The Fiscal Consequences of Special District Consolidation: Evidence from California*

Travis St.Clair

April 7, 2026

Abstract

The number of governments in the United States has increased steadily since the 1970s, largely due to the proliferation of special districts. This growth has fueled concerns over the efficiency of public service provision, as many metropolitan areas rely on a fragmented network of jurisdictions to deliver services. However, due to a lack of exogenous variation in the number of districts, identifying the causal impact of government fragmentation is challenging. To address this challenge, this paper exploits California's Cortese-Knox-Hertzberg Local Government Reorganization Act of 2000, which established procedures for the consolidation and annexation of cities and special districts. Using synthetic control methods, I show that the Act reduced the total number of local governments in the state by seven percent relative to its counterfactual trajectory, driven by a 15 percent decline in the number of special districts. Despite the relative decline in the number of governments, the total amount of local government spending in the state remained unchanged, suggesting that the slowing growth in jurisdictions was offset by higher spending per government. Efforts to reduce fragmentation may have limited fiscal impact unless they target districts where overlapping functions or scale inefficiencies are most pronounced.

JEL codes: H11, H72, R51

Keywords: government fragmentation, government consolidation, special districts, synthetic control

*St.Clair: Robert F. Wagner School of Public Service, New York University (NYU), 105 East 17th St, 3rd Floor, New York, NY, 10003 (e-mail: travis.stclair@nyu.edu). For helpful comments, I am grateful to Byron Lutz and conference participants at the National Tax Association (NTA) annual conference and the Association for Budgeting and Financial Management (ABFM) annual conference.

1 Introduction

Since the 1970s, the United States has seen a persistent rise in the number of its local governments, driven by the proliferation of special districts – such as park districts and water supply authorities – that carry out a narrow set of functions (Smaldone and Wright, 2024). This expansion has contributed to greater fragmentation in the provision of government services, particularly in large metropolitan areas such as Chicago, Pittsburgh and St. Louis, where hundreds of overlapping jurisdictions now share responsibility for basic public functions (Hendrick and Shi, 2015). Although school districts – the most common form of special purpose entity – have consolidated significantly over time, fragmentation among other special purpose entities has continued to increase, reflecting legal frameworks that encourage the creation of independent government entities to bypass general tax and debt limits (Foster, 1997; Goodman, 2019; Berry, 2009).

This trend has fueled ongoing debates about the potential benefits and harms of fragmentation. On one hand, an increase in the number of governments enhances consumer-voter choice by enabling residents to select jurisdictions that best match their preferences and leads to more efficient allocation of local public goods (Tiebout, 1956). Polycentric systems can provide enhanced local representation and responsiveness and generate healthy competition between multiple service providers (Ostrom et al., 1961). Recent empirical work exploiting boundary changes in developing countries finds that smaller, more decentralized general-purpose governments can improve service delivery through stronger political incentives and civic engagement (Narasimhan and Weaver, 2024; Dahis and Szerman, 2025). On the other hand, government fragmentation – be it vertical fragmentation, with multiple levels of governments, or horizontal fragmentation, with many governmental units at the same level – prevents governments from taking advantage of economies of scale and may lead to administrative redundancies, which in turn contribute to "overfishing" from a common tax base (Berry, 2008; Oates, 1972, 1999). Empirical studies focusing on the United States have

generally found that fragmentation leads to inefficient increases in the size of the public sector (Goodman, 2019, 2015; Berry, 2008; Dolan, 1990). However, some of this work finds important distinctions across forms of governments, with the expansion of special districts in particular an inefficient driver of bloat (Zax, 1989; Hendrick et al., 2011). International evidence from Indonesia and sub-Saharan Africa also finds support for the idea that high levels of regional fragmentation reduce administrative capacity and diminish scale efficiencies in service provision (Grossman and Lewis, 2014; Grossman et al., 2017; Lewis, 2017; Cassidy and Velayudhan, 2026).

One of the most commonly suggested urban reforms, consolidation, offers a potential remedy to these challenges, promising greater economies of scale and fewer redundancies. Consolidation can take many forms, including city-county (vertical) consolidation, school district consolidation, municipal consolidation, and the consolidation of special districts, each of which raises distinct concerns about administrative capacity and the effects on the quality of service provision (Jimenez and Hendrick, 2010). Amalgamation reforms suggest that mergers of general-purpose municipalities often succeed in reducing administrative expenditures (Blom-Hansen et al., 2014; Reingewertz, 2012; Blesse and Baskaran, 2016). Evidence from school district consolidations similarly points to cost efficiences from economies of scale in operating expenditures (Duncombe and Yinger, 2007). In contrast, studies of city-county consolidation find little impact on per capita spending (Carr and Feiock, 2004; Faulk and Grassmueck, 2012)

Despite its breadth, the consolidation literature has little to say about the consolidation of special districts, despite the fact that the growth of single-function governments has driven most of the expansion in the number of local governments over the past half century (Smaldone and Wright, 2024) and remains central to debates over cost efficiency. An important advantage of focusing on special-purpose districts is that changes in their boundaries tend to be orthogonal to political outcomes in other jurisdictions. This contrasts with much of the decentralization and federalism literature, which typically examines reforms that simul-

taneously alter fiscal authority, political accountability, and electoral structures (Treisman, 2007).

To address this gap, and to provide causal evidence, I exploit a state law that led to an exogenous change in the number of jurisdictions, and I leverage this change to study the effects of consolidation. In 2000, California passed the Cortese-Knox-Hertzberg (CKH) Government Reorganization Act, which established procedures for changes of organization to local governments, including city incorporations, annexations to a city or special district, and city and special district consolidations. This Act, as well as the commission that preceded it, was borne out of a concern for California's rapid population growth and the ensuing increase in overlapping service boundaries and urban sprawl. Among other things, the Act empowered local agency formation commissions (LAFCOs) in each county to review and initiate changes of organization so as to avoid excessive fragmentation and encourage local governments to be as effective and responsive as possible. In this regard California is not alone; several states have made efforts in line with California's to address fragmentation by streamlining the procedures by which local governments can reorganize. New York, Florida, and Colorado have all passed laws establishing procedures for the consolidation and dissolution of certain types of local governments, with the aim of minimizing duplication among service providers and avoiding double taxation.[1] However, the large number of local governments in California, and the amount of time that has passed since CKH, make the state an especially informative setting for evaluating the impact of state-led efforts at consolidation.

To assess the impact of the law, I use the synthetic control method (SCM). First developed by Abadie and co-authors as a means of assessing the causal impact of a policy change affecting a single treated region (Abadie and Gardeazabal, 2003; Abadie et al., 2010, 2015), the approach involves the construction of a synthetic control group – a weighted combination of comparison units chosen to minimize the root mean square error (RMSE)

[1]See New York's New N.Y. Government Reorganization and Citizen Empowerment Act (2010), N.Y. Gen. Mun. Law §§ 750–793, Florida's Uniform Special District Accountability Act (1989), Fla. Stat. §§ 189.01–189.076, and Colorado's Special District Act of 1981, Colo. Rev. Stat. §§ 32-1-101 to 32-1-1807.

of key predictor variables – which enables one to compare the evolution of key outcomes in the treated region with those in a counterfactual region over time. Indeed, the SCM is well-suited for evaluating the effect of state laws that apply only to one state and for which it is otherwise difficult to construct a counterfactual. In addition to the SCM, I also conduct a supplementary analyses that draws on a balanced panel of local governments.

Using these methods, I show first that the law reduced the growth rate of local governments in California, leading the number of governments to fall by an average of seven percent relative to its counterfactual trajectory. This decrease was driven primarily by a relative decline in the growth of special districts. However, the decline did not occur among special districts providing administrative functions, where consolidation would be most likely to generate efficiency gains. Next, I show that this decline had no discernible effect on the state's finances; per capita spending and own-source revenue at the local level remained unchanged relative to the counterfactual. These results are robust to a series of specification checks that vary the size of the donor pool and the number and type of predictor variables. They are also confirmed by the supplemental analysis, which shows that surviving governments in California increased spending relative to local governments in other states, offsetting the reduction in the number of governments. The results highlight how efforts to reduce fragmentation may have limited fiscal impact unless they target districts where overlapping functions or scale inefficiencies are most pronounced.

This paper proceeds as follows. Section 2 provides background on the growth in the number of governments in the United States and the California law that established procedures for the annexation and consolidation of local governments. Section 3 outlines a conceptual framework for thinking about jurisdictional overlap at the local level and its role in economies of scale. Section 4 outlines the data, while Section 5 discusses the methods. The results are presented in Section 6. Section 7 present the results from a complementary analysis. Section 8 concludes.

2 Background

2.1 Growth in the Number of Governments

Between 1972 and 2022, the number of local governments in the United States increased by 14 percent, from roughly 78,000 to nearly 89,000. As Figure 1 shows, this increase was entirely attributable to growth in the number of special districts, which grew by more than 50 percent even as the number of school districts and general-purpose governments declined. While special districts represented less than a third of total governments in 1972, there are now nearly as many special districts as there are general purpose governments. School districts, once the dominant form of local government in the country in the 1940s, continued a trend of consolidation that began in earlier decades (Smaldone and Wright, 2024).

This increase in the number of governments has translated to a high degree of fragmentation of government services. The Chicago metropolitan area alone includes more than 1,500 jurisdictions (McCasland et al., 2023), and taxpayers living in major metropolitan areas frequently belong to a large number of taxing jurisdictions. While scholarship attributes much of this fragmentation to a desire among households for homogenous communities and the exclusion of minorities and low-income households from access to high-quality public goods (McCasland et al., 2023; Alesina et al., 2004; Monarrez and Schönholzer, 2023), there is less consensus as to why the number of special districts in particular has continued to increase – or what gains may be achieved through consolidation (Goodman, 2019, 2020).

