Three Essays in the Public Finance of Education

by

Andrew Litten

**A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Public Policy & Economics)
in the University of Michigan
2017**

Doctoral Committee:

**Professor Susan M. Dynarski, Co-Chair
Professor James R. Hines Jr, Co-Chair
Professor Brian A. Jacob
Professor Michael F. Lovenheim**

Andrew Litten

alitten@umich.edu

ORCID iD: 0000-0001-5284-4852

ACKNOWLEDGEMENTS

I am grateful to the consistent and thoughtful help of Sue Dynarski, Jim Hines, Brian Jacob and Mike Lovenheim. This paper also benefited from discussions with John DiNardo, Jeff Smith, Kevin Stange, John Bound, Kim Rueben, and Doug Harris. This work would not have been possible without the help of my research assistant, Michael Doa. Thanks to feedback from Julian Hsu, Chris Sullivan, Sarah Johnston, Paul Brehm, Bryan Stuart, Jacob Bastian, Josh Hyman, and every conference and seminar attendee at the University of Michigan and the Assocation for Public Policy Analysis and Management annual conference. Finally, thanks to Kerwin Charles and Dan Black.

TABLE OF CONTENTS

LIST OF TABLES

Table

LIST OF FIGURES

Figure

LIST OF APPENDICES

Appendix

ABSTRACT

My dissertation addresses a variety of fiscal issues relating to the education sector in the United States. The funding of education in the Unities States is a complex topic which relies of a variety of programs at the Federal, state, and local level. This dissertation moves the literature forward by evaluating how these relationships respond to major policy changes.

My first chapter seeks to identify the effect that public sector unions have on compensation. Specifically, I look at the compensation premium associated with teachers' unions in Wisconsin. In 2011, Wisconsin passed a landmark law (Act 10) which significantly lowered the bargaining power of all public sector unions in the state. Using an event study framework, I exploit plausibly exogenous timing differences based on contract renewal dates, which caused districts to be first exposed to the new regulations in different years. I find that the reduction in union power associated with Act 10 reduced total teacher compensation by 8%, or \$6,500. Roughly two-thirds of this decline is driven through reduced fringe benefits. Subgroup analysis shows that the most experienced and highest paid teachers benefit most from unionization. I supplement the event study approach with synthetic control methods to find that regulatory limits on contract terms, rather than other mechanisms such as state financial aid cuts or union decertification, are driving the results.

My second chapter estimates the short-run flypaper effect in the education sector. From 2009-2012, the Federal government increased aid for K-12 education by a factor of 1.5. This expansion was channeled in large part through three formula-based programs - Title I, Special Education, and the State Fiscal Stabilization Fund. I exploit non-linearities in the Federal grant funding mechanisms to estimate the effects of these grants on local district spending and employment outcomes. I find that each \$108,000 in grants lead to one additional hire

per year. I also find that fiscal response to grants is sensitive to the nature of the program. Grants related to the temporary Stabilization Fund were nearly completely crowded out, while grants related to expanding permanent programs such as Title I or Special Education increased total spending nearly dollar to dollar.

The final chapter (with Morris Hamilton) seeks to identify the causal relationship between increased state Medicaid obligations and higher education spending. After several decades of Federal mandates and high rates of health cost inflation, Medicaid spending has taken an increasingly larger share of state budgets, forcing states to make offsetting cuts elsewhere. We argue that state governments are likely to cut higher education in response to these changes, as institutions of higher education have the capacity to find additional revenues elsewhere. We use Federally administered Supplemental Security Income (SSI) enrollments to instrument for state Medicaid spending. We find that a $1.00 increase in Medicaid costs leads to a decrease in higher education subsidies of 20 cents to 37 cents. Our approach provides estimates which are both more credible and more precise than those which have previously been used in the literature.

CHAPTER I

The Effects of Public Unions on Compensation: Evidence from Wisconsin

1.1 Introduction

This paper seeks to identify the compensation[1] premium associated with public sector unions. Specifically, I look at how teachers' unions in Wisconsin affect teacher compensation. In 2011, Wisconsin passed a landmark law (Act 10) which significantly lowered the bargaining power of all public sector unions in the state. I exploit plausibly exogenous timing differences based on contract renewal dates, which caused districts to be first exposed to the new regulations in different years. Using an event study framework, I exploit these timing differences to estimate the causal effect that limiting public union power had on teacher compensation.

The size (or even the existence) of the union wage premium is an empirical question. The cannonical model of rent-seeking suggests that labor unions organize in an effort to use market power to shift management choices towards member preferences for higher worker pay. There remains, however, a variety of reasons why unions might organize but still fail to increase wages. One possibility, described in Freeman and Kleiner (1990), is that union members care more about other types of working conditions than wages. In the context of

[1] "Compensation" throughout the paper will refer to the sum of salaries and fringe benefits.

teachers' unions, it is possible that unions act as an offset to regulatory capture from adminstrators, or are more motivated by student performance. In general, the broad "Voice" model described in Freeman and Medoff (1984) can describe unions which simply have different information than management. In general, a large union compensation premium is more likely to be associated with the rent-seeking model. Even this, however, depends not just on the unions' preferences, but also their institutional power.

The union wage premium is not simply an empirical question; it is an open empirical question. A number of well-identified, quasi-experimental studies have shown that unions have no causal effect on wages. In one study, DiNardo and Lee (2004), which follows newly certified private sector unions following a close election, the authors find no statistically significant union wage effects, and reject pay increases as large as 3%. Another study, Lovenheim (2009), is similar in spirit to DiNardo and Lee, and looks at the creation of new teachers' unions following the passage of union-friendly rules in three Midwestern states.[2] Lovenheim cannot identify statistically significant effects on wages either in the short or in the long run. Alternatively, Lewis (1990) shows a broad survey of 75 studies, and concludes the public union wage gap is 8-12%. Given this range of prior estimates, it is important to consider how the union compensation premium may be sensitive to institutional context.

This open question about the existence and magnitude of the union compensation premium has particular importance in the public sector. The public sector is more heavily unionized and more active in terms of policy than the private sector. As shown by Freeman (1988), private sector unionization has been falling and public sector unionization has been growing within the United States since the late 50s. As of 2016, unions represent 35% of workers in the public sector, relative to just 6.7% of workers in the private sector[3]. In addition, while private sector unions face a single national regulatory system, governed by the National Labor Review Board, public sector unions are governed by at the state level. Taxpayers and government watchdogs have had an interest in more closely regulating pub-

[2] Iowa, Minnesota, and Indiana.
[3] Hirsch and Macpherson (2016)

lic unions in pursuit of potential savings for some time[4], and state and local budget crises following the Great Recession have only served to increase this interest.

While Wisconsin is one of the most dramatic[5] examples of a renewed regulatory focus on public sector unions, it should be considered an important part of an ongoing trend, rather than a one-off event. Increased state activity in regulating state and local unionas been building over the last several years. One example of this is the charter school movement, which is partially motivated by the ability of charter schools to make hiring and compensation decisions unconstrained by teachers' unions[6]. Other examples include New Mexico, which had its teachers' bargaining rights lapse and subsequently restored between 1999-2003. Indiana restricted state union bargaining by executive order in 2005, and passed a set of reforms very similar to those in Wisconsin in 2012. Ohio signed a similar set of regulations into law, but the provisions were defeated by a last-minute ballot initiative before going into effect. Following Wisconsin's Act 10 and Ohio's referendum, other states passed more limited "right to work" legislation, including Indiana (2012), Michigan (2013) and West Virginia (2016).

Looking within public unions, teachers' unions are a useful sub-group to focus on. First, teachers' are widely employed, and needed in every state and local region of the country. Second, teachers' are highly unionized. CPS data shows teachers are almost 50% unionized nationally as of 2015. Third, the political and financial structure of school districts make teachers a convenient group to study in order to understand some of the secondary effects of the wage premium. The single-purpose of school districts makes it transparent how this price change causes substitution across other goods.

Prior efforts to investigate the teacher wage premium have faced three main challenges: accurate measurement of union status, exogenous variation in union status, and quality school financial data. The first, accurate union measurement, is important because many

[4] A useful discussion of this concern can be found in Courant, Gramlich and Rubinfeld (1979).

[5] Referencing not only the scope of regulatory change but also the political upheaval it caused. The exact circumstances surrounding the Wisconsin regulatory changes will be discussed in greater detail in section 1.2.

[6] See Abdulkadiroğlu et al. (2011) and Wilson (2009) for a discussion of how charter schools use this flexibility to their advantage.

studies have relied on the self-reported Current Population Survey, or the Survey of Governments to address this question. In one the first papers to use state regulatory changes to instrument for union status, Freeman and Valletta (1988), the authors used CPS data to look at union effects at the individual level. In a later comparison between self-reported union status and employer-reported union status, Card (1996) showed significant measurement problems with the CPS. If the true rate of transition between union and non-union status is low, as we assume would be the case for teachers, even small changes in measurement can lead to significant biases in a panel setting. Perhaps even more concerning, in his efforts to replicate the results of Hoxby (1996), Lovenheim (2009) showed that the Survey of Governments likely has non-classical measurement error. This challenge is particularly pronounced for teachers' unions, as many teachers operate in less formal trade organizations that established so-called"Meet and Confer" agreements (see Freeman and Han (2013)). Because these agreements occupy a middle ground between a non-union contract and a formal collective bargaining agreement, they may be less consistently reported. Thus, accurate measurement of union status is a critical part of any empirical study.

The second challenge is identifying exogenous variation in union status. There are a variety of reasons one might think there is endogenous self-selection into union formation. Unions might be more likely to form when management is particularly bad, or where the workers are treated particularly poorly. Two commonly used techniques to address this problem are regression discontinuity design and state law instruments. The regression discontinuity approach has been used in a variety of studies relating to unions in the public sector (examples include DiNardo and Lee (2004); Lee and Mas (2009); Frandsen (2012) and Sojourner et al. (2015)). This approach has been less common in the public sector, as elections are both less frequent and the public sector lacks a central, national organizing agency such as National Labor Relations Board. Instead, elections are managed by state boards, and tend to follow in bulk shortly after a major rule change. As a result, the tradition in studying public unions has been to use variation driven by changes in state regulations. This approach was first pioneered in Freeman and Valletta (1988) but has been used by many others such as Hoxby (1996) and Frandsen (2016).

The third major challenge is the collection of employment and district finance data. As teachers' unions grew primarily out of regulatory changes in the 1960s and 1970s, efforts to study the causal effects of these unions have been limited to the data available in this time frame. These data include the Survey of Governments, which is an incomplete sampling of school districts surrounding this time frame. Key studies such as Hoxby (1996) and Lovenheim (2009) have not been able to look at important outcomes such as teacher fringe benefits, when both historical surveys (Lewis (1990)) and more recent efforts (Frandsen (2016)) have shown this to be important.[7] Because this is a more recent policy change, I am able to take advantage of modern administrative data in Wisconsin and track the universe of school districts for the full period, while including key outcomes such as fringe benefits. Recent work Biasi (2017) has taken advantage of sorting which was driven by changes in teacher pay schedules, also following Wisconsin's Act 10.

Of course, in practice, unions could affect many things in addition to teacher compensation. Unions might also change the total factor productivity of various inputs, lobby to increase total spending on public goods, or exert preferences over other items in the production function.

I focus on the union compensation premium for three key reasons. First, while the role unions play on student outcomes is a crucial policy question, I do not expect to be able to offer a satisfying answer in the time frame I observe here. Student performance is cumulative and the cohorts which will be fully exposed to the post-union world are still in school. Those cohorts with currently measurable outcomes have only experienced treatment for a small fraction of their overall years of schooling. Second, school districts in Wisconsin

[7] Fringe benefits are conceptually, as well as empirically, important. First, fringe benefits are untaxed, so increasing these is an indirect way of increasing total income. Second, these are less transparent to the public than salary. Unions may wish to use fringe benefits to increase pay in a way which is less likely to meet with disapproval from the public. In both Lewis's survey and the Frandsen paper cited, large compensation effects are driven through fringe benefits rather than wages.

are institutionally constrained in their ability to increase total revenues. This makes lobbying to shift the demand for public goods an unlikely mechanism through which Wisconsin unions operate (I will discuss these limits - as well as the Wisconsin school finance system as a whole - in greater detail in section 1.2). Finally, while unions may have preferences in changing the distribution of inputs used on education (for example, teaching aids, textbooks, computers, etc.), compensation is a clear first-order effect. 40 percent[8] of school budgets are spent of teacher compensation. It would thus be difficult to measure the differences between union preferences over other inputs, versus simple substitution and income effects resulting from the lower price of the main input (teaching).

Thus, in this paper I contribute to the literature's understanding of the public union compensation premium. I do so in a modern context, looking at policy changes which are not only distinct from those which have occurred historically, but more similar to those most likely to be considered going forward. I address the key three challenges in the public union literature - measurement, endogeneity, and data quality - by using hand-collected contract data, plausibly exogenous contract termination dates, and rich administrative records on school operations and teacher-level pay. I find that the reduction in union borrowing power in Wisconsin is associated with an 8% reduction in total compensation, and that this effect is primarily driven by reductions in pay at the upper-end of the income distribution.

The rest of this section proceeds as follows. In section 1.2, I discuss institutional backgrounds related to teachers' unions at large, the legal changes in Wisconsin, and the Wisconsin public school system. In section 1.3, I describe the collective bargaining agreement data set and Wisconsin administrative data. In section 1.4, I discuss the contract timing event study. In section 1.5, I extend the analysis to heterogeneity in treatment by predicted teacher wage. In section 1.6, I consider the role financial institutions played in driving my results. Section 1.7 concludes.

[8] Author's calculations from Wisconsin Department of Public Instruction Data, and the NCES Common Core of Data.

