β_1
Superior	0.0313	0.00587	0.00556	0.0158
	(0.0277)	(0.0311)	(0.0313)	(0.0316)
North Coast	0.0861**	0.0621*	0.0616*	0.0631*
	(0.0341)	(0.0361)	(0.0365)	(0.0364)
Bay Area	0.0493*	0.0748**	0.0745**	0.0650**
	(0.0274)	(0.0303)	(0.0305)	(0.0308)
N San Joaquin	0.0268	0.00809	0.00768	0.0162
	(0.0309)	(0.0322)	(0.0325)	(0.0327)
Central Coast	0.0668**	0.0644**	0.0643**	0.0677**
	(0.0310)	(0.0312)	(0.0312)	(0.0312)
S San Joaquin	0.0291	0.0138	0.0134	0.0181
	(0.0285)	(0.0296)	(0.0299)	(0.0300)
Inland Empire	0.0122	0.0129	0.0131	0.0200
	(0.0318)	(0.0324)	(0.0325)	(0.0326)
Los Angeles Cty.	0.0177	0.0326	0.0326	0.0344
	(0.0291)	(0.0298)	(0.0299)	(0.0298)
Orange Cty.	0.0323	0.0416	0.0420	0.0418
	(0.0375)	(0.0384)	(0.0387)	(0.0386)
Elem. District	0.0171	0.00950	0.00892	0.0162
	(0.0134)	(0.0137)	(0.0150)	(0.0154)
HS District	-0.0259	-0.0200	-0.0203	-0.0296

	(0.0212)	(0.0218)	(0.0219)	(0.0223)
% FRPM		0.112*	0.112*	0.0890
		(0.0600)	(0.0605)	(0.0613)
% Other Race		-0.0818	-0.0849	-0.115
		(0.384)	(0.385)	(0.385)
% Amer. Ind.		-0.126	-0.126	-0.167
		(0.182)	(0.182)	(0.183)
% Asian		-0.121*	-0.119*	-0.106
		(0.0687)	(0.0719)	(0.0720)
% Latino		-0.100*	-0.0997*	-0.0906
		(0.0586)	(0.0592)	(0.0592)
% Black		-0.275**	-0.271**	-0.271**
		(0.131)	(0.137)	(0.137)
Log ADA			-0.000663	-0.0853**
			(0.00694)	(0.0418)
Log Rev.				0.0887**
				(0.0432)
Constant	0.762***	0.781***	0.786***	-0.0199
	(0.0244)	(0.0361)	(0.0694)	(0.399)
Observations	628	628	628	628
R-squared	0.027	0.040	0.040	0.046

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

3.6.1 Revenue Volatility and District Outcomes

The revenue volatility index measures the extent to which districts are able to dampen the link between the instability in revenues and that in expenditures—shielding students and staff from the disruption to the district's budget. We examine the relationship between this index and two

measures—test scores and class sizes—to get a better sense of the impacts of volatility. In both cases, we must aggregate the outcome measure over time to match the single, time-invariant revenue volatility index for each district. We measure the relationship between the district volatility index and district averages of our two outcomes, but we investigate this aggregation issue further when considering test scores.

Starting with class sizes, our data give us full-time equivalent, or FTE[24], counts of teachers in each district from 1999 to 2020. We obtain class sizes by taking the simple ratio of ADA to FTE teachers. Of course, this does not exactly capture class sizes in the sense that we have a count of the average number of students in a classroom in each district. Instead, the measure reflects the district's investment in personnel, scaled by enrollment to take into account the additional staff needed with greater levels of enrollment. In Figure 9, we see class sizes have hovered around 18 students per teacher. They rise during the 2008 recession and remain higher during recovery thereafter before falling again to the long-term average with the passage of the LCFF in 2014. The box plot shows that there has also been an increase in variability of class sizes in the later years.

[24] These counts are adjusted to reflect differences in the hours worked for various staff. For example, if a part time teacher only works half of the school day, that teacher is then counted as 0.5 FTE teachers.