2.2 The California Context

Following World War II, California experienced booming growth in its population, resulting in a proliferation of local governments. Many of the emergent districts were formed with little coordination or planning and with overlapping jurisdictional boundaries. In 1959, to address the problems arising from the growing fragmentation, Governor Edmund Brown appointed a

commission on Metropolitan Area Problems, whose recommendations resulted in the creation of Local Agency Formation Commissions, or LAFCOs, in each county (Bui and Ihrke, 2003). Between 1963 and 1985, LAFCOs administered three different laws, each of which established different procedures for local government boundary changes. In 1985, acknowledging that the three laws were not always consistent, the Cortese-Knox-Local Government Act of 1985 integrated the three laws (Bui and Ihrke, 2003). However, this integration did not result in procedures that were any less unclear or cumbersome (California Commission on Local Governance for the 21st Century, 2000). Moreover, LAFCOs did not possess sufficient powers to adequately address the growing number of special districts; in particular, they were unable to initiate proposals for dissolution on their own, and instead had to wait for a petition to be brought (Jeffries, 1971). In 1997, the Legislature formed the Commission on Local Governance for the 21st Century, which recommended strengthening LAFCO powers and streamlining and clarifying the procedures for local government boundary changes (California Commission on Local Governance for the 21st Century, 2000). These recommendations were formally adopted with the passage of the Cortese-Knox-Hertzberg Local Government Reorganization Act of 2000 (CKH, or the Act, from here forward).

The purpose of the Act was to clarify the procedures for changes of local government organization, but also more broadly to discourage urban sprawl and encourage intelligent growth. In particular, the Act found that "whether governmental services are proposed to be provided by a single-purpose agency, several agencies, or a multipurpose agency, responsibility should be given to the agency or agencies that can best provide government services" (California Assembly Committee on Local Government, 2024). The Act was not expressly intended to encourage consolidation per se; in fact, it clarified procedures for a range of changes of organization, including the formation of a district or incorporation of a city. However, underlying the Act's provisions was a broad concern about fragmentation that had given rise to LAFCOs and the Commission on Local Governance in the first place.

The Act formally delegated the state legislature's boundary powers to LAFCOs, em-

powering a LAFCO in each county (58 in total) to determine the boundaries of city and special districts in that county, and accordingly designate the land use authority of local governments in the county as well as their taxing and corporate powers (Bui and Ihrke, 2003).[2] LAFCOs have the power not only to review and approve proposals for changes of organization but also the power to initiate proposals for the consolidation, dissolution, merger, or formation of a district (California Assembly Committee on Local Government, 2024). Importantly, voters can't use an initiative or referendum to circumvent LAFCO approval, although elections may be required in some cases (Bui and Ihrke, 2003).

Figure 2 shows the number of local governments (scaled by area) (Figure 2A) and the number of special districts (Figure 2B) in California before and after CKH. While the number of governments in the state declined only slightly post-2000, the rate of growth fell off significantly from its prior trend. In contrast, the number of local governments in the rest of the country continued to grow during this period, though there is some suggestion of a drop-off relative to trend post 2010. This descriptive evidence suggests that CKH may have operated primarily by slowing the growth of special districts rather than eliminating existing governments per se. Thus, the consolidation examined in this paper should be understood as a reduction in the rate of fragmentation, not as the large scale merger of pre-existing districts. The analysis that follows assesses the fiscal consequences of this shift by comparing California's post-CKH trajectory to a counterfactual path in which the pre-existing growth in the number of governments would have continued.

3 Conceptual Framework

Let m_t represent the number of general-purpose governments in a region (counties, municipalities, etc.) at time t and n_t^C the number of single-function special districts that would

[2]LAFCOs regulate all city and most special district boundaries, including water districts, municipal utility districts, pest control districts, and recreation and park districts, but do not regulate boundaries for counties and a small number of local governments for which the state legislature has established alternative procedures (Bui and Ihrke, 2003).

prevail at time t absent reform. Let n_t^O represent the number of special districts observed at time t following a reform that curtails the growth of special districts. Consolidation reduces vertical fragmentation if it limits the growth of overlapping single-function districts. It reduces horizontal fragmentation if it slows the growth in the number of districts that compete with one another for residents at the same governmental tier.

Assume that all services are provided by either special districts or general purpose governments. Let C_{it}^d denote the cost of services provided by special district i at time t, and C_{jt}^g denote the total cost of services provided by general purpose government j at time t. The total cost of government services at time t along the counterfactual (no-reform) path is $C_t^C = \sum_{i=1}^{n_t^C} C_{it}^d + \sum_{j=1}^{m_t} C_{jt}^g$, while the total cost of services along the observed post-reform path is $C_t^O = \sum_{i=1}^{n_t^O} C_{it}^d + \sum_{j=1}^{m_t} C_{jt}^g$. The effect of consolidation on government spending at time t is $\Delta C_t = C_t^O - C_t^C$.

Assume a production cost function for each service, which depends on the quantity of service output (Q), managerial efficiency (E), and administrative overhead (A): $C = f(Q, E) + A$. If f exhibits economies of scale, i.e. marginal cost decreases with Q ($\frac{d^2 f}{dQ^2} < 0$), then limiting the proliferation of special districts may reduce per-unit costs relative to the counterfactual, implying $\Delta C_t < 0$. Similarly, slower growth in the number of districts may reduce duplication of administrative functions (A) and thereby limit the growth of fixed costs.

Alternatively, if consolidation does not generate meaningful cost savings ($\Delta C_t \approx 0$), then the total cost of government service provision may remain approximately constant relative to the counterfactual, even as the growth of fragmentation slows. This may occur if, for example, certain types of districts benefit from a more decentralized structure, such as those that must maintain geographically proximate capital assets (Duncombe and Yinger, 1993). Define average spending per government along each trajectory as $\bar{C_t}^C = \frac{C_t^C}{m_t + n_t^C}$ and $\bar{C_t}^O = \frac{C_t^O}{m_t + n_t^O}$. Since $m_t + n_t^O < m_t + n_t^C$, it follows mechanically that $\bar{C_t}^O > \bar{C_t}^C$. That is, consolidation that slows the growth of special districts will raise the average spending of

the surviving governments, reflecting the concentration of fiscal responsibilities that would otherwise have been distributed across a larger number of jurisdictions. In the absence of scale economies, consolidation may simply reallocate spending across governments rather than reduce the overall size of the public sector.

4 Data and Variables

4.1 Census of Governments

To study the effects of CKH on government fragmentation and spending, this paper draws on data from the Census of Governments (COG) and the Annual Survey of State and Local Government Finances. The U.S. Census Bureau conducts a complete census of the nation's state and local governments every five years (years ending in 2 and 7), collecting information on government finances and tabulating the total number of government entities. In the intervening years, it collects data on a stratified sample, with larger governments appearing in the data every year.

Counting the number of government entities requires defining the units of measurement. Accordingly, the Census defines a government as "an organized entity having governmental character and sufficient discretion in the management of its own affairs to distinguish it as separate from the administrative structure of any other government unit within that state" (U.S. Census Bureau, 2024, page 1). Rather than applying a single litmus test, the Census looks for indicators of these characteristics, including the possession of corporate powers, such as the right to sue, the power to levy taxes and provide services, and the ability to determine a budget without review by other local officials or governments.

In addition to tabulating the number of governments, the Census Bureau classifies local governments into five types: counties, municipalities, townships, school districts, and special districts. While the first three types of governments, all general purpose governments,

do not present significant difficulties for classification, the latter two, comprising special purpose governments, are less easily defined. Special districts in particular require careful delineation due to the variation in special district legislation across states. Accordingly, the Census Bureau defines special districts as "independent, special purpose governmental units that exist as separate entities with substantial administrative and fiscal independence from general-purpose local governments" (U.S. Census Bureau, 2024). This definition may or may not overlap with state-level definitions, where state statutes outline the relevant characteristics under state law. In California, special districts are "any agency of the state for the local performance of governmental or proprietary functions within limited boundaries" and have four distinguishing characteristics: 1) they are a form of government, 2) they have governing boards, 3) they provide services and facilities, and 4) they have defined boundaries (California Senate Local Government Committee, 2010). The key distinction between California's definition and that of the Census Bureau appears to be the degree of independence from general purpose governments (Goodman, 2020).

In addition to tallying the number of governments, the Census and Annual Survey also collect information on revenues, spending, and outstanding debt, making it the only source of nationwide data on local government finances. While the financial information is reported by districts themselves, the Census cleans the data and compares the reported information to financial statements.

For the analysis that follows, I utilize count data on the number of local governments in each state-year and aggregated state-level fiscal outcomes for the period 1977-2022. For census years (years ending in 2 and 7), the fiscal outcomes and counts of governments are tabulated directly from the complete Census of Governments, while for non-census years, they are estimates.[3] Together, this yields 24 annual pre-intervention observations and 22

[3]For fiscal outcomes in non-census years, I rely on the Census Bureau's published state-level aggregates, as compiled by the Tax Policy Center, which provide annual estimates of total local government revenues and expenditures. Because these published estimates are not always available in every year for all subgroups (general-purpose governments and special districts), I allocate the Census totals across these subgroups using their observed shares in adjacent Census of Governments years. For the count of governments, I construct estimates for non-census years by combining unit-level information from the Annual Surveys with the census

post-intervention observations for the main outcomes, covering a 45 year time span. In a robustness test, I also re-estimate my results using only data from the complete censuses.