1.2 Background

In this section, I discuss institutional features of both teachers' unions and the Wisconsin education system which will be useful to the reader in understanding the nature of the main treatment effect I estimate in this paper. First, I describe the history of public union regulation and formation. Second, I describe the nature of the regulatory changes as implemented in Wisconsin. Finally, I describe the roles played by the state and local governments in the Wisconsin education system.

1.2.1 Collective Bargaining History for Teachers

Teachers in the United States are heavily unionized. According to CPS data[9], 50 percent of teachers in the US claim to be union members. In states with favorable regulations, such as Wisconsin in 2011 and earlier, virtually every teacher is covered by a union contract. Compared to the private sector, or other public sector employees, teachers have one of the highest rates of unionization in the country.

This unionization began in the 1960s, as states (beginning with Wisconsin) began passing laws which created a more favorable regulatory environment. These include laws granting state and local unions the right to strike, or the right to mandatory arbitration. One widely adopted law was the so-called "Duty to bargain" rule, which required employers to recognize their unions and bargain with them in good faith. 32 states adopted such rules by 1980, and 36 by 1990. By the 80s, most states which were going to adopt duty to bargain laws had already done so. During the 90s and 2000s, only one state experienced any change to its duty to bargain rules.[10]

Public unions were able to operate in these relatively favorable environments for decades following their initial formation. Increasingly, however, state governments are beginning to

[9] See Hirsch and Macpherson (2016).

[10] New Mexico experienced a number of changes relating to a provision lapsing and being subsequently renewed, discussed in Lindy (2011). Discussion on the timing of changes sourced to summary tables from Hoxby (1996) and Frandsen (2016). See also Murphy (1990) for a useful history.

reevaluate the public interest in these regulations. Interestingly, modern policy changes tend not to operate by rolling back the same types of benefits that lead to public unions initially forming, but rather by considering new regulations which are outside of the historical regulatory legacy. One example of this is "right to work" legislation, which limits union financial power by prohibiting mandatory dues collection. That is, typically unions raise funds by collecting fees from all employees covered by their contracts. Right to work laws prohibit this, making it so unions only receive voluntary contributions from members. Recently states have been adopting right to work laws that affect all unions.[11] Others, such as Indiana in 2005 and Illinois in 2015, have implemented right to work rules for state employees by executive order. In 2011, Wisconsin went further than any state had at the time, in terms of implementing not only right to work for all public employees, but also a variety of new limitations.

1.2.2 Wisconsin's Act 10

Wisconsin's Budget Repair Bill of 2011, or Act 10, was a broad revision to public union power in Wisconsin. It operated by lowering the benefits and increasing the costs to unionization. It lowered the benefits through a series of measures which limited a public unions' contractible options. It increased costs both by implementing a public right to work measure, and by creating a new system of recertification elections.

The law limited a union's contractible options in a number of ways. First, it limited contracts to a one-year period. Second, it limited "base" pay increases to a cost-of-living adjustment, and required a referendum in order to exceed this limit.[12] Third, and perhaps most importantly, it limited collective bargaining to salary compensation only.

[11] Indiana and Michigan, 2012. West Virginia, 2016.

[12] In practical terms, no such referenda have ever been held. This is in part due to the legal definition of "base pay" and the flexibility it affords to school districts. For example, pay increases due to experience, education, or other incentive/merit pay are not subject to this limitation. So a hypothetical school district which had full retention and several teachers earn masters degrees would see its total pay increase by more than cost of living, and would not need to hold a referendum. So while this requirement affords the districts a great deal of flexibility, it has been effective in terms of forcing districts to motivate pay increases by other means than simply broad base increases.

Act 10 also created a new system of recertification elections. In all states, as well as in Wisconsin prior to Act 10, unions were only required to hold a vote for an initial certification, or for a decertification in the rare event that a union chose to end formally operating. Act 10 changed this by requiring that unions hold an annual affirmative vote to maintain legal status. Elections for unions operating within school districts are held every November beginning in 2011[13]. Thus, taken as a whole, the law subjected unions to a challenging administrative process, which made it more difficult to maintain legal status, limited their fund-raising abilities, and reduced the benefits to organizing.[14]

Following the announcement and ultimate passage of Act 10, Wisconsin unions and their political allies organized an intense response. The intensity of this response is noteworthy because it shows that the act was regarded as a serious threat. In addition, it is helpful in understanding why only Wisconsin and Indiana have implemented such far-reaching reforms. The suddenness of the reaction upon the act's announcement also helps to demonstrate that Act 10 was considered surprising - an important characteristic for my identification strategy, which I will return to discuss in greater detail later.

The response to the law included protests, lawsuits, recall elections, as well as unorthodox political maneuvering by members of the Wisconsin Senate. Upon the announcement of the law, all 14 Democratic members of the Wisconsin State Senate fled to Illinois to deny the senate a quorum.[15] Protests that began outside (and indeed, at times inside) the state house in Madison continued from February until June, and at their peak included as many as 100,000 protesters. A total of sixteen state senators were subject to recall elections following these events, including 8 Republicans who supported the bill, and 8 counter-recalls targeting Democrats for fleeing the state. Three Republicans lost their recall election. Governor

[13] November 2012 elections were canceled because of an ongoing legal dispute - the law was temporarily overturned during the window in which they would normally be held.

[14] I attempted a regression discontinuity analysis to determine the effect of losing these elections, but the low density of districts near the election cutoff meant that there was insufficient power to draw reliable conclusions, so the regression discontinuity analysis is not reported in this paper.

[15] Senate Democrats would remain in Illinois for about three weeks, but proponents of the bill ultimately found a way to restructure it so that it could bypass quorum requirements.

Walker was targeted for recall as well, although he was safely reelected.

On top of the above changes, Wisconsin reduced its total state aid to local school districts as well. In public statements, the architects of Act 10 explained that Act 10 and the budget cuts were interrelated components of a single policy goal. The plan was to allow local government savings via reduced union power, and then translate that to state government savings via reduced aid. During the financial crisis in the years preceding Act 10, the Wisconsin state budget was stretched extremely thin. The state relied on short term Federal stimulus grants in order to maintain funding for public schools. In addition it initiated emergency transfers from the transportation fund to educational use on three separate occasions.

1.2.3 Wisconsin School Finance

As I estimate the causal effect of union regulations on teacher compensation one key threat is Wisconsin's reduction in state aid to local districts. As this change was concurrent with Act 10, it will be important to consider both changes in the particular context of Wisconsin's school finance system. I describe here two key features to Wisconsin's state-local funding system for school districts. The first is progressive state equalization aid. The second is state-set revenue limits. Both of these systems are mechanically determined by the state, and limit district operations in various ways.

Every year, Wisconsin calculates state "equalization revenue", which is a form of general aid from the state to local districts. Equalization revenue is designed so that districts with higher property values receive less in aid, and districts with lower property values receive more. This equalization revenue is the sum of two specific components - the first is per-pupil property values in the district, and the second is prior year spending per pupil. The state every year sets a total budget and maps these two components to the total district aid. Wisconsin also provides "categorical" aid in addition to general equalization aid. Categorical aid includes special education funding, and transportation aid for busing students. Equalization revenue was reduced by 5% in aggregate between 2011 and 2012.

The equalization aid provided by the state operates in a system of revenue limits. The revenue limits are binding on the combined district equalization revenue and local property tax revenue. Other forms of aid, such as Federal grants, or state categorical aid, are exempt from these limits. The total revenue limits faced by districts are driven largely by inertia; they were set based on spending in the 1993 year and are primarily only adjusted by either long term changes in enrollment, or with a district passing a referendum to increase its revenue limit. In 2012, Wisconsin cut revenue limits along with its equalization aid. On average, districts faced a 5.5% lower limit.

1.3 Data

1.3.1 Wisconsin Collective Bargaining Agreements

In my main analysis in section 1.4, I compare districts based on the expiration date of the collective bargaining agreement in the period prior to the passage of Act 10. In order to make this comparison, I collect finalized contracts which have been agreed upon by both union leadership and the district administration.[16] These agreements allow me to accurately measure the timing with which a given district is required to renegotiate its contract. Because Act 10 grandfathered collective bargaining agreements signed prior to implementation, I am able to use these contracts to determine the exact point in time at which Act 10 affected different districts.

School district contracts traditionally cover a period of two years, beginning and ending in odd-numbered years. This is primarily driven by district reliance on state aid. Because the state government uses a two year budgeting cycle, it is natural for districts to contract based on the same time frame. Thus, the overwhelming majority of Wisconsin unions have contracts that expired on June 30, 2011. In the course of my contract review, there were

[16] The majority of these contracts were provided to me by an organization which had collected these contracts for their own purposes. Said organization requested anonymity as a condition of providing these contracts.

16 districts which were off-cycle in this time frame, in such a way that 1) spanned the 2012 school year, or longer, and 2) were finalized prior to any public announcement of Act 10. The main reason a district would be off-cycle is because a dispute in a previous cycle lead to the all parties agreeing to a longer-than average contract at some point. I discuss summary state comparisons between these districts and the universe of Wisconsin schools in greater detail in section 1.4.

My data collection effort sought all contracts covering the years between 2007 and 2013. Because of Act 10, not all districts have contracts following 2011. As a result, I supplemented the information included in the contracts by contacting district administration in certain cases. A full taxonomy of contract timing patterns is included in the data appendix. I use the signing dates and termination dates of the contracts to construct a variable measuring the initial period of treatment.

1.3.2 School Finance

Finally, I use administrative Wisconsin school district data to look at compensation outcomes and observable district characteristics. The Wisconsin data is collected by the Wisconsin Department of Public Instruction ("DPI"), and includes teacher compensation, school finance, school enrollment, and other characteristics.

The DPI compensation data I use here is maintained at the teacher level. It records the teacher level salary, value of fringe benefits, years of experience, education level, and employing district. Fringe benefits in this case refers to the value of the premiums paid by the district on the employees behalf. It is primarily driven by health and retirement premiums. I collect this data from 2003-2016.

Other DPI school district data includes financial and enrollment reports, which are collected for the years available (2005-2015). This includes school district demographics, revenues by source (state, Federal, local), spending by category (instructional, administrative, support), and other administrative information such as the district revenue limits.

In section 1.6 I include national state-level school data from the Common Core of Data ("CCD"), an administrative K-12 database of all US school districts, maintained by the National Center for Education Statistics at the US Department of Education. The CCD includes financial data, such as teacher compensation, teacher benefits, grants by source, and total spending. It also includes population characteristics, such as enrollment by race, special needs enrollment, and total hiring. Summary characteristics are included in Table 1.11.

1.4 Wisconsin Event Study

In this section, I describe my main analysis: the causal effect of unions on teacher compensation, driven by plausibly exogenous contract timing. I start by informally describing the key variation - the first year after Act 10 in which a Wisconsin teachers' union is required to renegotiate its contract. Next, I justify this variation as exogenous to other factors which would determine wages. I then formally describe the empirical specification, and present results.

1.4.1 Identification Strategy

In estimating the causal effect of unions on compensation, one key challenge is to identify exogenous variation in union status. For example, if unions are more likely to form in districts where pay is particularly low, but these low-paying districts are only observed in their treated, unionized state, that could bias my estimates downward.

The literature has typically taken one of two approaches to address this problem. One approach, more common in evaluating private sector unions, is regression discontinuity. Because unions must hold a vote to be certified, there will be a discontinuity in the voting results around 50%. As long as some component of the vote is unpredictable, this will identify the causal effect of the union certification. The main challenge with this approach is that it requires a large sample of votes near the cutoff point.

The other main approach, more common in the public union literature, is to look at new districts formed after exogenous state rule changes. This is the approach of both Hoxby (1996) and Lovenheim (2009). This approach treats a state's policy choices as exogenous. Hoxby provides a useful argument for the validity of this approach. Many states made what were ultimately very similar policy choices, but had timing differences for the exact period of implementation. These timing differences do not represent discrete changes in underlying state preferences. Rather, there is a continuous underlying function which manifests itself in a discrete pattern.

The main approach I use in this study is a new strategy, somewhat related to the state law change approach, which exploits plausibly exogenous contract timing. School districts in Wisconsin typically adopt a two year contract cycle. As a result, most districts are immediately exposed to Act 10 in the first year following its passage (as Act 10 passed in early 2011, this means 2012 is the first year of full treatment for these districts). The pattern is only a matter of convenience for districts, however, and there is no requirement that unions and districts contract according to this cycle. As mentioned in section 1.3, I review case by case individual collective bargaining agreements to identify cases where districts were off cycle in such a way that 1) their existing contract covered a period after most districts were exposed to Act 10, and 2) the provisions of these contracts were negotiated prior to any knowledge districts would have had of the impending passage of Act 10.

I argue these off-cycle districts form a natural exogenous comparison group for districts immediately exposed to Act 10. I base this on two key reasons. First, I argue that off-cycle districts, while rare, are off cycle for idiosyncratic reasons which do induce any meaningful selection problems. The main reason a district would be off cycle is because it had problems in negotiating contracts in a prior cycle (for example, in the 2007-2009 period or earlier). Because prior negotiations delayed a contract, both the union and the district administration would agree that the expedient thing to do is sign a longer, off-cycle contract, simply to avoid excessive negotiation. One concern might be that districts which have longer negotiations in prior periods are likely to be different than those which have shorter negotiation periods in prior years. I will discuss the empirical differences between these groups in greater detail in section 1.4.2 below.