Figure 3.8 Class sizes over time

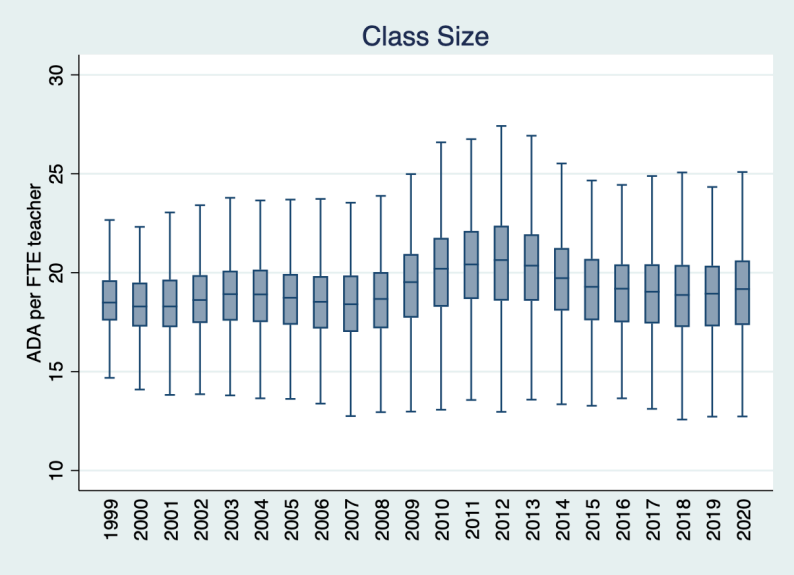

In order to aggregate class sizes across the years of data available, we standardize class sizes within year—describing class sizes in terms of the difference (in terms of standard deviations) from the state average each year. We then take the average over all the years of data, arriving at a single number the captures the relative differences between districts in their typical class sizes. We regress this average class size on the district revenue volatility index. Regression results are shown in Figure 10. We add a second specification that includes district type because class sizes are generally smaller in the earlier grades—a state program even provides funding for districts to reduce class sizes up to grade 3. In both specifications we see that there is a negative relationship between the volatility index and average class sizes. The coefficient is statistically significant in both models as well. District revenue volatility indices range from 0.71 to 0.90 between the 25th and 75th percentile. The coefficient -0.479 in the second specification would then imply that class sizes would be about 0.09 standard deviations—or 0.21 students per teacher—smaller in districts at the 75th percentile compared to those at the 25th.

Table 3.2 Revenue volatility index and class sizes

VARIABLES	Avg. Class Size	Avg. Class Size
β_1	-0.527**	-0.479**
	(0.214)	(0.214)
Elem. District		-0.0388
		(0.0701)
HS District		0.287**
		(0.112)
Constant	0.408**	0.356**
	(0.175)	(0.177)
Observations	626	626
R-squared	0.010	0.023

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The result here is the opposite of what we would predict. The districts with the higher volatility indices—those that smooth less—tend to have class sizes that are smaller rather than larger. Despite the relative instability in their spending over time, these districts can maintain lower ratio of students to teachers than other districts. Perhaps these districts make cuts elsewhere when shortfalls occur. Or, alternatively, they invest more in teacher staffing when revenues are high, and average class sizes over time just turn out to be smaller than other districts in the years we observe.

The next set of analyses focuses on student performance. We measure this by using standardized test scores from both the California Standards Test (CST), which was administered from 2002 to 2012, and the Smarter Balanced Assessments (SBAC) that were introduced with the

LCFF and have been used every year since 2014. Much like the aggregation method used with class sizes, we normalize test scores, describing them as deviations from the mean in units of standard deviations. An added complication here though is that test scores are normalized against means for specific combinations of year, grade, test subject (math or reading), and demographic group (race, low-income status and English learner status). We take district averages of these normalized scores in order to get a measure of student performance that is not sensitive to differences across districts in their distribution of enrollment across grade span and student group. Test score performance also varies by grade level (Cascio & Staiger 2012), so this adjustment mitigates some of those concerns as well. While tests and their scores vary over time[25], our intention with the aggregation here is to produce a measure that indicates the relative performance between districts on average. The results in Figure 11 show a positive, but statistically insignificant relationship between volatility index and student performance. These suggest that districts with less consumption smoothing (i.e. a higher volatility index) have better performance on average than other districts.

[25] Test scores vary on two dimensions here. Of course, variations in student performance over time are observed, with scores rising in most years as students learn more and adjust to the test-taking system. However, scores also differ due to the change from the CST to the SBAC in 2014. The new tests were based on new state standards and were administered on computers rather than with pencil and paper. There was a period of adjustment during the transition, with proficiency rates falling in the first years of the exam before recovering subsequently. The standardization of the scores allows for comparability across time by preserving the interpretation of the relative performance of students across districts.

Table 3.3 Revenue volatility index and student achievement

VARIABLES	Mean test scores	Grade 11 cohort
β_1	0.0709	
	(0.180)	
Elem. District	0.0929	
	(0.0590)	
HS District	0.185*	
	(0.0944)	
β_1^*		0.0528
		(0.0984)
Constant	-0.00711	0.134*
	(0.149)	(0.0767)
Observations	628	4,299
R-squared	0.008	0.000

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Despite normalizing and averaging the two outcomes reviewed above, the aggregation of outcomes over time is still problematic. We address this concern with an alternative specification. Students take their final state tests in grade 11. Rather than using the full 27 years of data for each district to calculate the revenue volatility index, we instead use only the 12 years covering kindergarten up until grade 11 for each cohort of students. We then regress grade 11 performance

on the revenue volatility index generated based on the prior 12 years of data. This approach has some advantages and disadvantages. We are able to account for how our outcome, student test scores, evolve over time, but with fewer years of data with which to generate the volatility index, the measures are somewhat more imprecise. Figure 13 shows a comparison of the distribution of district volatility indices over time using either all of the years for a single measure or using only the years in a given cohort. We see that the range of the volatility index is generally larger when we use just the cohort years. Further, the volatility index appears to be sensitive to the set of years it is based on. The regression results again show a positive relationship between the volatility index and student performance, though this is statistically insignificant as well.