4.2 Outcomes

This paper's focus is understanding the effects of government consolidation on cost efficiencies. Thus, to understand whether CKH successfully addressed government fragmentation, I first look at the impact that the Act had on the total number of governments in California as well as the total number of special districts, which I scale by land area (number of governments per thousand square miles). To better understand the exact nature of the reform, I further drill down to investigate what types of special districts were most affected. Following this, I examine the effect of the Act on government spending and revenues. Specifically, I focus on the aggregate amount of local government spending per capita as well as the aggregate amount of own-source revenues per capita collected by local governments in the state. I focus on own-source, rather than total revenues, to measure the extent to which local governments draw separately from a common pool of fiscal resources rather than rely on higher-level governments (Berry, 2008).

5 Synthetic Control

To evaluate the impact of CKH, I use the synthetic control method (SCM). Developed as a means of estimating the causal impact of a policy change when there is only one treated unit,

benchmarks. Specifically, I assume that any government that appears in the Annual Survey has existed continuously between its first and last observed appearances. Using this assumption, I first reconstruct raw annual counts by state from unit-level lifespans. I then construct estimates for the non-census years by incorporating a scaled version of the variation implied by the raw counts into a linear interpolation between the census benchmarks.

such as a state, the SCM compares the evolution of outcomes in the treated unit with the evolution of outcomes in a "synthetic control" unit, constructed by minimizing the RMSE of predictor variables among a pool of comparison units (the "donor pool"). In this case, I evaluate the log of the number of local governments in California per thousand square miles - along with associated fiscal metrics, such as log per capita local spending and log per capita own-source revenues – before and after CKH, and I compare these outcomes to those of the synthetic control constructed using a weighted combination of other states. The treatment effect estimates reflect the average difference in outcomes between California and its synthetic control over the post-treatment observations.

Following recent developments in synthetic control methodology (Sun et al., 2025), rather than estimating separate donor weights for each outcome, I estimate a common set of weights by balancing an index of the outcomes. Specifically, I construct an index of the outcomes of interest – log number of local governments per thousand square miles, log spending per capita, log own source revenues per capita – by standardizing each variable using its pre-treatment variation and computing an average. I then use this average to estimate a set of weights for the synthetic control, and I use the same set of weights to estimate treatment effects for all of the outcomes in the paper. This approach leads to a consistent comparison group (synthetic state), exploits the latent factor structure across outcomes, and – relative to estimating outcome-specific weights – reduces bias from imperfect pre-treatment fit when outcomes are correlated (Sun et al., 2025).

5.1 Predictors

Although data-driven, the SCM does require various specification choices on the part of the researchers. First among these is the set of predictor variables. Although it is common to use only the set of pre-intervention outcomes as predictors to reduce specification searching (McClelland and Mucciolo, 2022; Ferman et al., 2020), using all outcome lags does not necessarily

produce the most accurate prediction of posttreatment values and can increase the bound on the bias (Abadie, 2021; Kaul et al., 2022). Thus, to predict the size of the local government sector – as captured by the index of outcomes – I use a combination of pre-treatment outcomes and other variables. Specifically, I use four years of pre-treatment outcomes (1982, 1987, 1992, and 1997), per capita income (averaged over the pre-treatment period), the percentage of families with children (averaged across the 1990 and 2000 censuses), and the level of population (averaged over the pre-treatment period). These covariates capture key structural determinants of local government formation and thus help improve the fit of the synthetic control by accounting for relevant cross-sectional variation; per-capita income accounts for the fact that demand for government services increases with income (Shelton, 2007), the percentage of families with children reflects the fact that K-12 education makes up a high component of local government budgets (Poterba, 1997), and the level of the population captures scale effects in the provision of local public goods. In additional tests, I vary the sets of predictor variables to assess robustness.

5.2 Donor Pool

The SCM also requires the researcher to select the donor pool. In this case, the set of state governments provides a natural comparison set. However, in order to ensure that each unit in the donor pool is a reasonable control for California, I exclude those states that have enacted laws similar to CKH. This excludes the states discussed above in Section 2 – New York, Florida, and Colorado. I also exclude Georgia, which passed a law, the Service Delivery Strategy Act (1997), aimed at minimizing the duplication of services, as well as Oregon, cited by Clarke (2014) for having especially clear merger statutes. In robustness tests, I explore specifications with alternative restrictions on the donor pool.

5.3 Bias-Correction

One recent development in the SCM literature is the use of bias-correction procedures that adjust for differences in the predictors of the treated unit and its synthetic control donors (Abadie and L'hour, 2021; Ben-Michael et al., 2021; Wiltshire, 2022). If the predictor values for the treated unit are not closely reproduced by the synthetic control, then bias corrections can improve on the classical SCM by reducing the estimated bias from predictor discrepancies. Bias-correction is especially relevant when using a fixed set of weights for SCM estimation, as the fit of the synthetic control will necessarily be weaker than it would be under outcome-specific weights (Eliason and Lutz, 2018). Thus, in addition to using a fixed set of weights, I also perform two simple bias-correction procedures based on pre-treatment data. First, I apply a level correction that centers the pre-treatment gap at zero. Second, when a pre-defined diagnostic indicates a meaningful trend in the pre-period, I residualize the gap with respect to a time trend.[4] As in all applications of SCM, the adequacy of the synthetic control and the bias correction can be directly evaluated by inspecting the treatment effect estimates in the pre-period.

5.4 Inference

There is no one agreed upon inference procedure for SCM with a single treated unit; in fact, this continues to be an active area of research (Li, 2020; Cattaneo et al., 2021; Chernozhukov et al., 2021). I follow Sun et al. (2025) by using the conformal inference approach proposed by Chernozhukov et al. (2021). This approach compares post-treatment differences between the treated unit and its synthetic control to the distribution of pre-treatment differences, obtained through time series permutation of the differences. Under the null hypothesis of no

[4]Specifically, using pre-treatment observations, I regress the gap on a linear time trend and apply detrending when the drift – the implied change in the gap over the pre-period under the fitted trend – exceeds a pre-specified threshold (0.10) and the slope is estimated with reasonable precision (t-statistic ≥ 6).

treatment effect, the post-treatment gaps should be statistically indistinguishable from those in the pre-treatment period. Thus, I report p-values for the average treatment effects that are based on the rank of the observed (bias-corrected) post-treatment discrepancy relative to this permutation distribution.

6 Main Results

6.1 Weights

Before presenting the main results, I first examine which units in the donor pool receive positive weights in the synthetic control estimation. Examining the weights helps to ensure that the comparison pool is qualitatively similar to the treated unit with respect to the outcome of interest. As discussed above, rather than using a separate set of weights for each outcome, I use a fixed set of weights for all outcomes constructed using an index of the main outcomes of interest. Table 1 shows the states that receive positive weight in the synthetic control. Consistent with expectations, the synthetic control assigns the greatest weight to states such as Texas and Illinois, which – like California – exhibit relatively fragmentated local government structures and a high degree of fiscal decentralization, with substantial spending and revenue collection occurring at the local level.

6.2 Number of local governments

As a first step, I estimate the effect of CKH on the log of the number of local governments in California per thousand square miles. Figure 3A shows the bias-corrected treatment effect estimates of CKH. Prior to the law passing in the year 2000, the treatment effect

estimates hover close to zero, indicating that there is little deviation between California and its synthetic control in the pre-treatment period. After CKH, the treatment effect estimates exhibit a modest, approximately linear downward trend and remain consistently negative. Table 2 provides the average treatment effect estimates of the intervention over the post-intervention years (2001-2022) along with the associated p-values. On average, CKH reduced the number of local governments by 7 percent, an effect that is significant with a p-value of 0.022.

Figures 3B and 3C show the effect of CKH on the log of the number of general purpose governments and the number special districts respectively (also scaled by square miles). Because special districts are the majority of local governments in the state as well as the primary driver of growth in local governments nationally, any law that incentivizes merging or consolidating districts should act primarily through its effect on special districts. Figure 3B once again shows treatment effect estimates that hover extremely close to zero in the pre-treatment period but that begin to slope slightly downward below zero following CKH in 2000. Figure 3C shows this same general pattern, though there is more deviation in the pre-treatment period from the null of zero treatment effect, and also a greater decline post-treatment. On average, CKH reduced the number of general purpose governments in the state by six percent (Table 2) and the number of special districts by 15 percent. Both effects are statistically significant, with p-values of 0.022. Thus, Figure 3 shows that CKH was effective in reducing the number of local governments in the state relative to its counterfactual trajectory, and that it did so primarily by reducing the number of special districts.

In order to further unpack the results depicted in Figure 3, I also examine the effect of CKH on special districts of different types. I assign special districts in the state to one of two groups: 1) those that perform administrative, financing, or coordinating functions where consolidation might be most likely to produce efficiency gains[5], or 2) all other special

[5]Included in this category are authorities and districts performing functions related to: risk management and insurance, economic development, housing finance, housing and community development, parking, school

districts. In principle, districts that are highly administrative in nature and have overlapping missions across jurisdictions, such as risk management and insurance authorities or housing finance authorities, may be able to reduce duplicative governance without materially affecting service delivery or spending levels. In contrast, other types of districts, especially those with highly local functions and geographically specific assets, such as fire protection districts, may be less likely to benefit from consolidation. Table A1 shows the average treatment effect for the two categories. Perhaps surprisingly, the reform had little impact on administrative districts, and instead primarily impacted other types of districts. The average treatment effect on administrative districts is actually positive, though not statistically significant, meaning that the bulk of the effect fell on districts that may not have been the most likely to benefit from consolidation.