The second reason I argue that these are an exogenous comparison group is that I select only those districts which have a contract agreed upon prior to any public announcement of Act 10. While some 150 districts had contracts covering 2012 or 2013 which were signed prior to the enactment of Act 10, this represents an endogenous response which would have been influenced by knowledge of the bill's impending passage. The group I select suffers from no such contamination. I use February 14, 2011, the date Act 10 was publicly announced, as the main reference point in terms of prior knowledge. An additional possible concern is that districts had some knowledge of this prior to announcement. I argue that this is unlikely. First, the reaction to the law (as described in section 1.2.2) was immediate and intense. Had the law been anticipated, there would not have been such a stark change in behavior from unions and union allies on the 14th. Furthermore, independent fact-checkers have made the case that Governor Walker did not campaign on the provisions of Act 10.[17] Finally, many of the districts which signed extensions did so in 2009 or early 2010, well before it was clear that Republicans would take complete control of the state government.

This identification strategy thus uses plausibly exogenous timing exposure to Act 10 based on contract renegotiation dates. That is, off-cycle districts are comparable to those on-cycle, but the similarity of these groups changes in 2012 as a direct consequence of Act 10. Because the date of initial exposure varies by district and year, I can include district and year fixed effects to address concerns about selection on district or time specific differences. I will discuss the main threats to this identification strategy in the section below.

Because all districts must eventually renegotiate, the estimates in the section should be considered an average treatment effect of the provisions which effect all Wisconsin districts

[17] A series on Governor Walker and Act 10 by Politifact and the Milwaukee Journal Sentinel makes the case succinctly: "During the campaign, Walker prided himself on presenting many specific proposals to voters...But nowhere in our search did we find any such detailed discussion of collective bargaining changes as sweeping as Walker proposed."

at different times. This includes such provisions as the limited contract options, the limited ability for unions to collect resources from covered employees, or a reduced union "threat power" because of new possibility of forced decertification. It excludes the effects of state financial aid cuts, as these effects all districts at the same time.

1.4.2 Comparison of Treatment Groups and Pre-trend Analysis

In this analysis, I assume that off-cycle districts make a valid comparison group for those districts immediately treated by Act 10. This analysis relies on two key assumptions. The first is that the immediately treated districts and the off-cycle districts have parallel trends in the outcome variables. If off-cycle groups had a higher rate of pay growth, perhaps because their unions were more willing to fight for larger pay increases, I would risk conflating the effects of Act 10 with this secular trend. The second threat is that these two groups do not face any differential shocks in the treatment period, other than the union power effects of Act 10.

I look at the trends in district compensation in Figure 1.1. This Figure looks at the nominal[18] total compensation by treatment group (meaning, initially treated vs "early extender" or "off cycle" districts) between 2005 and 2016. At the beginning of this period, from 2005-2011, the two groups have extremely similar levels and trends in teacher compensation. This relationship discretely breaks beginning in 2012, when the early extenders start to have a higher levels of total compensation. Eventually, all districts are required to renegotiate their contracts, and the outcome variables re-converge. Trends in salaries and benefits are also shown separately in Figures 1.2 and 1.3. Figure 1.2, depicting salaries, shows that salaries continue to rise for the early extender districts after Act 10. The trend in benefits in 1.3 shows a partial decline in benefits in 2012 even for those early extender districts. This suggests a non-contractible component of benefits which was affected by Act 10 for all districts, regardless of contract cycle. One possibility is that district administrators were able to find

[18] The real compensation trend demonstrates an identical story. I present the nominal numbers because districts and unions bargain over nominal, rather than real wages.

less expensive health benefit providers. As a result, many of the effects I show for benefits may in fact be under-estimates. Taken as whole, these figures are remarkably consistent with contract timing being the main driver of differences, and suggest little to no overall selection.

I also look at balance in observable characteristics between these groups in Table 1.1. This Table presents differences in district mean characteristics.[19] Across most characteristics, such as the share of rural districts, the share of black and Hispanic students, the share of students economically distressed, and total enrollment, there is no statistically significant difference across the two treatment groups. This is both because point estimates are not far off, and also because there is broad variation in each of these characteristics within each group. On average, the early extender districts are somewhat wealthier, and thus receive less in state aid.

I look for differential shocks by treatment status in Figures 1.4 and 1.5, and Table 1.2. Figure 1.4 looks at per pupil revenue limits in each group up to and following Act 10. Figure 1.5 does the same for external (state+Federal)[20] grants. Figure 1.4 shows that revenue limits moves similarly prior to Act 10 and experienced a similar shock. Figure 1.5 shows a similar shock for external funds, and a slightly steeper increase in state funds over time for districts immediately treated. In Table 1.2 I formally test for a difference-in-differences in these characteristics and find no evidence of a differential shock. The changes in revenue limits and external grants are mechanically determined and unlikely to be an outcome of the other provisions of Act 10.

Overall, my assumption of parallel trends appears to hold. There are some differences in levels of observable characteristics, though no apparent threat of a differential fiscal shock.

[19] I take for each district the mean characteristics for the three years prior to treatment, 2009-2011, and average these across districts within the group which is immediately treated in 2012 and the group treated in later years.

[20] There are two reasons why it is important to look at the combined trend in state and Federal grants. First, these funds serve a similar purpose of lowering a district's reliance on local property tax revenues. Second, Wisconsin used Federal funds from the State Fiscal Stablization Fund, a major program of the American Recovery and Reinvestment Act of 2009, to supplement state grants during the Great Recession. Thus the combined number is the best measure of what changed in 2012.

As an added precaution, I run specifications with district-specific linear trends, and with districts reweighted based on the observable characteristics in Table 1.1.

1.4.3 Empirical Specification

Here, I describe the semi-parametric and parametric event study specifications used to estimate the effects of Act 10 within Wisconsin. I also describe the district matching and reweighting equation.

The full event study specification is as follows:

$$Y_{dt} = \sum_{j=-8}^{3} \gamma_j I(t - year = j) + \eta X_{dt} + \pi_d + \delta_t + \varepsilon_{dt} \tag{1.1}$$

I estimate this equation at the district-year level, running a panel of 418[21] school districts from 2004 to 2016. Districts at denoted by d and year by t. Formally, I construct a balanced panel of all districts using the observations 8 years prior to contract negotiation though 3 years after.

The key treatment variables are defined by the function $I(t - year = j)$. This function is equal to 1 if the year is j years from the initial treatment, and zero otherwise. In this case, "treatment" is defined by being required to renegotiate a contract following Act 10. I omit the year t-1, so all treatments are defined by the year immediately preceding renegotiation. This approach allows me to track the differences between the early and late groups by year (γ_j) before and after the law is implemented. I include treatments for up to three years after treatment as going outside this window causes sample size to drop for districts which had longer pre-period contracts. This approach has the advantage of identifying treatment effects by year. In addition, the pre-treatment variables offer a specification test. I expect all γ_j for $j < 0$ to jointly be zero.

This specification includes district fixed effects (π_d) to address union effects based on

[21] There are 424 public school districts in Wisconsin in 2016, but the number is reduced as I consolidate all mergers and splits over this time frame into single districts.

changes within a district relative to its mean. Additionally, I include year (δ_t) fixed effects to address shocks that effect all of Wisconsin, such as the effects of the Great Recession. The variable X_{dt} keeps track of district-year variables which are typically not a choice of the district. These include district enrollment, property values, and student characteristics such as race and poverty. Finally, outcomes Y_{dt} include average teacher salary, teacher fringe benefits, and teacher compensation (salary + benefits) in a given district-year.

In addition to the flexible, semi-parametric event study, I run the more parametric specification below:

$$Y_{dt} = \gamma PostAct10_{dt} + \eta X_{dt} + \pi_d + \delta_t + f(dt) + \varepsilon_{dt} \tag{1.2}$$

The primary difference between equations 1.2 and 1.1 is that 1.2 replaces the flexible treatment-by-year effects with the single binary variable $PostAct10$. $PostAct10$ is a zero until the first period in which a district has to renegotiate a contract after the passage of Act 10, and is one thereafter. This equation is more parametric because it averages together annual effects into a single treatment effect. Including equation 1.2 simplifies comparisons across districts, and the single treatment parameter allows me to to include a panel of districts balanced by school/fiscal year.

I add several measures of robustness to equation 1.2. First, $f(dt)$ is a set of district-specific linear trends, which allows me to test of the possibility of differential pre-trends. In this specification, X_{dt} allows for a richer set of possible covariates[22] including state and Federal grants, property values, and district enrollment. As a final precaution, I re-weight all districts based on observable characteristics. The reweighting equation is a logit[23] model which assesses the probably of a district being a early-extender based on its pre-period observable characteristics[24]. Balance in reweighted observable characteristics are included in

[22] The semi-parametric model in equation 1.1 requires 2016 teacher salary to see three years of post treatment. As I have teacher salary but not the full set of district finance data for 2016, I must omit some of these variables to fit equation 1.1.

[23] Probit models perform equally well.

[24] Matching characteristics include suburban/rural, levels in state grants, Federal grants, district property

Table 1.3.

The within Wisconsin analysis has a number of useful features. All districts face a common legal environment, which reduces the risk of endogenous policy selection, a concern common to cross-state comparisons. That is, in a cross state comparison one might worry that Wisconsin passed this law because it was on a different trajectory than other states. In this environment, the variation is at the district level, while districts themselves are not the actors choosing the policy environment. In addition, every district within Wisconsin faces the effects of reduced state grants, which allows this difference to be interpreted as a pure union power effect, rather than conflating changes in union power with changes in district finance.[25]

1.4.4 Results

The main results from equation 1.1 are shown in Figures 1.7 and 1.8. These figures show the point estimates and 95% confidence interval of the treatment parameter γ_j leading to and immediately following treatment. Figure 1.7 measures compensation in logs and 1.8 in levels. Prior to Act 10 contract renegotiation, treatment parameters are individually and jointly statistically insignificant (F statistic .23, when using the log compensation specification). Point estimates are near zero but slightly positive, perhaps reflecting the non-contractible change in benefits driven by the law, as seen in Figure 1.3.

Following a district renegotiating their contract under the new regime, the effect on total compensation is immediate, economically significant, and statistically significant. Teachers see their total pay drop by 8% or $6,000 in year 1, with slight increases over time. The post-event parameters are individually and jointly (F statistic 12.8 in the log specification) significant, although I cannot reject the assumption that these are equal by year in the post period (F statistic .96 in the log specification).

values, and 2012 shocks in external funding

[25] One final concern might be that this approach conflates the effect of changes in a legal regime and the effect of a given union decertifying. I also run a specification which explicitly controls for lost elections, though I do not include these here as it causes no perceptible change to the point estimates.

As the treatment variables are seemingly constant by year, or close enough together that it makes sense to average them, I focus the rest of the discussion section on Tables 1.4 and 1.5, which estimate the parametric equation 1.2 in log and level measures of teacher compensation, respectively. This Table shows that results are robust across a variety of specifications and show a similar pattern of the separate contributions salary and wages make to the combined effect. Each cell in Table 1.4 is a separate regression, with different outcome variables by row and different specifications by column.

This Table checks for robustness by making three main changes. First, this Table adds district-specific linear trends. Next, it drops observable district covariates, both out of concern that these could be interpreted as outcomes, and to add additional years of data where the covariates are unavailable. Finally, I repeat these three analyses while reweighting districts to more closely match off-cycle districts, to address the concern that these districts are meaningfully different from those on-cycle. Across all specifications, the point estimates stay close to one another.

Looking within a given column in Table 1.5, it is clear reductions in benefits are driving results more so than teacher compensation. Roughly two-thirds of the total changes are being driven by benefit reductions, rather than salary reductions. Because salaries are larger than benefits overall, this corresponds to a much larger percentage drop in benefits relative to salaries, as shown in Table 1.4. Salaries fall by about 4% as districts are required to renegotiate, and benefits fall by roughly 18.5%[26]. Combined, these show that districts are able to reduce teacher compensation by about 8%. This alone translates to 3.2% lower operating costs per district.

The point estimates here are larger than those estimates in the quasi-experimental public unions literature. My salary effects are comparable to those estimated by Hoxby, and larger than those estimated by Lovenheim. Including teacher benefits, my point estimates are much

[26] I report coefficient point estimates in tables, but translate these to percent changes in the discussion. Because benefit changes are far from zero, they must be transformed for a percentage interpretation. Here, I pick .17 as the benefits effect, meaning a $100 * (e^{.17} - 1) = 18.5\%$ effect on benefits.

larger than both. This last point is unsurprising as Hoxby and Lovenheim both do not have data on fringe benefits, but this result emphasizes the importance of including this measure when estimating the effect of public unions on compensation.

The 8% compensation effect I estimate here is determined by the components of Act 10 that affect all districts at different times based on their renegotiation window. That is, it would include limits on contract length, limits on bargaining over benefits, but exclude the effects of losing a recertification vote.[27] Additionally, it excludes the effects of cuts in state aid or revenue limits, as this is something which all districts are exposed to at the same time, unlike the contract provisions. As Wisconsin had a generous regulatory system for public unions, and near-universal union representation prior to Act 10, the point estimates here are most generalizable to other states which have a generous regulatory system for public unions.

1.5 Distributional Effects

In this section I discuss the effects of unions on the distribution of wages. I begin by discussing the policy motivation of this question. I then discuss the empirical methodology, and end with a discussion of results.

1.5.1 Policy Motivation

Understanding the distributional effects of unions is an important policy question because it offers insight into which members unions are more likely to reward, as well as how unions might affect the production function for student learning. Unions, rent-seeking or otherwise, would all have reasons to want to increase total teacher compensation. But compared to a voice union, a rent-seeking union may be more interested in directing these increases towards long-standing members. Alternatively, compared to a rent-seeking union, a voice union may

[27] I note that additional specifications which include measures of recertification vote losses do not change my point estimates.

be more interested in offering more incentives to hire the best young teachers.

One policy problem in education is that math and science teachers are considered valuable, due to their specialized and technical knowledge. But these candidates will also often have higher outside options that the average teacher (see West (2013) for a useful discussion). As a result, union contracts which limit pay raises to years of experience and education, regardless of field, may distort a district's ability to recruit this type of candidate.