3.7 Conclusion

Faced with volatile revenues that are largely out of their control, we find that California school districts have had mixed success managing unstable revenues by preventing similar swings in their annual spending. On average, about 90 percent of revenue changes are translated to shifts in spending. Though not pure myopia, districts in the state typically match their spending to whatever windfall or contraction they receive in revenues each year. This can be a boon in good times, but disruptive when local or statewide economic conditions sour.

There is considerable variation in how different districts manage the volatility they face. Some are able to nearly perfectly smooth expenditures while others match spending changes to revenues one-to-one (and a few even more so). Districts in the Bay Area and other counties along the coast tend appear to have higher volatility indices—that is, less expenditure smoothing. Also districts that have greater shares of low-income students smooth less, while those with more Asian, Latino and Black students smooth more. This suggests that this question of management is both a matter of district capacity and the impact of various conditions facing different districts.

Despite the hypothesized downsides of budget instability and the lack of expenditure smoothing, we find in our regression analysis, that the outcomes we measure—class sizes and test scores—had the opposite relationship to the revenue volatility index. Districts that smoothed less saw, on average, smaller class sizes and better test scores. This result could have something to do with the particulars of revenue changes in the recent history of the state. When districts are given additional funding described as one-time funds as opposed to ongoing funding, then perhaps it is prudent simply spend the new money immediately. There is no potential unexpected shortfall over the horizon, but instead the simple planned expiration of a revenue program. This, and other dynamics like it, can be explored in future work.

Given the structure of revenues for California school districts, as well as the broader state budget, revenue volatility continues to be a concern. When the seemingly persistent fears of a recession for the state and the nation begin to look more prescient, the costs of volatility—in addition to the negative effects of declines per se—threaten to be significant. With little means of their own beside prudent savings in times of revenue growth, California districts must rely on efforts from the state that encourage stability like Proposition 2's rainy day fund or hope for broader reforms to the state's funding structure.

References

Atteberry, A., Loeb, S., & Wyckoff, J. (2017). Teacher churning: Reassignment rates and implications for student achievement. Educational Evaluation and Policy Analysis, 39(1), 3-30.

Cascio, E. U., & Staiger, D. O. (2012). Knowledge, tests, and fadeout in educational interventions (No. w18038). National Bureau of Economic Research.

Chakrabarti, R., Livingston, M., & Roy, J. (2014). Did cuts in state aid during the Great Recession lead to changes in local property taxes?. Education Finance and Policy, 9(4), 383-416.

Figlio, D. N. (1997). Did the "tax revolt" reduce school performance?. Journal of Public Economics, 65(3), 245-269.

Flavin, M. A. (1981). The adjustment of consumption to changing expectations about future income. Journal of political economy, 89(5), 974-1009.

Gruber, J. (1997). The Consumption Smoothing Benefits of Unemployment Insurance. The American Economic Review, 87(1), 192–205. http://www.jstor.org/stable/2950862

Hall, J. C., & Koumpias, A. M. (2018). Growth and variability of school district income tax revenues: Is tax base diversification a good idea for school financing?. Contemporary Economic Policy, 36(4), 678-691.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence. Journal of political economy, 86(6), 971-987.

Hill, Elizabeth G. "Revenue volatility in California." Report, Legislative Analyst's Office of California, Sacramento (2005).

Kapphahn, K., & Kuhn, J. (2017). A historical review of Proposition 98. Legislative Analyst's Office.

Lavertu, S., & Clair, T. S. (2018). Beyond spending levels: Revenue uncertainty and the performance of local governments. Journal of Urban Economics, 106, 59-80.

Poterba, J. M. (1994). State responses to fiscal crises: The effects of budgetary institutions and politics. Journal of political Economy, 102(4), 799-821.

Rivkin, S. G., Hanushek, E. A., & Kain, J. F. (2005). Teachers, schools, and academic achievement. Econometrica, 73(2), 417-458.

Rockoff, J. E. (2004). The impact of individual teachers on student achievement: Evidence from panel data. American economic review, 94(2), 247-252.

Sargent, T. J. (1978). Rational expectations, econometric exogeneity, and consumption. Journal of Political Economy, 86(4), 673-700.

Shambaugh, L., Kitmitto, S., Parrish, T., Arellanes, M., & Nakashima, N. (2011). California's K-12 education system during a fiscal crisis. American Institute for Research.

Silva, Fabio, and Jon Sonstelie. "Did Serrano cause a decline in school spending?." National Tax Journal 48.2 (1995): 199-215.

Vanyolos, I. (2005). Fiscal reaction to state aid uncertainties: evidence from New York State. Public Budgeting & Finance, 25(4), 44-58.

Zeldes, S. P. (1989). Consumption and liquidity constraints: an empirical investigation. Journal of political economy, 97(2), 305-346.