6.3 Fiscal Outcomes

The above results confirm that CKH was successful in reining in fragmentation in California, reducing the number of local governments relative to the state's counterfactual trajectory. Was this reduction in the number of districts accompanied by a reduction in government spending? Figure 4 shows the impact of CKH on the log of total per capita local government spending. As in Figure 3, the treatment effect estimates are shown for all districts as well as general purpose governments and special districts. In Figures 4A, 4B, and 4C, the treatment effect estimates in the pre-period are all close to zero, although there is some volatility, with some of the annual estimates rising to 0.10 log points. In the post-treatment period, in all three figures, the trajectory of the treatment effect estimates effectively remains unchanged. If anything, in Figures 4A and 4C, there is a slight *upward* trend in the estimates at the end of the post-treatment period. Nevertheless, none of the estimates are significant, with p-values of 0.41, 0.22, and 0.11 (Table 2). Thus, there is no indication that CKH lowered

building, transit, transportation, highways, sea and inland port facilities, and air transportation.

total spending among local governments in the state, despite substantially decreasing the number of special districts.

Figure 5 shows the results for the own sources revenues collected by local governments. Own-source revenue may be a more precise way of measuring the effects of CKH as it is uncontaminated by transfers from higher level governments. As in Figure 4, other than the earliest pre-treatment years, the pre-treatment estimate are very close to zero. In the post-treatment period, the estimates are slightly more volatile than the spending results. Once again, however, there is little evidence of an effect. None of the estimates are significant, with p-values of 0.41, 0.24, and 0.13. The absence of a significant decline in spending or revenues suggests that the relative slowdown in government growth did not translate into lower aggregate fiscal activity, with average spending instead increasing among existing jurisdictions.

6.4 Robustness

In this section, I assess the robustness of the above estimates. Specifically, I vary several aspects of the synthetic control design and compare the resulting estimates to my baseline results. First, following the recommendation of Ferman et al. (2020), I include a specification with only pre-treatment outcomes as predictors, a benchmark that avoids specification searching. Next, in order to reduce the potential for over-fitting and interpolation biases, I restrict the donor pool to states with predictor values that are close to those of California (Abadie and Vives-i Bastida, 2022). To accomplish this, given California's size, I exclude states with populations that place them in the bottom ten percent of the distribution. Next, to address the possibility that some states may already make it simple to dissolve special districts that are not serving the public purpose, even if they have not passed a specific law to that effect, I exclude states from the donor pool with clear dissolution procedures.

Specifically, I exclude the states that Bauroth (2010) identifies as having clear dissolution procedures for all types of special districts. Finally, I limit my analysis to years in which a full census of governments was conducted. In each case, I use a common set of weights for all outcomes as in the baseline analysis.

Table 3 presents the results. The baseline estimates are presented alongside the alternative specifications for comparison. The first set of estimates pertaining to the number of governments shows that almost all of the alternative specifications are consistent with the baseline results. 13 of the 15 estimates are negative, and a majority are statistically significant. There is one outlier - a positive and statistically significant finding for the number of special districts when using only pre-treatment outcomes as predictors. However, in this case the pre-treatment estimates show some variability, suggesting that the fixed set of weights may not provide a strong comparison group for the outcome with this particular set of predictors. The second and third sets of estimates, pertaining to per-capita spending and own source revenues, also affirm the baseline results. There is no evidence that CKH decreased per capita spending or own-source revenues. In fact, all of the specifications in columns 1-5 for the fiscal outcomes are positive, indicating that, if anything, the local government sector in California expanded during this period relative to the counterfactual rather than shrank.

As one additional check on the findings, I assess the quality of the synthetic control by evaluating its predictive power in the pre-intervention period. This follows Abadie and Vives-i Bastida (2022)'s suggestion to use out-of-sample prediction as an important validity check. I divide the pre-intervention periods into two halves, using the first three observations to fit the synthetic control, and the remaining three to evaluate the fit. The predictor variables are the same as those outlined above for the main results, except they include only two lags of the outcome variable (in 1982 and 1987) and omit all covariates from after 1988. Figure 6 shows the bias-corrected treatment effect estimates for the three main outcomes of interest – the log of the total number of local governments per thousand square miles, the log of total local spending per capita, and the log of total local own-source revenues per capita. In all

cases, the treatment effect estimates are close to zero in the pre-treatment period, with the exception of own source revenues in the earliest pre-treatment periods. This findings of a zero effect continues into the post-treatment period, correctly estimating a zero effect in the absence of an intervention. The associated p-values are 0.57, 0.48, and 0.39.

7 Difference-in-Difference Analysis

How could the reform rein in the growth of governments but fail to achieve cost savings? If California's number of governments stabilized while fragmentation continued elsewhere, and total local spending did not decline relative to other states, then average spending per government in California must have risen relative to its counterparts. In order to evaluate the change in spending among California local governments that remained post-CKH, I assemble a balanced panel of local governments with a full set of observations in the period of time spanning CKH. Specifically, I assemble a sample of governments that appear every year in the Census' Annual Survey of State and Local Finances between 1994-2022. I use this sample to construct a matched panel of treated (California) and untreated local governments, and I estimate the change in fiscal outcomes using a simple difference-in-difference (d-i-d) design of the following form:

$$Y_{it} = \beta_1 CA_i * Post_t + \delta_i + \gamma_t + \varepsilon_{it} \tag{1}$$

where CA denotes an indicator variable for local governments in California subject to CKH, $Post$ indicates the time period after 2000, and δ_i and γ_t represent government and year fixed effects respectively. The Census does not report population statistics for special districts and school districts and thus the outcome variables cannot be scaled by population as in the earlier fiscal analysis, however in principle this should not be necessary since, unlike with the SCM, there is no concern about assigning negative weights.

To assemble the matched comparison group, I perform nearest neighbor matching by

computing the Mahalanobis distance between each treated unit and all governments in the comparison pool on the basis of the outcome of interest in six pre-intervention years (1994-1999). Because California contained many of the highest spending local governments in the country in 1999, each of which cannot be easily paired with a match, I use a caliper of 0.6 to exclude matches of poor quality, lowering the number of treated districts in the sample from 323 to 254 for the expenditure analysis. As with the synthetic control estimation above, I exclude from the comparison pool governments from New York, Florida, Colorado, Georgia, and Oregon. Because the treated units all come from the same state, the conventional approach of clustering standard errors at the state level does not provide reliable inference. Thus, I follow the approach developed by Ferman and Pinto (2019) and calculate p-values using a heteroskedasticity-adjusted bootstrap that accounts for differences in sampling variability across states.

Figure 7 presents the graphical results of the matching process. Figure 7A shows average log total expenditures among local governments in California (of all types) relative to the matched sample, while Figure 7B shows the results for own source revenues. It is clear from Figure 7A that average total expenditures among local governments in California begins to outstrip the matched sample shortly after the passage of CKH in 2000, with the difference continuing to grow over time. Figure 7B shows similar results for own-source revenues. Columns 1 and 3 in Table 4 show the corresponding d-i-d coefficients, with expenditures and own source revenues increasing by an average of 16 and 9 percent respectively, effects that are significant at the 1 percent level. The results are consistent with a model in which local governments in other states that experienced greater fragmentation shed fiscal responsibilities while the surviving governments in California did not.

To further validate the findings presented in the earlier sections and to isolate the effects of CKH by type of government, I also estimate specifications with an interaction term for special districts:

$$Y_{it} = \beta_1 CA_i * Post_t + \beta_2 CA_i * Post_t * SD_i + \delta_i + \gamma_t + \varepsilon_{it} \quad (2)$$

where SD denotes an indicator variable for special districts. These results are in columns 2 and 4. As expected, the interaction term for special district is positive for both outcomes (though only statistically significant in one case), indicating that the spending increase is especially pronounced among special districts, which saw the greatest reduction in number relative to the counterfactual.

8 Discussion and Conclusion

This paper revisits a long-standing debate about government fragmentation and efficiency. By leveraging a state reform in California, it seeks to determine whether or not consolidation in the number of governments can lead to cost savings at the local level. Unlike most prior studies, which have concentrated on consolidations among general-purpose governments, such as cities or counties, California's Cortese-Knox-Hertzberg (CKH) Government Reorganization Act of 2000 had its largest effect on special districts, the entities most responsible for the growth in the number of local governments in the United States over the last half century. Using synthetic control methods, I show that the reform reduced the number of special districts in the state by an average of 15 percent relative to its counterfactual trajectory over the period 2001-2022. However, this reduction had no impact on overall spending or own-sources revenues collected by local governments in the state. A separate difference-in-difference analysis using a panel of local governments supports this result by showing that spending increased among the surviving governments during this period.

These results stand in contrast to the common presumption – derived largely from studies of amalgamation reforms that targeted general purpose municipalities – that consolidation necessarily delivers cost efficiencies. In the case of California, the evidence suggests

that curbing the growth of special districts has failed to reduce costs. These findings also challenge the view that special districts are prone to inefficiency due to low visibility or limited accountability; other states that experienced greater fragmentation did not experience large increases in spending.