A number of papers have shown that unions affect not only the level of wages, but the entire distribution. DiNardo, Fortin and Lemieux (1995) ("DFL") and Card (2001) are two examples, both of which have shown a strong compression effect on wages. Both papers used CPS data and made comparison across industries, although each employed different methods. Card constructed predicted wages from the non-unionized sector based on employee characteristics, such as years of experience and education level. He then calculated the union wage premium by difference-in-means with each predicted non-union wage decile, and showed that this was largest in the lowest predicted wage decide, and negative in the highest decile. DFL used the non-parametic approach of reweighting individuals in the CPS to match union and non-union employees on observable characteristics, and then comparing the total density of wages across each of these distributions. Similar to Card, the main finding is that unions compress the distribution of wages.

1.5.2 Methods

I replicate below the approach used by Card. It is important to note, however, that there are a number of differences between the context of that study and my own. Here, I have access to the universe of observations within a given industry, whereas Card relied upon cross-industry comparison and randomly selected observations within each industry. In addition, I have variation in union status which is attributable to a specific policy change, whereas Card relied on a selection on observables comparison across sectors. Thus, in this setting my distributional comparisons will be richer and better identified, but also more industry-specific and more difficult to generalize.

First, I empirically test the distributional effects of unions by constructing a predicted wage quintile. The predicted compensation equation is based off three variables: teacher experience, teacher education, and their interaction.[28]

$$Y_i = \beta_1 Exp_i + \beta_2 Ed_i + \beta_3 Exp * Ed_i + \varepsilon_i$$

I estimate the coefficients on this model based on 2003 data, which is subsequently dropped, and define the quintile for each year based on out-of-sample predictions of $\hat{Y}_i$. The final step is to run the Wisconsin event study separately within each quintile of teacher predicted compensation. The following equation is identical to equation 1.2, with the main difference that I construct within-district means by predicted wage quintile and run the regression separately by each predicted wage quintile.[29]

$$Y_{qdt} = \gamma PostAct10_{qdt} + \eta X_{qdt} + \pi_d + \delta_t + f(dt) + \varepsilon_{idt} \tag{1.3}$$

1.5.3 Results

Results from estimating equation 1.3 are included in Table 1.7. In all specification in this section, I include district and year fixed effects, district-specific linear trends, observable characteristics, and reweight to match off-cycle districts. The leftmost column (Predicted quintile 1) in Table 1.7 shows the teachers who earn the least under the 2003 wage distribution (that is, the least experienced teachers, and those without masters degrees). The rightmost column (predicted quintile 5) shows those teachers who earn the most under the old regime. Moving from left to right, total compensation falls for all teachers, but it falls most for the highest earning teachers under 2003 pay distribution. This is again driven by both salary and wages, but the point estimates for salary are not statistically significant in each quintile.

[28] According to union contracts, teachers are paid as a fairly strict function of education and experience.

[29] As an alternative specification, I run this regression at the teacher level. This approach creates the same distributional patterns I describe below.

In fact, it is possible that salaries are increasing for the least experienced teachers in column 1.

Taken literally, this suggests that teachers' unions have a compensation dispersal effect, rather than the compensation compression effect identified by Card, or DiNardo et al. It is important to note that these are not perfect comparisons, as other estimates of union wages considered comparisons across industries and sectors. What Table 1.7 shows should be interpreted as a reduction in the return to tenure. This pattern is most consistent with a rent-seeking model, where union members seek to direct rents to the most-established members, and newer members are less able to share in the rents. This is not, however, the only possible interpretation. A voice union may wish to pay less to newer teachers because the marginal product of less experienced teachers is lower.

1.6 The Role of Financial Instituions

In this section, I address a major concern relating to a concurrent policy change. That is, Wisconsin cut state aid to school districts at the same time as they reduced union power. While I implicitly address this in the balance testing and reweighting exercises in section 1.4, in this section I attempt to estimate the direct effects of the financial aid cuts which occurred concurrently with the reduction in union power. First, I categorize the schools by quartile of fiscal shock size in 2012, and estimate treatment effects separately by each quartile. Next, I considering cross-state comparisons in compensation following Act 10.

1.6.1 Treatment effects by fiscal shock quartile

As a major concurrent policy change, this reduction in state grants in 2012 is the primary threat to this analysis. In interpreting the results in section 1.4, it is important that I distinguish between the effects of union power and the effects of reduced district funding. While Table 1.1 shows no statistically significant differences in terms of fiscal shocks for on-cycle versus early extending districts, here I explore this threat in greater detail. I do this by ranking school districts by the size of their fiscal shock (defined as the change in state and Federal grants per pupil between 2011 and 2012), and then calculating the event study treatment effect from equation 1.2 separately by each quintile.

Before presenting heterogeneity results, I want to consider the selection into the defined fiscal shock quartiles. The change in grants between 2011 and 2012 is not random; it is determined by Wisconsin's aid formula. While the aid formula limits the ability of legislators to target specific districts, it does explicitly consider property values and prior year spending. As a result, major changes in this formula are likely to be correlated with these factors. Table 1.8 presents summary statistics by each fiscal shock quartile. A few points stand out. Districts which lost less money in 2012 tend to be wealthier districts which were already less reliant on state aid. In addition, these districts tended to be smaller, more rural, and have slightly higher revenue limits per pupil.

As the size of fiscal shocks are not distributed randomly across districts, I assess how this selection may affect teacher compensation in Figure 1.9. Figure 1.9 shows the trend in real teacher compensation between 2005 and 2010, by quartile of fiscal shock in 2012. Two things stand out. First, these groups tend to move together from year to year. The gap between the highest and lowest paid group is relatively constant between 2005 and 2010. The only major change is that quartile 4 and quartile 3 occasionally trade ranks. Second, there is no obvious relationship between the size of the fiscal shock and total compensation. The highest paid group is the second quartile of fiscal shock, while the lowest paid group is the lowest quartile of fiscal shock. Overall, it does not seem likely that differential pre-trends would drive differential union compensation effects by quartile.

Actual treatment heterogeneity by fiscal shock quartiles are included in Table 1.9. Each regression in Table 1.9 compares those districts immediately treated by Act 10 in a given fiscal quartile to all districts which were off cycle.[30] Across columns, we see that total compensation effects range between 7-8%, and that cross quartile comparisons produce similar estimates

[30] Due to the limited number of off-cycle districts, limiting this category by fiscal shock quartile as well was not feasible.

to the results when using the full panel of districts. It thus appears that fiscal shocks are less meaningful than union regulations in driving my results.

1.6.2 Cross-state synthetic control approach

In this section, I take the final step of looking at cross-state comparisons in teacher compensation. Because Wisconsin faced additional regulations at the same time as it faced reduced state aid, the cross state comparisons should provide the combined effect of regulations and state aid. If these estimates are similar to those obtained in the within-Wisconsin event study, that suggests that cuts in state aid played a limited role in reducing teacher compensation. Alternatively, if cross state estimates are larger than within Wisconsin estimates, that suggests state aid may have played an important role.

To make this comparison, I use the synthetic control methods of Abadie, Diamond and Hainmueller (2012). This approach is meant as a case study analysis - here, we are interested in the unique case of Wisconsin teachers. In this approach, states are weighted to construct a "sythetic Wisconsin" which mimics the time series of teacher compensation in Wisconsin. After the treatment period, the untreated synthetic Wisconsin and the treated real Wisconsin will diverge. I construct a synthetic Wisconsin using the Common Core of Data from 2005 to 2013, allowing two years of post-treatment.

Predictor variables used in this construction of synthetic Wisconsin include: teacher compensation in odd years between 2005 and 2011, the state trend in teacher compensation, special needs enrollment, race and ethnicity of the state student population, the share of students in the state attending rural school districts, and the share of students in the state attending urban districts. The weights used in the construction of synthetic Wisconsin are in Table 1.10. I include a comparison in characteristics between Wisconsin, synthetic Wisconsin, and the unweighted mean of all states other than Wisconsin in Table 1.11.

The trend in teacher compensation in Wisconsin relative to synthetic Wisconsin is included in Figure 1.10. In this figure, the synthetic comparison tracks Wisconsin closely until the treatment period, at which point a difference of \$7,430 emerges. This is similar, but somewhat larger than the differences of \$6,461 estimated in the event study. The somewhat larger gap is unsurprising, given that the amount determined in the event study was known to be an under-estimate.[31] Overall, the comparable numbers from the event study and synthetic control analysis suggest that the regulatory limits on contracting, rather than state aid cuts, are driving the falling teacher compensation.

In Figure 1.11, I replicate the synthetic control study while asserting a hypothetical treatment for each of 50 states.[32] I then take the $(syntheic - treatment)$ state differences and show the time series of differences for every simulation. When ranking all 50 replications by the size of the difference in 2012 and 2013, Wisconsin has the largest total drop. Thus, it is unlikely that the the differences produced by the synthetic control approach are driven by chance.

1.7 Conclusion

In this paper, I estimate the causal effects of union power on teacher compensation. I do this by comparing districts within Wisconsin based on the first year in which they are forced to renegotiate contracts following a major regulatory change in the state. I find that districts on different contract cycles follow similar patterns prior to the implementation of Act 10.

The package of reforms adopted in Wisconsin reduces total teacher compensation by about 8%. This is driven by a roughly equal share of salaries and fringes benefits, though as fringe benefits make up a smaller share of total compensation than salaries, these fall by a larger percentage of their relative value. These findings are robust to a variety of specifications and are driven by the union regulations, rather than concurrent policy changes

[31] See Figure 1.3. There is a component of benefits which drops following Act 10 even for the early extender districts, which the event study misses. The additional \$1,000 pay drop identified by the synthetic control comparison relative to the event study may be attributable to this drop.

[32] New York state is dropped from the analysis as its compensation is consistently higher than the rest of the country and was unable to form a comparable synthetic match.

such as a reduction in state aid to local districts. As union power falls, the associated reduction in pay fall disproportionately on the most experienced and highest paid teachers.

Overall, my estimates are larger than prior causal studies of the effects of unions on teacher unions. This is in part because my data includes the value of fringe benefits. There are a number of other important differences between my results here and those in prior studies. Perhaps most importantly, the setting of Wisconsin's Act 10 is policy relevant in a number of crucial ways. First, it looks at a reduction in power of existing unions, whereas the literature to date has primarily focused on the creation of new unions. Second, the set of policies used in Wisconsin's Act 10 have been considered elsewhere. They have been enacted in Indiana and were narrowly defeated in Ohio. Tightening state and local budgets make it likely that similar plans will be implemented elsewhere in the future. To the extent this occurs in states with favorable union regulations, such as those Wisconsin had prior to Act 10, these results are likely to be externally valid.

Finally, while compensation is important, both from the perspective of taxpayers and teachers, a great deal of additional work is needed to fully understand all aspects of Act 10. In particular, it is important to understand how the savings involved in Act 10 translate into student outcomes. If Wisconsin students meet similar levels of achievement following Act 10, the union regulations would appear to be a success. Alternatively, it is possible that these regulations demoralized teachers into putting in less effort, or otherwise negatively affected the K-12 production function. Ultimately a deeper understanding of how these payroll savings translated into either other inputs or student learning is necessary to both understand the program itself, and to understand the behavior of public sector unions more broadly.

1.8 Figures

Figure 1.1: Parallel trend comparison: nominal compensation

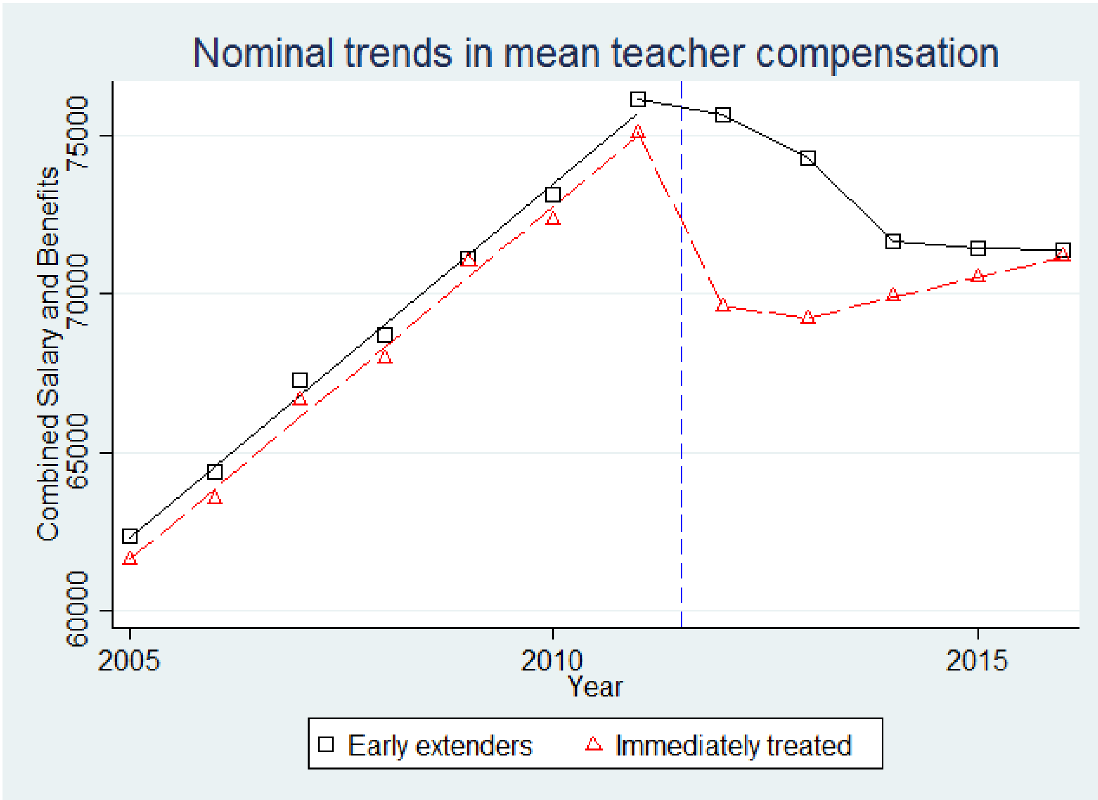

Figure shows the trends in teacher compensation (salary + benefits) for two types of districts over time. "Immediately treated" refers to those districts which were required to renegotiate a contract in the Post Act 10 regime beginning in 2012. "Early extender" districts are those which have a contract in place prior to Act 10's announcement which expired in 2013 or later. Points on the graph are means across districts (district weighted) within each group. The two groups are comparable on levels and trends in the outcome variable until 2012. At this point, the immediately treated districts sharply decline relative to the other districts. By 2015, all districts are required to renegotiate, and the two groups re-converge. Figure is consistent with little selection on unobservables; differences appear to be driven by contract status only.

Figure 1.2: Parallel trend comparison: nominal salary

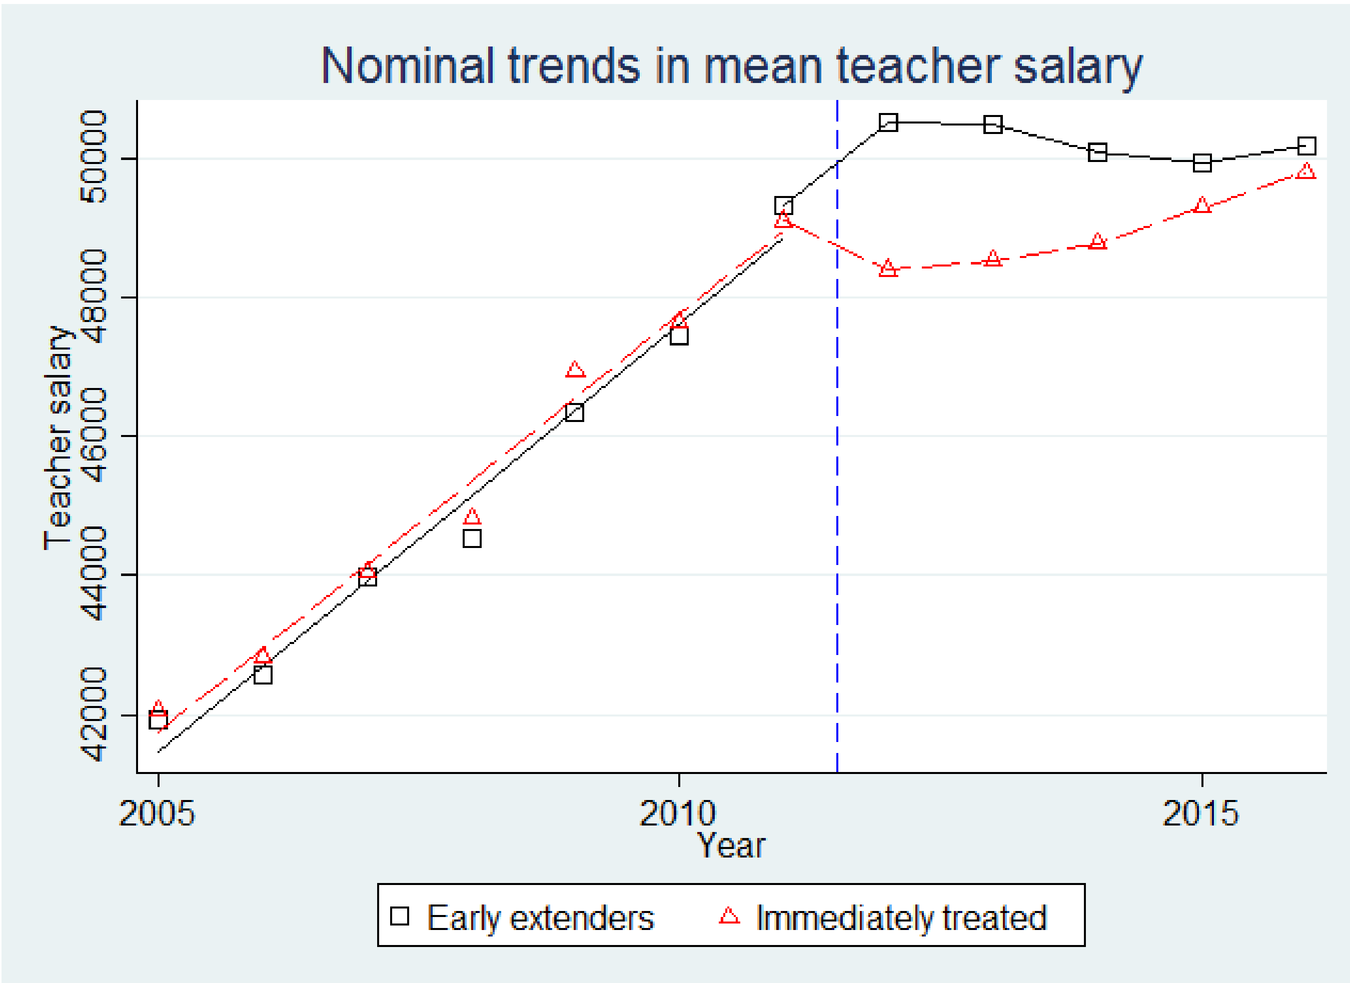

Figure shows the trends in teacher salary for two types of districts over time. "Immediately treated" refers to those districts which were required to renegotiate a contract in the Post Act 10 regime beginning in 2012. "Early extender" districts are those which have a contract in place prior to Act 10's announcement which expired in 2013 or later. Points on the graph are means across districts (district weighted) within each group. This Figure is broadly similar to Figure 1.1 above, which combines salary and benefits. The primary difference in this Figure is that salary continues to rise for the early extender group even after 2012. This suggests that having a contracts offered a more complete protection with regards to salaries than benefits.

Figure 1.3: Parallel trend comparison: nominal benefits

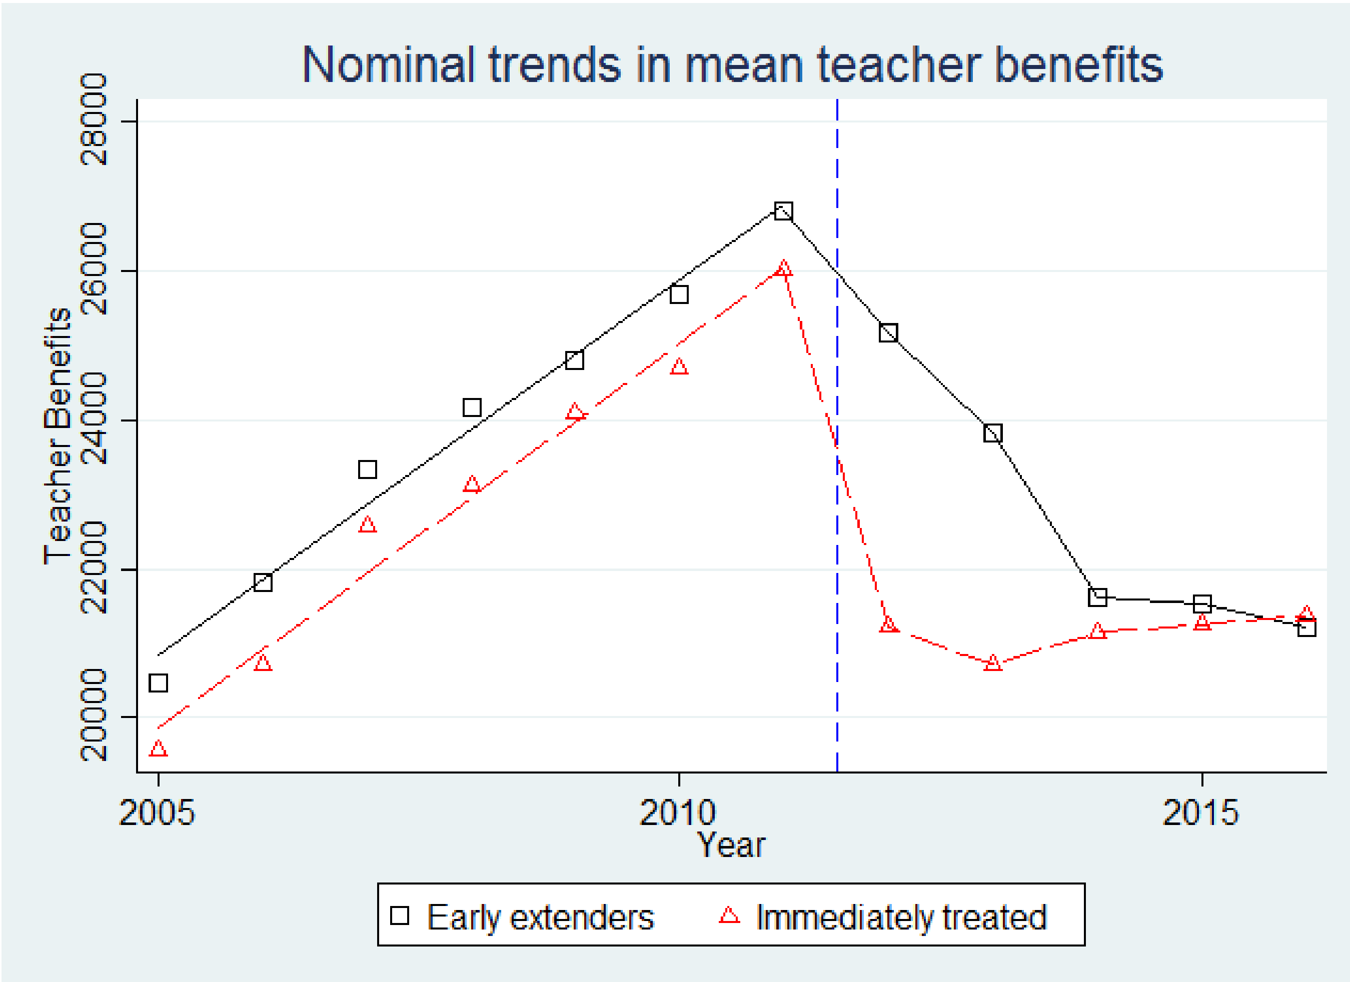

Figure shows the trends in teacher salary for two types of districts over time. "Immediately treated" refers to those districts which were required to renegotiate a contract in the Post Act 10 regime beginning in 2012. "Early extender" districts are those which have a contract in place prior to Act 10's announcement which expired in 2013 or later. Points on the graph are means across districts (district weighted) within each group. This Figure is broadly similar to Figure 1.2 above, which looks at the trend in salary by group. The primary difference in this Figure is that benefits fall immediately in early extender districts (albeit less than they do for immediately treated districts), while salary continues to rise. This suggests that there is a non-contractible component of benefits which was still affected by Act 10. As a result, my estimates will be an under-estimate.