An important limitation of the analysis in this paper is that it focuses on fiscal aggregates and does not directly assess changes in the *quality* of service provision. It is possible that curtailing the growth of local governments in California led to improvements in the quality of public service provision, even if it did not yield spending reductions. Recent international evidence suggests that reforms affecting jurisdictional boundaries can sometimes generate meaningful changes in political incentives and access to public goods even when fiscal outcomes remain unchanged (Tricaud, 2025; Narasimhan and Weaver, 2024; Dahis and Szerman, 2025). This suggests one possible avenue for future work on special districts: assessing whether consolidation alters service quality or responsiveness along margins that are not well captured by aggregate spending measures.

Regardless of potential changes in service quality, an important question remains: Why might it be difficult to realize economies of scale with single function districts? The findings in this paper indicate that consolidation efforts did not necessarily target those districts that might have been the most likely to benefit, i.e. districts with a high share of administrative spending and overlapping responsibilities across districts. Many single function districts offer services that are place-based and capital intensive; as Duncombe and Yinger (1993) note, economies of scale in the production of services such as fire protection flatten out quickly as their core operating costs are tied to geography and response-time standards rather than administrative costs that can be easily trimmed. Thus, the findings in this paper suggest that consolidation among single-function districts will not necessarily generate the kinds of cost savings observed in mergers of general purpose governments or school districts, who can often pool administrative and financial functions across multiple areas. If reform efforts do not successfully target districts where overlapping functions or scale inefficiencies are most

pronounced, they are unlikely to have much fiscal impact.

References

Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature 59*(2), 391–425.

Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American Statistical Association 105*(490), 493–505.

Abadie, A., A. Diamond, and J. Hainmueller (2015). Comparative politics and the synthetic control method. *American Journal of Political Science 59*(2), 495–510.

Abadie, A. and J. Gardeazabal (2003). The economic costs of conflict: A case study of the basque country. *American Economic Review 93*(1), 113–132.

Abadie, A. and J. L'hour (2021). A penalized synthetic control estimator for disaggregated data. *Journal of the American Statistical Association 116*(536), 1817–1834.

Abadie, A. and J. Vives-i Bastida (2022). Synthetic controls in action. *arXiv preprint arXiv:2203.06279*.

Alesina, A., R. Baqir, and C. Hoxby (2004). Political jurisdictions in heterogeneous communities. *Journal of Political Economy 112*(2), 348–396.

Bauroth, N. G. (2010). The strange case of the disappearing special districts: Toward a theory of dissolution. *The American Review of Public Administration 40*(5), 568–592.

Ben-Michael, E., A. Feller, and J. Rothstein (2021). The augmented synthetic control method. *Journal of the American Statistical Association 116*(536), 1789–1803.

Berry, C. (2008). Piling on: Multilevel government and the fiscal common-pool. *American Journal of Political Science 52*(4), 802–820.

Berry, C. R. (2009). *Imperfect union: Representation and taxation in multilevel governments*. Cambridge University Press.

Blesse, S. and T. Baskaran (2016). Do municipal mergers reduce costs? Evidence from a German federal state. *Regional Science and Urban Economics 59*, 54–74.

Blom-Hansen, J., K. Houlberg, and S. Serritzlew (2014). Size, democracy, and the economic costs of running the political system. *American Journal of Political Science 58*(4), 790–803.

Bui, T. and B. Ihrke (2003). It's time to draw the line: A citizen's guide to lafcos, second edition. *California Senate Committee on Local Government*.

California Assembly Committee on Local Government (2024). Guide to the cortese-knox-hertzberg local government reorganization act of 2000. https://calafco.org/images/downloads/Cortese_Knox_Hertzberg_Act/ckh_2024_update.pdf.

California Commission on Local Governance for the 21st Century (2000). *Growth within Bounds: Planning California Governance for the 21st Century*. Diane Publishing.

California Senate Local Government Committee (2010). What's so special about special districts? A citizen's guide to special districts in california, fourth edition.

Carr, J. B. and R. C. Feiock (2004). *Perspectives on city-county consolidation and its alternatives*. Routledge.

Cassidy, T. and T. Velayudhan (2026, 01). Government fragmentation and economic growth. *The Review of Economics and Statistics*, 1–18.

Cattaneo, M. D., Y. Feng, and R. Titiunik (2021). Prediction intervals for synthetic control methods. *Journal of the American Statistical Association 116*(536), 1865–1880.

Chernozhukov, V., K. Wüthrich, and Y. Zhu (2021). An exact and robust conformal inference method for counterfactual and synthetic controls. *Journal of the American Statistical Association 116*(536), 1849–1864.

Clarke, C. (2014). Merging and dissolving special districts. *Yale J. on Reg. 31*, 493.

Dahis, R. and C. Szerman (2025). Decentralizing development: The economic impacts of goverment splits. *American Economic Journal: Applied Economics*.

Dolan, D. A. (1990). Local government fragmentation: Does it drive up the cost of government? *Urban Affairs Quarterly 26*(1), 28–45.

Duncombe, W. and J. Yinger (1993). An analysis of returns to scale in public production, with an application to fire protection. *Journal of Public Economics 52*(1), 49–72.

Duncombe, W. and J. Yinger (2007). Does school district consolidation cut costs? *Education Finance and Policy 2*(4), 341–375.

Eliason, P. and B. Lutz (2018). Can fiscal rules constrain the size of government? an analysis of the "crown jewel" of tax and expenditure limitations. *Journal of Public Economics 166*, 115–144.

Faulk, D. and G. Grassmueck (2012). City-county consolidation and local government expenditures. *State and Local Government Review 44*(3), 196–205.

Ferman, B. and C. Pinto (2019). Inference in differences-in-differences with few treated groups and heteroskedasticity. *Review of Economics and Statistics 101*(3), 452–467.

Ferman, B., C. Pinto, and V. Possebom (2020). Cherry picking with synthetic controls. *Journal of Policy Analysis and Management 39*(2), 510–532.

Foster, K. A. (1997). *The political economy of special-purpose government*. Georgetown University Press.

Goodman, C. (2020). Patterns in special district creation and dissolution.

Goodman, C. B. (2015). Local government fragmentation and the local public sector: A panel data analysis. *Public Finance Review 43*(1), 82–107.

Goodman, C. B. (2019). Local government fragmentation: What do we know? *State and Local Government Review 51*(2), 134–144.

Grossman, G. and J. I. Lewis (2014). Administrative unit proliferation. *American Political Science Review 108*(1), 196–217.

Grossman, G., J. H. Pierskalla, and E. Boswell Dean (2017). Government fragmentation and public goods provision. *Journal of Politics 79*(3), 823–840.

Hendrick, R. and Y. Shi (2015). Macro-level determinants of local government interaction: How metropolitan regions in the united states compare. *Urban Affairs Review 51*(3), 414–438.

Hendrick, R. M., B. S. Jimenez, and K. Lal (2011). Does local government fragmentation reduce local spending? *Urban Affairs Review 47*(4), 467–510.

Jeffries, D. W. (1971). Lafco: Is it in control of special districts. *Hastings LJ 23*, 913.

Jimenez, B. S. and R. Hendrick (2010). Is government consolidation the answer? *State and Local Government Review 42*(3), 258–270.

Kaul, A., S. Klößner, G. Pfeifer, and M. Schieler (2022). Standard synthetic control methods: The case of using all preintervention outcomes together with covariates. *Journal of Business & Economic Statistics 40*(3), 1362–1376.

Lewis, B. D. (2017). Does local government proliferation improve public service delivery? evidence from indonesia. *Journal of Urban Affairs 39*(8), 1047–1065.

Li, K. T. (2020). Statistical inference for average treatment effects estimated by synthetic control methods. *Journal of the American Statistical Association 115*(532), 2068–2083.

McCasland, J., T. Monarrez, D. Schonholzer, and E. Stamm (2023). Shattered metropolis: The great migration and the fragmentation of political jurisdictions. *Working paper*.

McClelland, R. and L. Mucciolo (2022). An update on the synthetic control method as a tool to understand state policy. https://www.taxpolicycenter.org/sites/default/files/publication/163919/an_update_on_the_synthetic_control_method_as_a_tool_to_understand_state_policy.pdf. Urban Institute Research Report.

Monarrez, T. and D. Schönholzer (2023). Dividing lines: racial segregation across local government boundaries. *Journal of Economic Literature 61*(3), 863–887.

Narasimhan, V. and J. Weaver (2024). Polity size and local government performance: evidence from india. *American Economic Review 114*(11), 3385–3426.

Oates, W. E. (1972). *Fiscal Federalism*. New York: Harcourt Brace Jovanovich.

Oates, W. E. (1999). An essay on fiscal federalism. *Journal of Economic Perspectives 13*(2), 112–128.

Ostrom, V., C. M. Tiebout, and R. Warren (1961). The organization of government in metropolitan areas: a theoretical inquiry. *American Political Science Review 55*(4), 831–842.

Poterba, J. M. (1997). Demographic structure and the political economy of public education. *Journal of Policy Analysis and Management 16*(1), 48–66.

Reingewertz, Y. (2012). Do municipal amalgamations work? Evidence from municipalities in Israel. *Journal of Urban Economics 72*(2-3), 240–251.

Shelton, C. A. (2007). The size and composition of government expenditure. *Journal of Public Economics 91*(11-12), 2230–2260.

Smaldone, A. and M. L. Wright (2024). Local governments in the us: A breakdown by number and type. *The Regional Economist*.

Sun, L., E. Ben-Michael, and A. Feller (2025). Using multiple outcomes to improve the synthetic control method. *Review of Economics and Statistics*, 1–29.

Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy 64*(5), 416–424.

Treisman, D. (2007). *The architecture of government: Rethinking political decentralization*. Cambridge University Press.

Tricaud, C. (2025). Better alone? evidence on the costs of intermunicipal cooperation. *American Economic Journal: Applied Economics 17*(1), 160–207.

U.S. Census Bureau (2024). Individual state descriptions: 2022.

Wiltshire, J. C. (2022). allsynth:(stacked) synthetic control bias-correction utilities for stata.

Zax, J. S. (1989). Is there a leviathan in your neighborhood? *American Economic Review 79*(3), 560–567.

Figure 1: Number of Local Governments in the United States

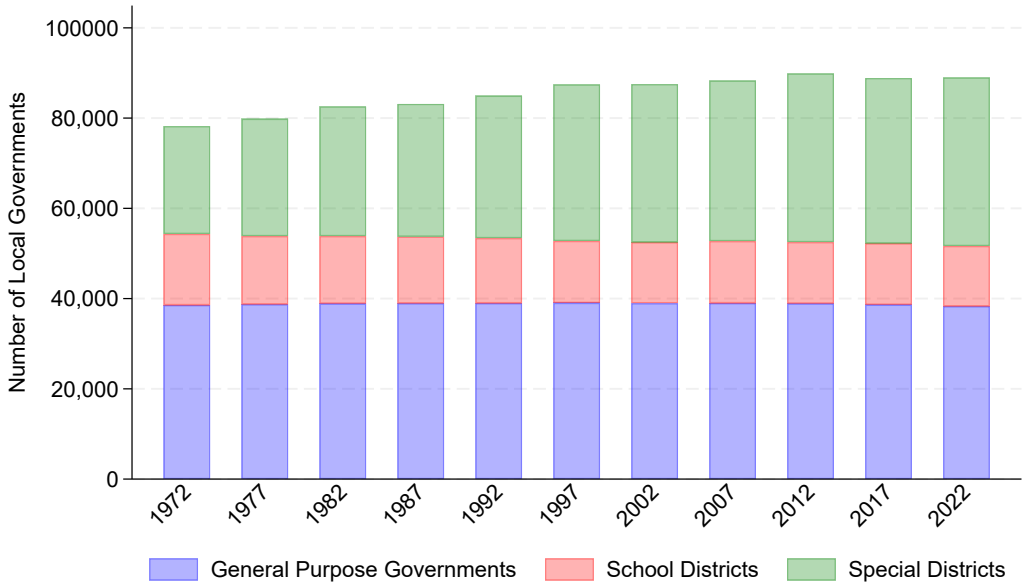

Source: Census of Governments and author's calculations.

Figure 2: Number of Governments in California Pre- and Post-CKH

Figure 2A: Total Number of Local Governments Per Thousand Square Miles

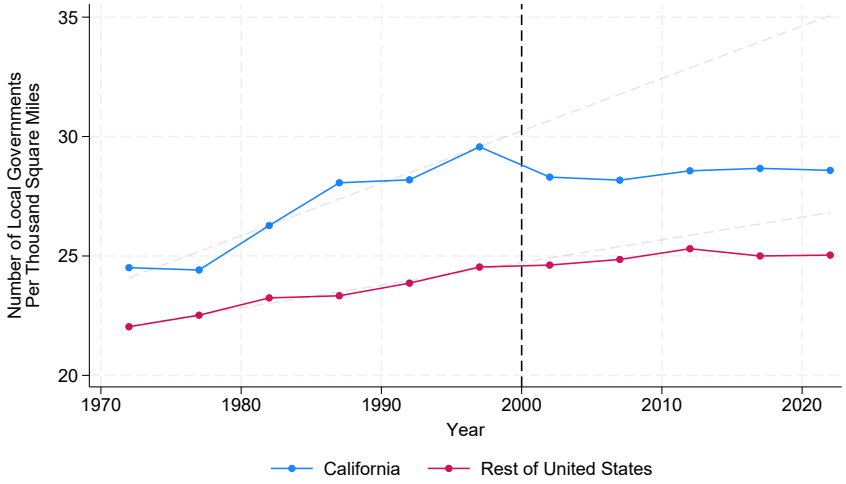

Figure 2B: Number of Special Districts Per Thousand Square Miles

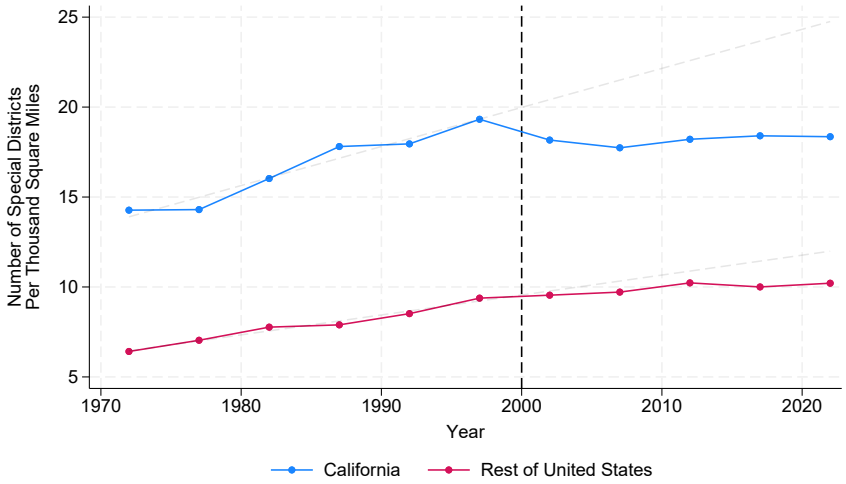

Source: Census of Governments and author's calculations.

Figure 3: Treatment Effects - Number of Governments

Figure 3A: Log Total Number of Governments Per Thousand Square Miles

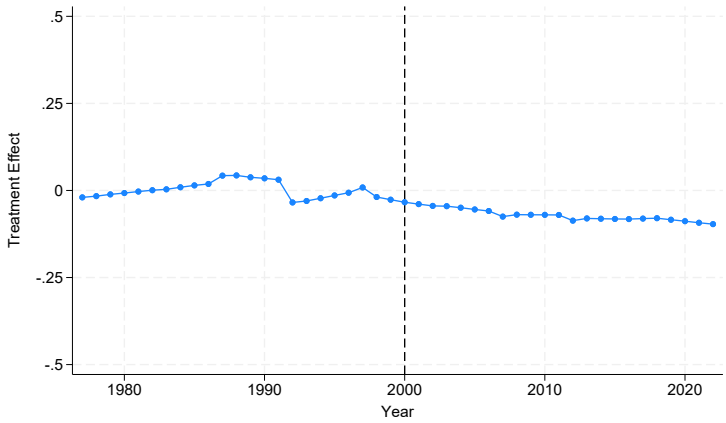

Figure 3B: Log Number of General Purpose Governments Per Thousand Square Miles

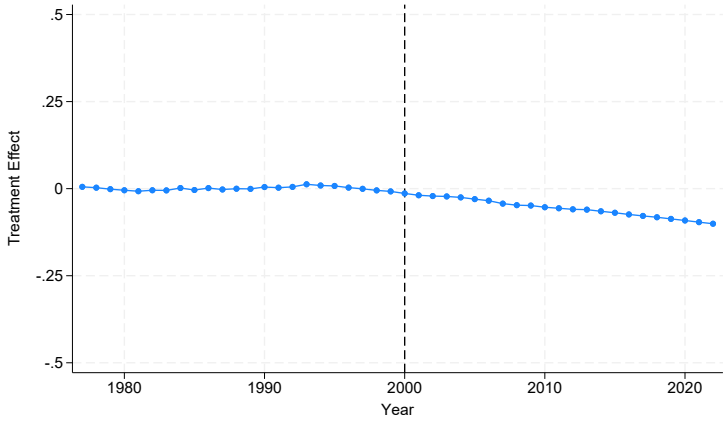

Figure 3C: Log Number of Special Districts Per Thousand Square Miles

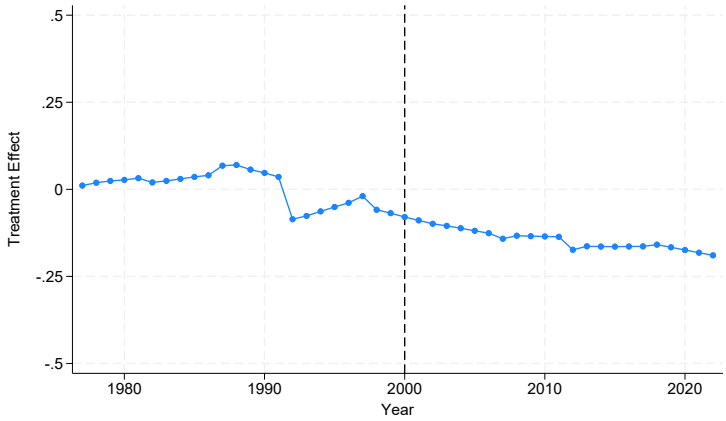

Note: Figure 3A plots the synthetic control treatment effect estimates for the effect of CKH on the logged total number of governments per thousand square miles. The plot shows the bias-corrected gap between California and the synthethic control before and after the law was passed in 2000. Figures 3B and 3C plot the treatment effect estimates for the logged number of general purpose governments and special districts respectively.