Figure 1.4: Parallel trend comparison: revenue limits

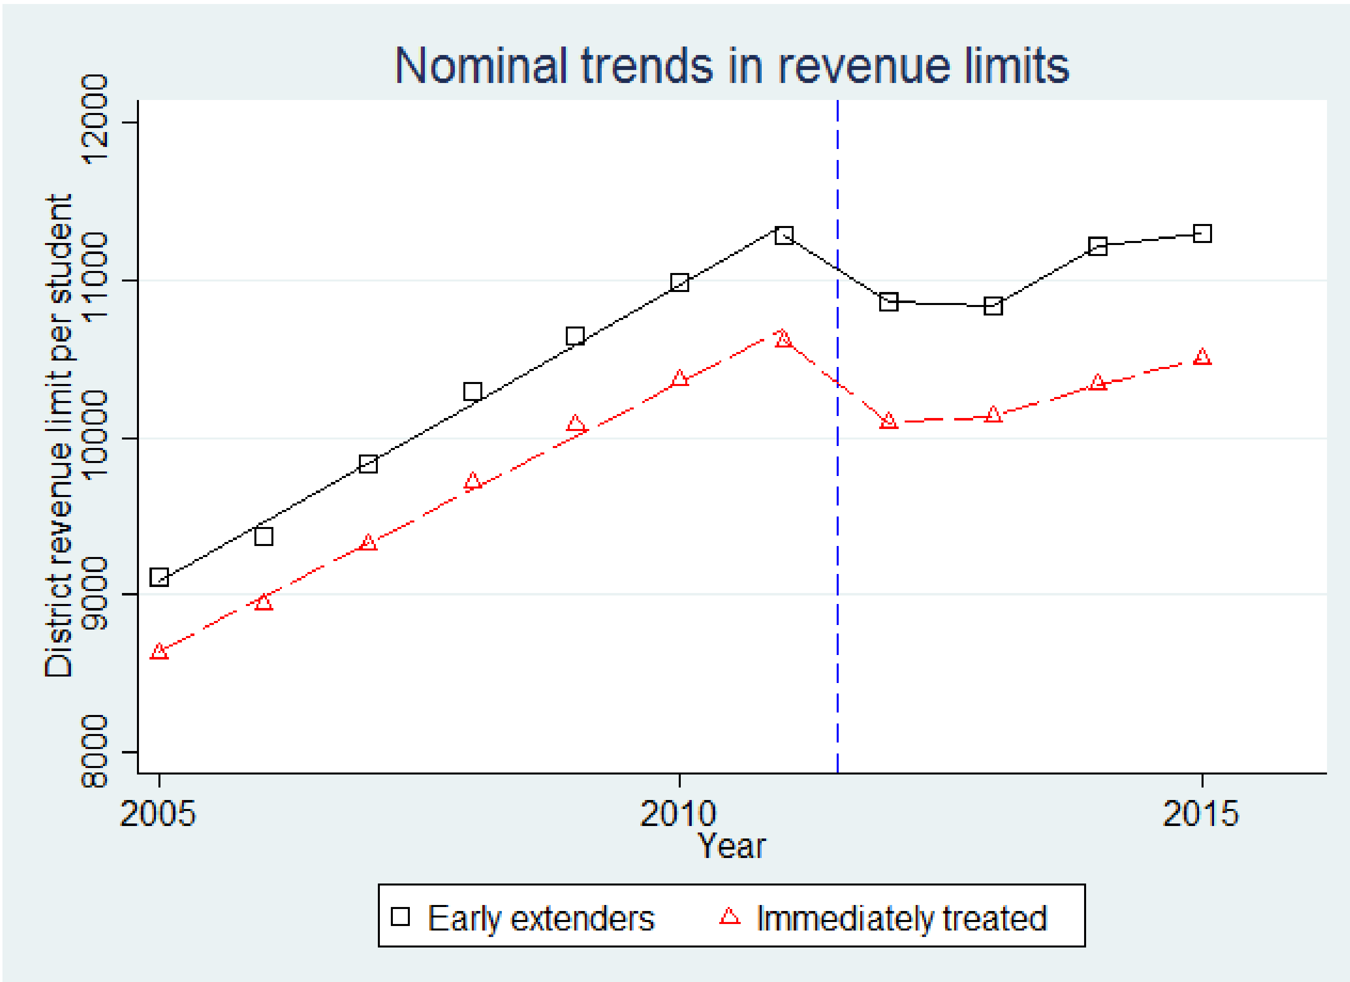

Figure shows the trends in state-assigned district revenue limits. "Immediately treated" refers to those districts which were required to renegotiate a contract in the Post Act 10 regime beginning in 2012. "Early extender" districts are those which have a contract in place prior to Act 10's announcement which expired in 2013 or later. Points on the graph are means across districts (district weighted) within each group. This Figure investigates whether there is a distinct shock to limits coincident with Act 10, across these groups. While both types of districts receive a fall in revenue limits, this fall is roughly equal in each group. Districts have similar growth trends in the pre-period. Refer to Table 1.2 for a formal difference-in-differences test of this characteristic.

Figure 1.5: Parallel trend comparison: state and Federal grants

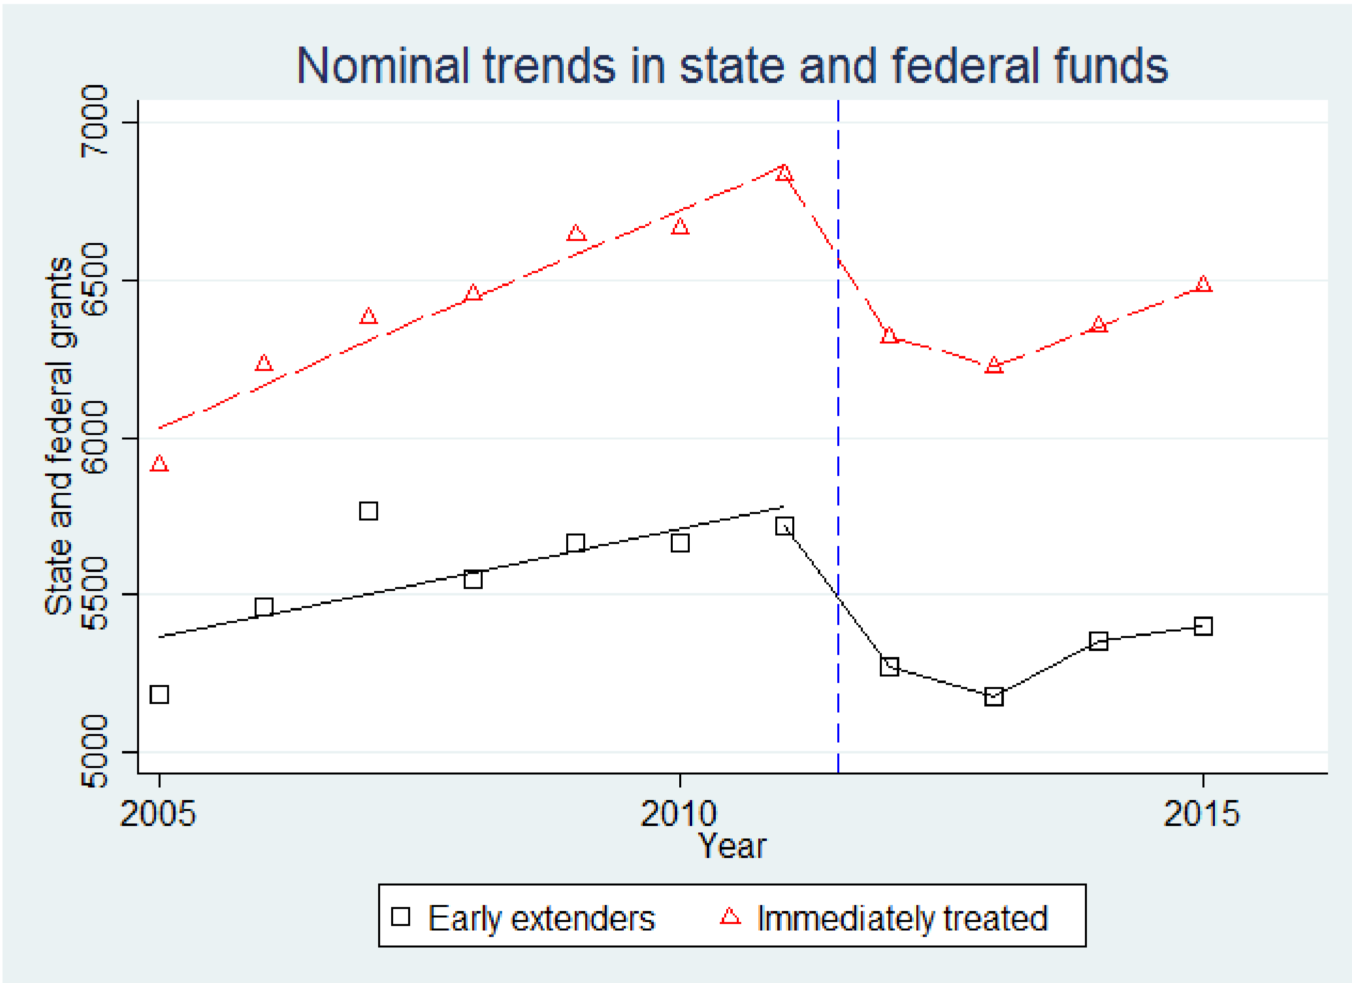

Figure shows the trends in externally sourced (state+Federal) funds for two types of districts over time. "Immediately treated" refers to those districts which were required to renegotiate a contract in the Post Act 10 regime beginning in 2012. "Early extender" districts are those which have a contract in place prior to Act 10's announcement, which expired in 2013 or later. Points on the graph are means across districts (district weighted) within each group. This Figure investigates whether there is a distinct shock to external funds coincident with Act 10. While both types of districts receive a fall in external funds, this fall is roughly equal in each group. Early extender district also have a slower growth trend in external fund. Refer to Table 1.2 for a formal difference-in-differences test of this characteristic.

Figure 1.6: Parallel trend comparison: student teacher ratio

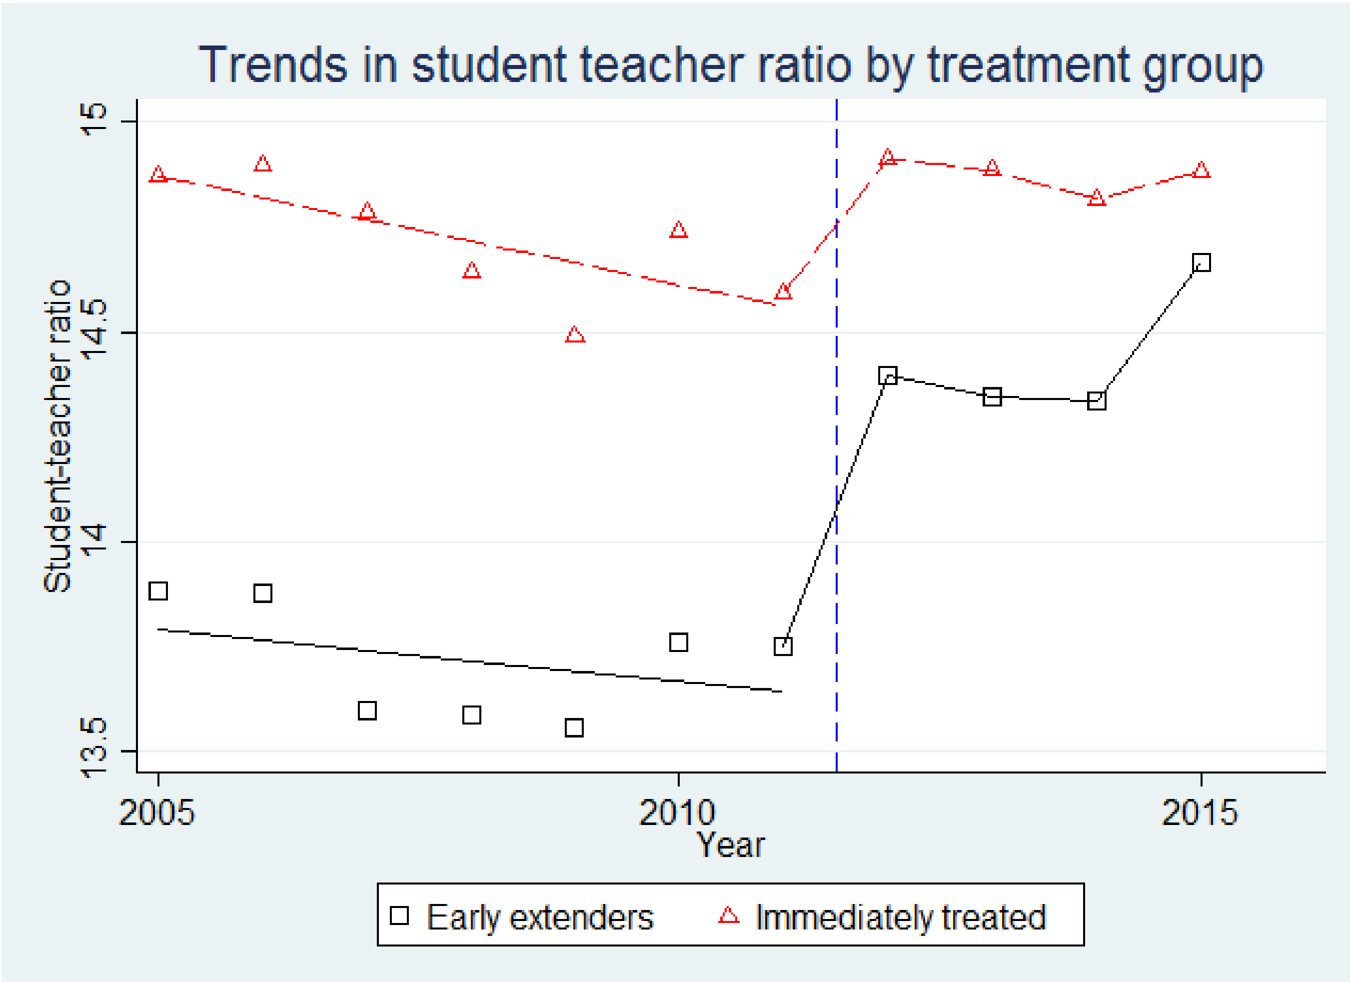

Figure shows the trends in the student-teacher ratio for two types of districts over time. "Immediately treated" refers to those districts which were required to renegotiate a contract in the Post Act 10 regime beginning in 2012. "Early extender" districts are those which have a contract in place prior to Act 10's announcement, which expired in 2013 or later. Points on the graph are means across districts (district weighted) within each group. Prior to Act 10, both districts had a similar trend in student-teacher ratios, although earlier extenders had fewer students per teachers. Figure 1.6 shows that early extender districts had a slightly larger jump in the student-teacher ratio following Act 10, which suggests that layoffs were necessary to balance district budgets while salaries remained high, although this effect is not statistically significant.