Figure 4: Treatment Effects - Log Per Capita Local Government Spending

Figure 4A: All Governments

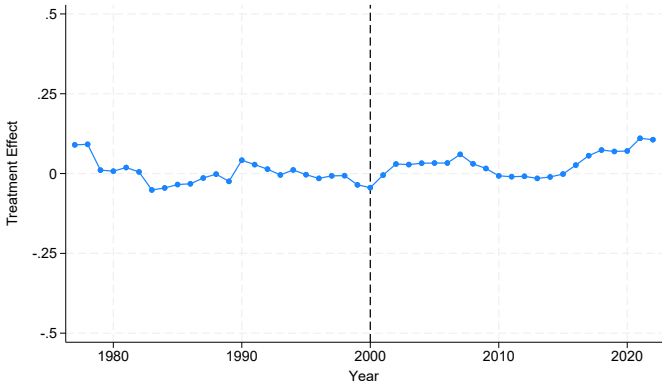

Figure 4B: General Purpose Governments

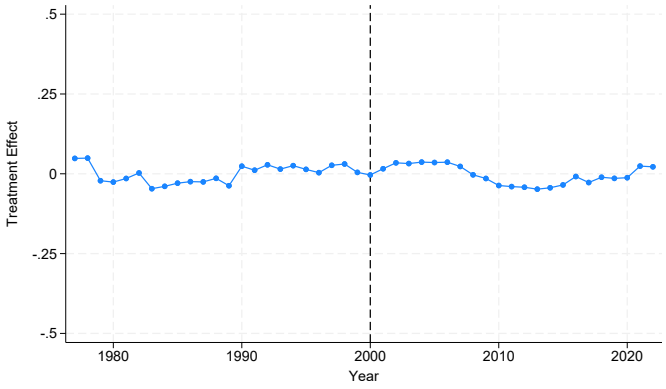

Figure 4C: Special Districts

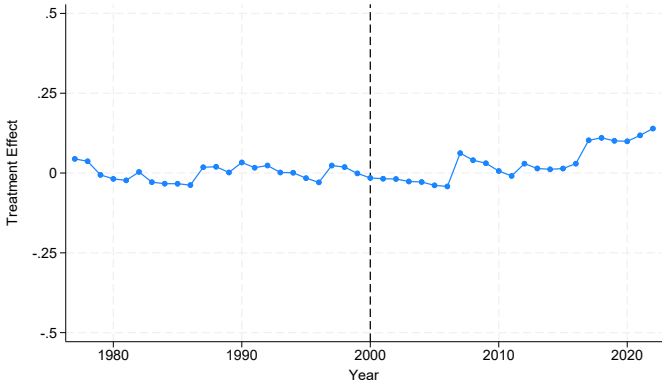

Note: Figure 4A plots the synthetic control treatment effect estimates for the effect of CKH on the log per capita spending by local governments in California. The plot shows the bias-corrected gap between California and the synthetic control before and after the law was passed in 2000. Figures 4 B and C plot the treatment effect estimates for the effect of CKH on the log of per capita spending by general purpose governments and special districts respectively.

Figure 5: Treatment Effects - Log Per Capita Own-Source Local Government Revenue

Figure 5A: All Governments

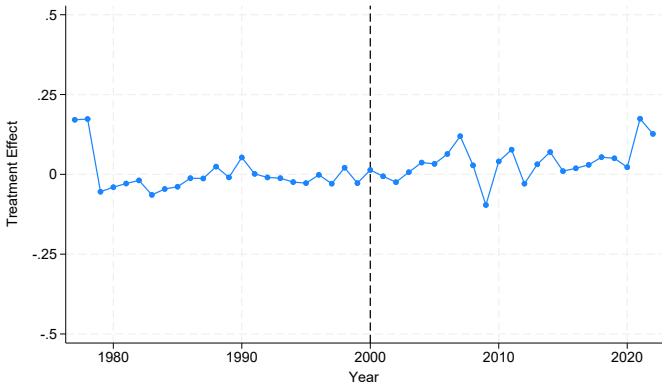

Figure 5B: General Purpose Governments

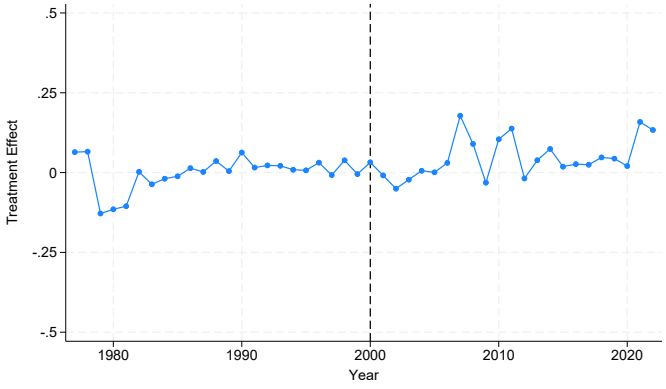

Figure 5C: Special Districts

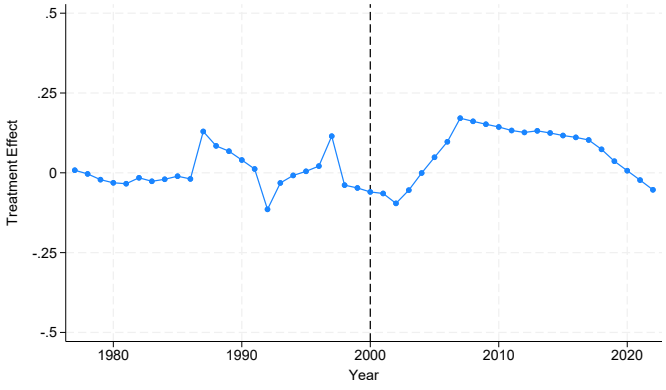

Note: Figure 5A plots the synthetic control treatment effect estimates for the effect of CKH on the log per capita own source revenues collected by local governments in California. The plot shows the bias-corrected gap between California and the synthetic control before and after the law was passed in 2000. Figures 5 B and C plot the treatment effect estimates for the effect of CKH on the log per capita own sources revenues collected by general purpose governments and special districts respectively.

Figure 6: Placebo in Time

Figure 6A: Log Number of Governments Per Thousands Square Mile

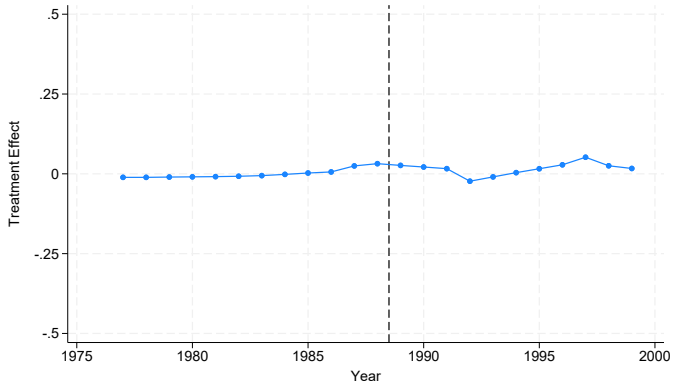

Figure 6B: Log Per Capita Local Spending

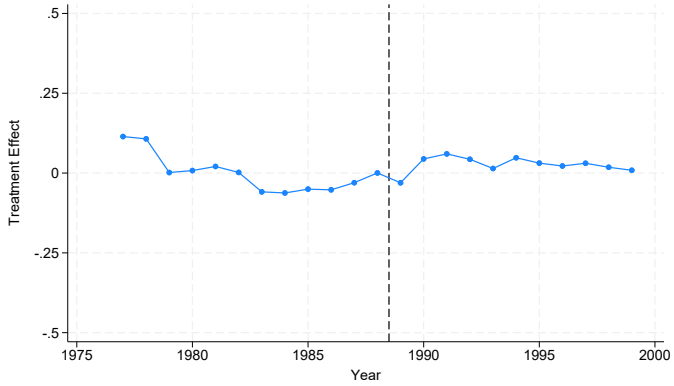

Figure 6C: Log Per Capita Own-Source Revenues

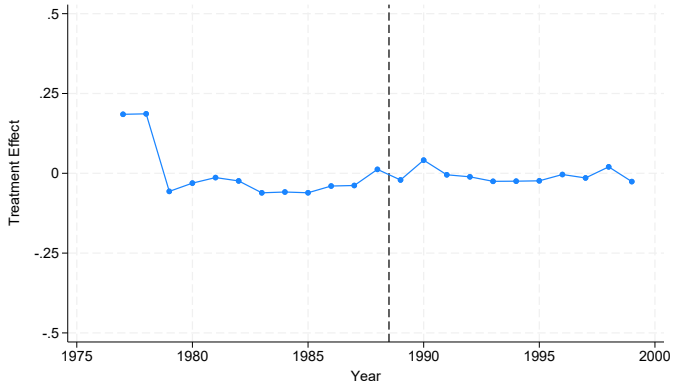

Note: Figure 6A plots the synthetic control treatment effect estimates for the effect of a placebo on the logged total number of governments. The plot shows the bias-corrected gap between California and the synthetic control before and after the placebo intervention in 1988. Figures 6B and C plot the treatment effect estimates for the effect of the placebo on log per capita local government spending and log per capita own source revenues respectively.

Figure 7: Fiscal Outcomes in a Matched Panel of Local Governments

Figure 7A: Log Total Expenditures

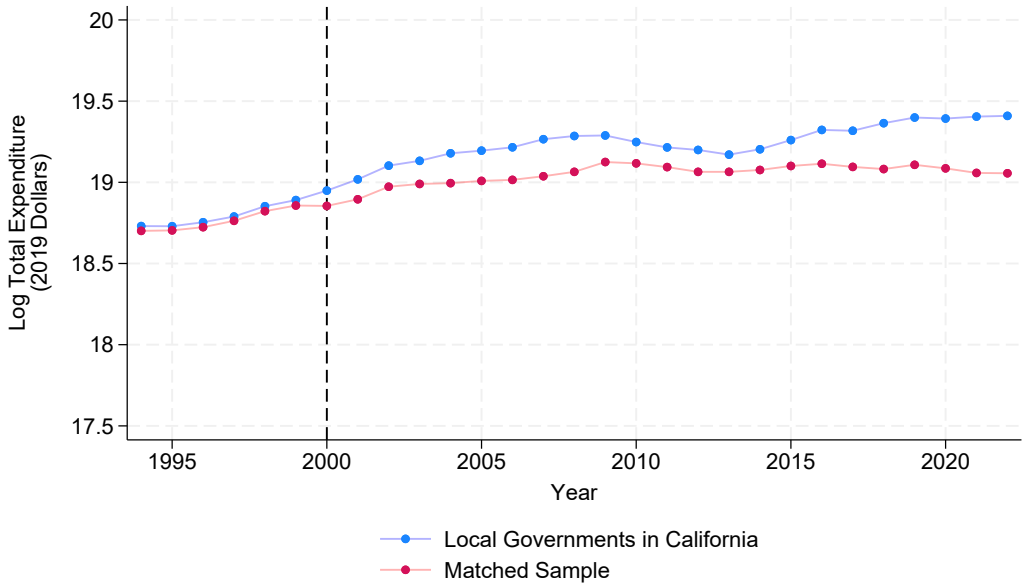

Figure 7B: Log Own-Source Revenues

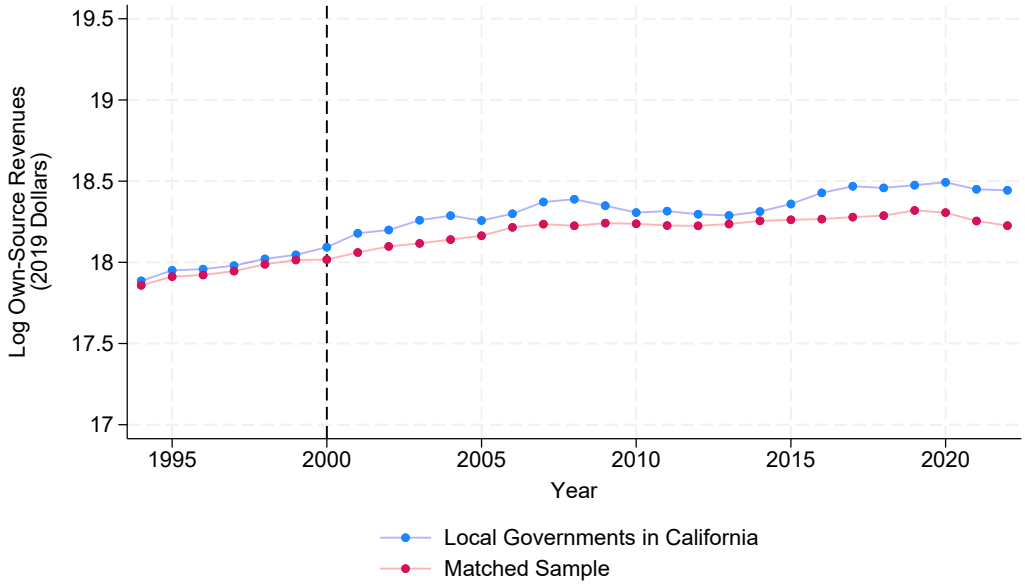

Note: Figure 7 A plots the mean amount of spending by local governments in California relative to a sample of matched governments before and after the passage of the Cortese-Knox-Hertzberg (CKH) Local Government Reorganization Act of 2000. The sample is limited to governments that appear in the Annual Survey of State and Local Finances every year over the sample period. The matched comparison group is constructed using nearest neighbor matching on the Mahalanobis distance of the pre-intervention outcomes. The final matched sample excludes matches with a distance greater than 0.6. Figure 7B plots the corresponding means for own-source revenues.

Table 1: Synthetic Control Weights

State	Weight	State	Weight
Alabama	0	Nebraska	0
Alaska	0	Nevada	0
Arizona	0	New Hampshire	0
Arkansas	0	New Jersey	0.017
Connecticut	0.247	New Mexico	0
Delaware	0	North Carolina	0
Hawaii	0	North Dakota	0
Idaho	0	Ohio	0
Illinois	0.314	Oklahoma	0
Indiana	0	Pennsylvania	0
Iowa	0	Rhode Island	0
Kansas	0	South Carolina	0
Kentucky	0	Douth Dakota	0
Louisiana	0	Tennessee	0
Maine	0	Texas	0.422
Maryland	0	Utah	0
Massachusetts	0	Vermont	0
Michigan	0	Virginia	0
Minnesota	0	Washington	0
Mississippi	0	West Virginia	0
Missouri	0	Wisconsin	0
Montana	0	Wyoming	0

Note: This table lists the states in the donor pool along with the weights used in the synthetic control. The donor pool excludes several states based on the discussion in section 5.2.

Table 2: Synthetic Control Estimates

		Average Post-Treatment Effect		
		All Governments	General Purpose Governments	Special Districts
		(1)	(2)	(3)
Log Number of Governments Per Thousand Square Miles	Estimate	-0.07	-0.06	-0.15
	P-value	0.022	0.022	0.022
Log Per Capita Total Local Government Spending	Estimate	0.03	-0.00	0.03
	P-value	0.41	0.22	0.11
Log Per Capita Total Local Own-Source Revenue	Estimate	0.04	0.05	0.00
	P-value	0.41	0.24	0.13

Note: The table presents synthetic control estimates of the average effect of CKH over the post-treatment period (2001-2022). The p-values are based on the inference procedure described in Section 5.4.

Table 3: Alternative Specifications

		Average Post-Treatment Effect														
		All Governments					General Purpose Governments					Special Districts				
		Baseline	Lags Only	Smaller DP	Clear Proc	Census Only	Baseline	Lags Only	Smaller DP	Clear Proc	Census Only	Baseline	Lags Only	Smaller DP	Clear Proc	Census Only
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Log # of Governments	Estimate	-0.07	-0.09	-0.08	-0.03	-0.07	-0.06	-0.08	-0.06	0.04	-0.07	-0.15	0.03	-0.15	-0.11	-0.21
per Thousand Sq Miles	P-value	0.022	0.15	0.022	0.26	0.10	0.022	0.022	0.022	0.15	0.10	0.022	0.043	0.022	0.022	0.10
Log Per Capita Total	Estimate	0.03	0.02	0.04	0.03	0.06	-0.00	0.02	-0.00	0.00	0.01	0.03	-0.06	0.02	0.07	0.06
Local Gov Spending	P-value	0.41	0.26	0.41	0.46	0.20	0.22	0.24	0.20	0.83	0.20	0.11	0.13	0.13	0.15	0.10
Log Per Capita Total	Estimate	0.04	0.15	0.04	0.05	0.06	0.05	0.03	0.04	0.07	0.04	0.00	-0.10	-0.00	0.15	0.16
Own Source Revenue	P-value	0.41	0.09	0.39	0.39	0.50	0.24	0.17	0.24	0.24	0.10	0.13	0.15	0.13	0.022	0.10

Note: The table presents the baseline synthetic control estimates for the effect of CKH alongside the results of four different robustness checks. The first, "Lags Only" includes all pre-treatment outcomes as the only predictors. The second, "Smaller DP" (donor pool), removes states in the bottom 10 percent of the population distribution (based on 1997 population numbers). The third, "Clear Proc," removes states from the donor pool that have clear dissolution procedures for special districts, based on Bauroth (2010). The fourth, "Census Only" limits the sample to years in which a complete census of governments was conducted (years ending in 2 and 7).

Table 4: Difference-in-Differences

		Log Total Expenditures		Log Own-Source Revenues	
		(1)	(2)	(3)	(4)
CA*Post	Coefficient	0.16***	0.15***	0.09***	0.07***
	P-Value	0.001	0.001	0.001	0.001
CA*Post*Special_District	Coefficient		0.03		0.07***
	P-value		0.63		0.001
Year Fixed Effects		Yes	Yes	Yes	Yes
Gov Fixed Effects		Yes	Yes	Yes	Yes
Clusters		36	36	39	39
Observations		13,427	13,427	10,728	10,728

Note: *** $p < 0.01$. This table shows the results of difference-in-difference estimation using the panel of governments depicted in Figure 7. Financial variables are in 2019 dollars. P-values calculated using the method described in Ferman and Pinto (2019).

9 Appendices

Table A1: Synthetic Control Estimates of the Number of Special Districts by Type

		Average Treatment Effect
Log Number of Administrative Districts	Estimate	0.07
	P-value	0.13
Log Number of Non-Administrative Districts	Estimate	-0.22
	P-value	0.022

Note: The table presents synthetic control estimates of the average effect of CKH on the number of special districts in California by type over the period 2001-2022 . "Administrative Districts" are those whose functions primarily involve administrative coordination, financing, or system-level infrastructure management and include risk management and insurance, economic development authorities, housing finance authorities, housing and community development, parking authorities, school building authorities, transit authorities, transportation authorities, highway authorities, sea and inland port facilities, and air transportation authorities. Non-administrative districts include all other special districts. The p-values are based on the permutation test described in Section 5.4.