Figure 1.7: Event study: Act 10 on log teacher compensation

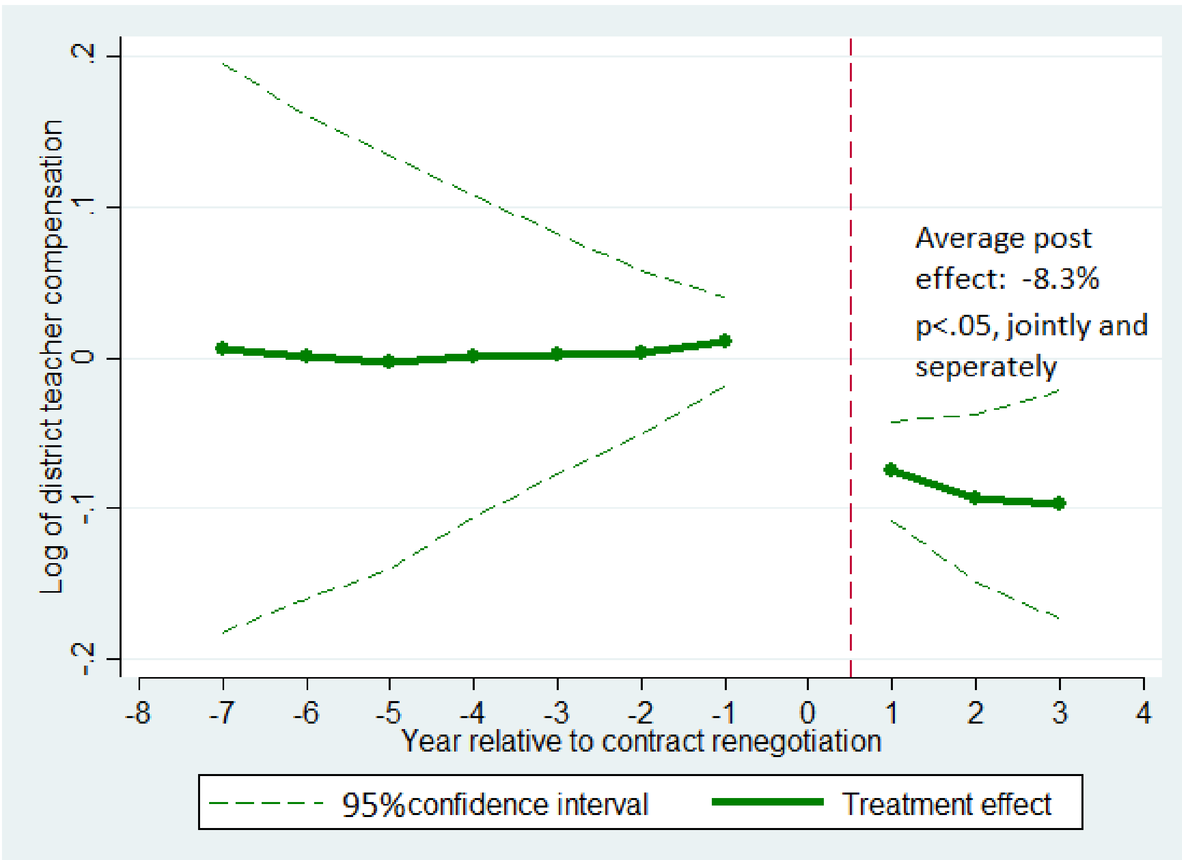

Figure shows the estimated results of equation 1.1 on log teacher compensation. Point estimates and confidence intervals are included for each period relative to the year immediately prior to treatment. All pre-treatment coefficients are jointly insignificant (F statistic .23) while all post-treatment coefficients are jointly significant (F statistic 12.8). In terms of total compensation, differences are steady and minor in the pre period, while districts see a large (9-10%) drop in pay immediately following Act 10 contract renegotiation. Standard errors clustered by district.

Figure 1.8: Event study: Act 10 on level of teacher compensation

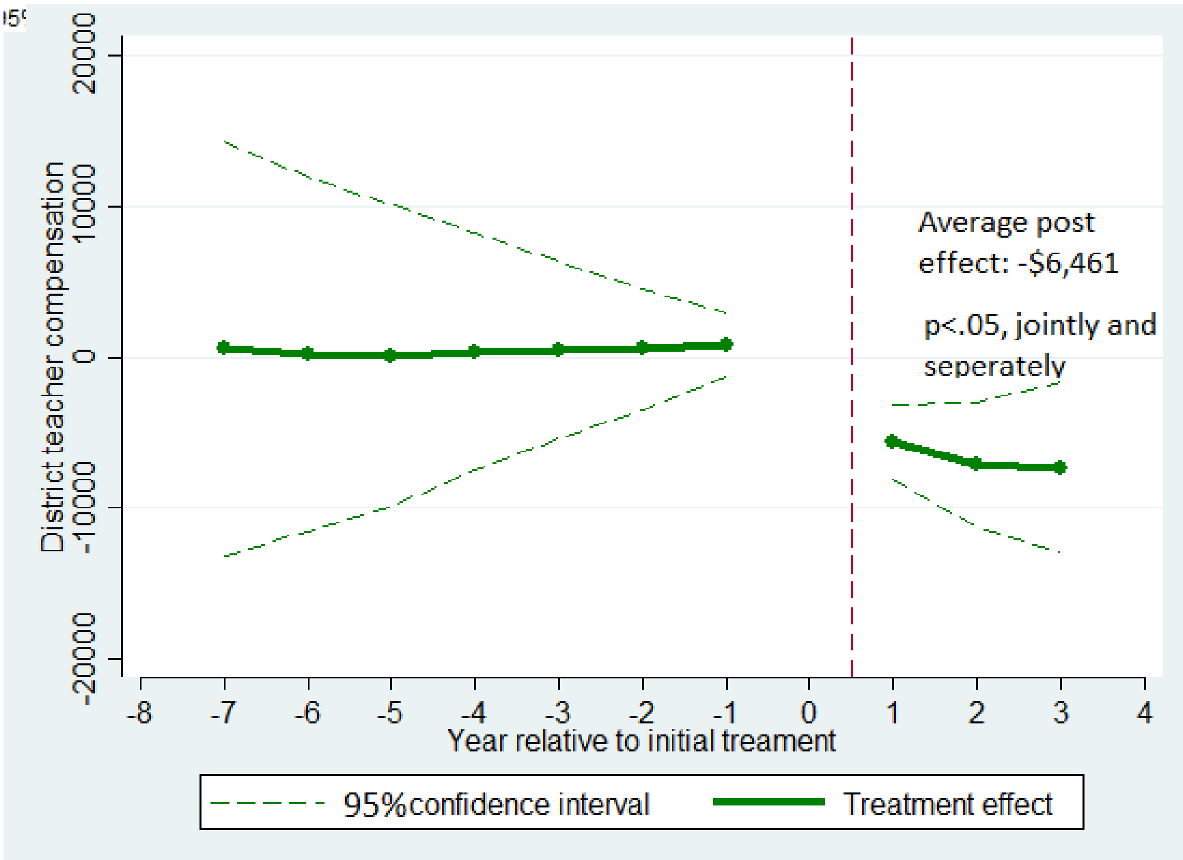

Figure shows the estimated results of equation 1.1 on the dollar value of teacher compensation. Point estimates and confidence intervals are included for each period relative to the year immediately prior to treatment. All pre-treatment coefficients are jointly insignificant (F statistic .23) while all post-treatment coefficients are jointly significant (F statistic 12.8). In terms of total compensation, differences are steady and minor in the pre period, while districts see a large (9-10%) drop in pay immediately following Act 10 contract renegotiation. Standard errors clustered by district.

Figure 1.9: Pre-trends in compensation by quartile of 2012 fiscal shock

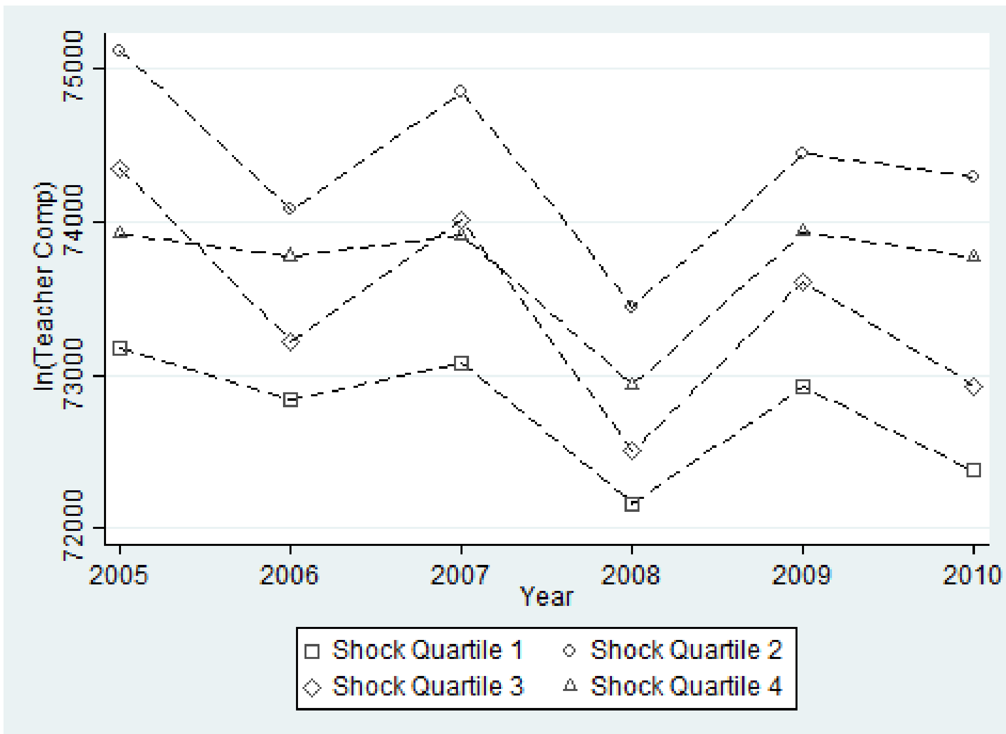

Figure shows the trends in real teacher compensation between 2005 and 2010, grouping districts into 4 seperate groups based on the change in external (state+Federal) grants between 2011 and 2012. This Figure seeks to identify strong secular trends by each group before using this categorization to check for heterogeneous treatment effects using the event study analysis. The fourth quartile, which had the smallest decline in external grants, does appear to be on a different trajectory than the other quintiles, however this change of $1,000 real dollars in compensation over 5 years is unlikely to substantially bias results.

Figure 1.10: Synthetic control analysis trend in total teacher compensation

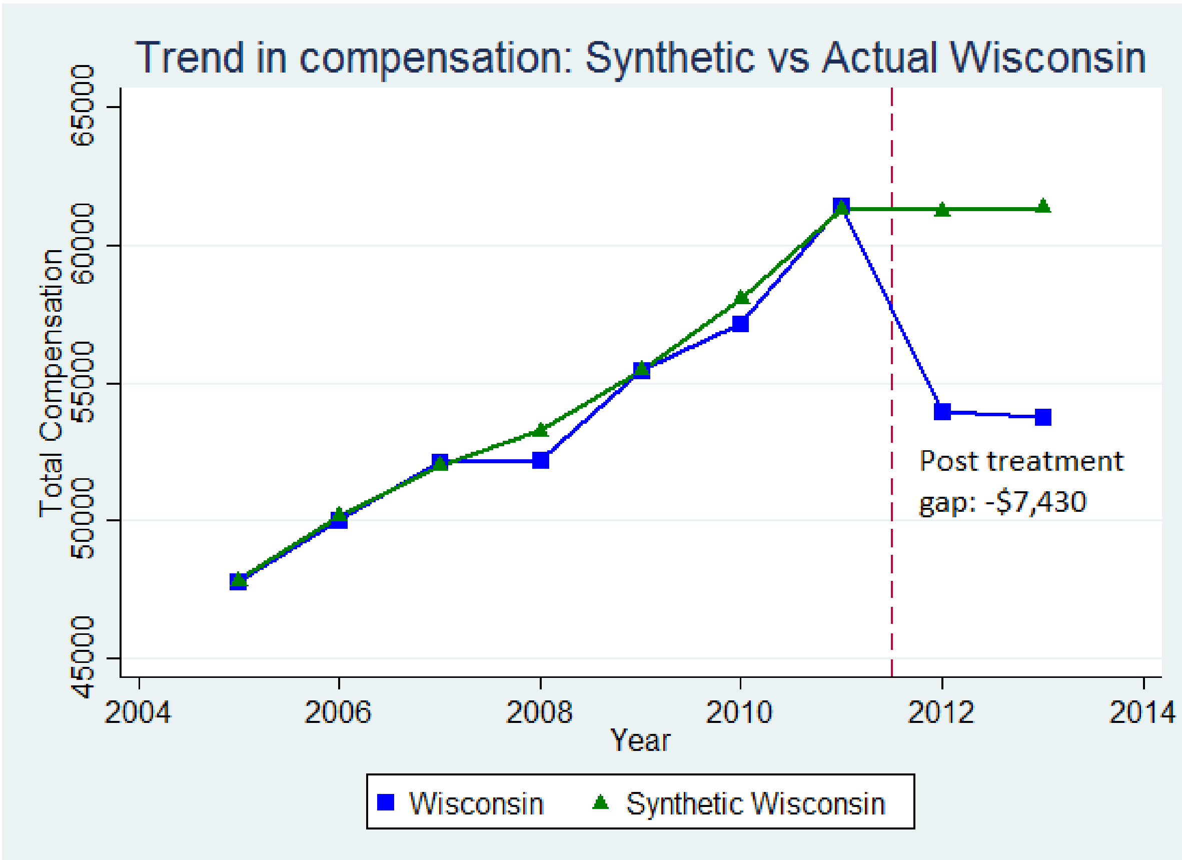

Figure shows the time series in teacher compensation in Wisconsin and in synthetic Wisconsin. Synthetic Wisconsin closely tracks total compensation in Wisconsin until the treatment year in 2012, at which point a gap of $7,400 emerges. Matching properties are based on teacher compensation in odd years from 2005-2011, the state trend in teacher compensation, special needs enrollment, race and ethnicity of the state student population, the share of students in the state attending rural school districts, and the share of students in the state attending urban districts. Compensation figures are taken from the Common Core of Data. There are multiple measures of teacher compensation in the Common Core of Data. I show here a measure which is consistently lower than the Wisconsin administrative data, but is available for all states from 2005-2013. Results show in figures 1.10 and 1.11 are robust to other measures which more closely track administrative data but are unavailable for all states over this period.

Figure 1.11: Synthetic control analysis trend in total teacher compensation

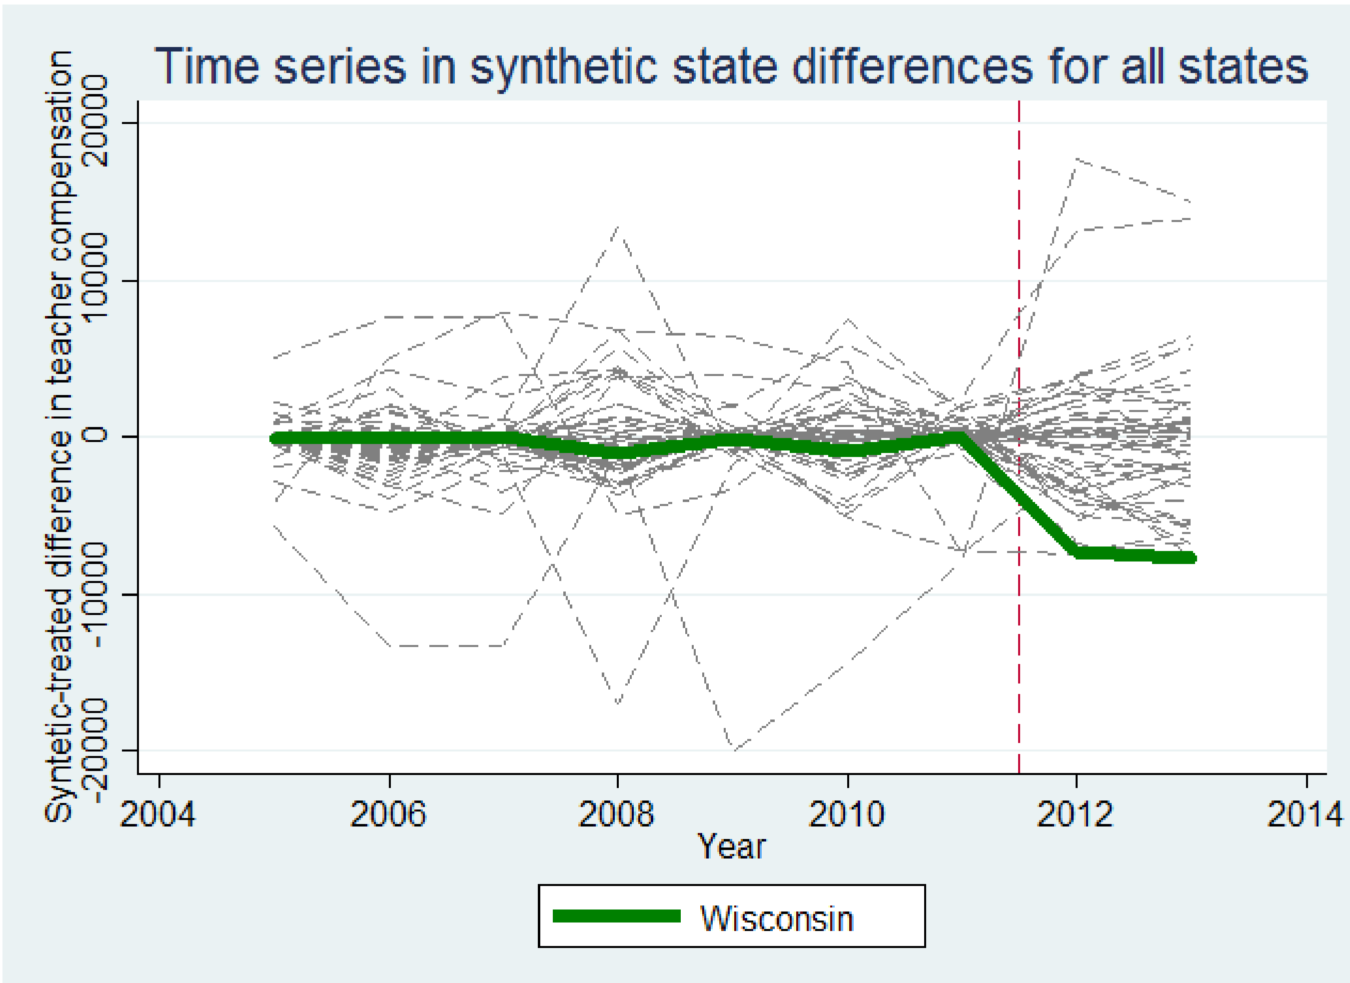

Table shows the simulated synthetic control gaps for 49 states (excluding New York, which is too high at baseline to form a match), asserting a treatment for each state in 2012. Wisconsin, as shown, generates the largest post-period gap between the treatment state and synthetic control. Additionally, it is one of the only cases to show the anticipated effect of an immediate and persistent drop beginning in 2012.

1.9 Tables

Table 1.1: Balance in district characteristics across treatment groups

District Characteristic	Immediately Treated	Early Extenders	Differences
Share of rural districts	.491	.375	-.116
	(.5)	(.5)	(.127)
Share of black and hispanic students	.0691	.0804	.0113
	(.0826)	(.0761)	(.021)
Share of economically distressed students	.329	.393	.0636*
	(.15)	(.169)	(.0385)
Federal grants (per student)	1,186	1,305	119
	(658)	(716)	(168)
State grants (per student)	5,529	4,377	-1,152***
	(2,000)	(2,667)	(517)
Property value (per student)	815,881	1,326,870	510,989*
	(1,038,198)	(1,478,939)	(269,531)
Total enrollment	2,064	1,584	-480
	(5,046)	(2,424)	(1,268)

Table shows mean values in selected district characteristics, comparing districts immediately treated by Act 10 in 2012 to those which had a longer contract in place. Cell values are computed by taking each district average in the three years prior to Act 10 (2009-2011) and averaging them with each treatment category. Standard deviations in parentheses in colums 1 and 2. Column 3 shows the difference in means, and the standard error on the regression coefficient comparing these two groups. Districts are fairly similar across most observable traits. Early extenders tend to be somewhat smaller and wealthier, and thus receive less state aid. However both groups experience similar fiscal shocks in external funds between 2011 and 2012.

*** p<.01, ** p<.05, * p<.10

Table 1.2: Double difference in district characteristics

Covariate	Initially Treated Districts	Early Extenders	Difference
State and federal grants per student	6,715 (2,261)	5,682 (3,019)	1,033
Revenue limit per student	10,355 (1,284)	10,974 (2,173)	-618
Property value per student	815,853 (1,039,888)	1,326,870 (1,448,593)	-511,017
Enrollment	2,064 (5,041)	1,584 (2,371)	480

	Initially Treated Districts	Early Extenders	Difference	Difference in differences
State and federal grants per student	6,276 (2,196)	5,224 (2,811)	1,052	19 (527)
Revenue limit per student	10,116 (1,305)	10,848 (2,365)	-733	-114 (312)
Property value per student	816,059 (1,078,512)	1,324,343 (1,461,601)	-508,284	2,734 (24,9717)
Enrollment	2,058 (4,945)	1,556 (2,346)	502	22 (1,147)

This tables tests for differential shocks in district characteristics concurrent with Act 10. Particularly important are external (state+Federal) grants, and district revenue limits, both of which were changed in the 2012 budget. I estimate the values for the two years immediately prior to and following Act 10, for the immediately treated and early extender group districts. Standard deviations by groups are in parentheses, and standard errors on the differences in differences shown in column 4. Based on the difference-in-differences summary in the bottom right, there is no differential shock per group.

Table 1.3: Balance in reweighted district characteristics across treatment groups

District Characteristic	Immediately Treated	Early Extenders	Differences
Share of rural districts	.487	.375	-.112
	(.497)	(.699)	(.129)
Share of black and hispanic students	.0695	.0625	-.00707
	(.0841)	(.0624)	(.0213)
Share of economically distressed students	.332	.325	-.00696
	(.159)	(.272)	(.042)
Federal grants (per student)	1,192	1,119	-73
	(705)	(918)	(182)
State grants (per student)	5,485	4,638	-847
	(1,935)	(5,171)	(545)
Property value (per student)	837,585	906,875	69,290
	(1,138,496)	(1,052,236)	(289,469)
Total enrollment	2,046	1,894	-152
	(4,890)	(4,005)	(1,239)

Table shows mean values in selected district characteristics, comparing districts immediately treated by Act 10 in 2012 to those which had a longer contract in place. This Table differs from Table 1.1 as districts have been reweighted based on the propensity of being an early extended district. Cell values are computed by taking each district average in the three years prior to Act 10 (2009-2011) and averaging them with each treatment category. Standard deviations in parentheses in columns 1 and 2. Column 3 shows the difference in means, and the standard error on the regression coefficient comparing these two groups. Following the reweighting process, the two groups of districts are much more comparable on observed characteristics.
*** p<.01, ** p<.05, * p<.10

Table 1.4: Robustness of average effect of Act 10 on log teacher compensation

	Model					
Outcome	(1)	(2)	(3)	(4)	(5)	(6)
ln(Teacher Comp)	-0.084***	-0.077***	-0.083***	-0.084***	-0.074***	-0.087***
	(0.014)	(0.018)	(0.016)	(0.019)	(0.021)	(0.019)
ln(Teacher Salary)	-0.045***	-0.039***	-0.041***	-0.047***	-0.036***	-0.043***
	(0.010)	(0.012)	(0.009)	(0.009)	(0.012)	(0.009)
ln(Teacher Benefits)	-0.171***	-0.160***	-0.179***	-0.163***	-0.156***	-0.189***
	(0.035)	(0.040)	(0.042)	(0.045)	(0.048)	(0.064)
N	4,180	4,180	5,432	4,180	4,180	5,432
Years	2006-15	2006-15	2004-16	2006-15	2006-15	2004-16
Year FE	x	x	x	x	x	x
District FE	x	x	x	x	x	x
District-trends		x	x		x	x
District characteristi	x	x		x	x	
Reweighted				x	x	x

Table shows the effect of Act 10 on log teacher compensation (reporting the coefficient on *PostAct10* from equation 1.2, following a district's first exposure to the law based on contract expiration). Standard errors in parentheses. Each cell is a separate regression, where rows are outcomes and columns are specifications. All regressions include district and year fixed effects. Columns 2, 3, 5 and 6 add district-specific linear trends. Columns 4-6 reweight observations by rural status, race, enrollment, property values, and external grants. Columns 3 and 6 remove district-characteristics to include observation-years where these are not available. Standard errors clustered at the district level. Results are additionally robust to dropping urban districts, including observable characteristics, and including recertification election loss measures. Overall, salary effects are about 4%, benefits 20 and total compensation 8.5%.
*** p<.01, ** p<.05, * p<.10

Table 1.5: Robustness of average effect of Act 10 on teacher compensation levels

	Model					
Outcome	(1)	(2)	(3)	(4)	(5)	(6)
Teacher Comp	-6,179.265***	-5,526.109***	-6,272.434***	-6,229.786***	-5,214.449***	-6,461.641***
	(1,035.203)	(1,304.152)	(1,151.916)	(1,376.261)	(1,630.930)	(1,384.556)
Teacher Salary	-2,272.236***	-1,894.496***	-2,134.417***	-2,531.384***	-1,764.879***	-2,248.094***
	(451.719)	(561.662)	(453.310)	(411.225)	(568.341)	(478.605)
Teacher Benefits	-3,907.029***	-3,631.613***	-4,138.016***	-3,698.403***	-3,449.571***	-4,213.547***
	(849.501)	(977.014)	(908.419)	(1,094.364)	(1,194.925)	(1,308.607)
N	4,180	4,180	5,432	4,180	4,180	5,432
Years	2006-15	2006-15	2004-16	2006-15	2006-15	2004-16
Year FE	x	x	x	x	x	x
District FE	x	x	x	x	x	x
District-trends		x	x		x	x
District characteristics	x	x		x	x	
Reweighted				x	x	x

Table shows the effect of Act 10 on teacher compensation levels (reporting the coefficient on *PostAct10* from equation 1.2, following a district's first exposure to the law based on contract expiration). Standard errors in parentheses. Each cell is a separate regression, where rows are outcomes and columns are specifications. All regressions include district and year fixed effects. Columns 2, 3, 5 and 6 add district-specific linear trends. Columns 4-6 reweight observations by rural status, race, enrollment, property values, and external grants. Columns 3 and 6 remove district characteristics to include observation-years where these are not available. Standard errors clustered at the district level. Results are additionally robust to dropping urban districts, including observable characteristics, and including recertification election loss measures. Overall, compensation effects are about \$6,500, with over \$4,000 over that coming out of benefits.
*** p<.01, ** p<.05, * p<.10

Table 1.6: Distributional effect of Act 10 by predicted wage (level regression results)

	Compensation Quintile				
Outcome	(1)	(2)	(3)	(4)	(5)
Teacher Comp	-2,773**	-2,611*	-5,183***	-6,008***	-6,801***
	(1,109)	(1,552)	(1,417)	(1,873)	(1,696)
Teacher Salary	-490	80	-1,684***	-2,060***	-2,707***
	(756)	(688)	(633)	(671)	(807)
Teacher Benefits	-2,282***	-2,691**	-3,562***	-3,947***	-4,093***
	(810)	(1,149)	(1,194)	(1,324)	(1,413)
N	4,180	4,180	4,180	4,180	4,180
Years	2006-15	2006-15	2006-15	2006-15	2006-15
Year FE	x	x	x	x	x
District FE	x	x	x	x	x
District-trends	x	x	x	x	x
District characteristics	x	x	x	x	x
Reweighted	x	x	x	x	x

Table shows the effect of Act 10 by predicted compensation quintile. Each column limits the sample to a different predicted wage quintile, reporting the coefficient on *PostAct10* from equation 1.3. Predicted compensation is determined by out of sample prediction from 2003 payroll data. The prediction equation regresses total compensation on teacher experience and education, and their interaction. Results are shown for teacher-weighted regressions, though results are robust to regressing on district-level averages within each predicted quintile. Standard errors are clustered at the district level. All regressions include district and year fixed effects, district trends, and reweight observations to match off-cycle districts. Each cell is a separate regression, using different outcome variables by row.
*** p<.01, ** p<.05, * p<.10

Table 1.7: Distributional effect of Act 10 by predicted wage (log regression results)

	Compensation Quintile				
Outcome	(1)	(2)	(3)	(4)	(5)
ln(Teacher Comp)	-0.050**	-0.037*	-0.069***	-0.078***	-0.089***
	(0.022)	(0.021)	(0.020)	(0.024)	(0.022)
ln(Teacher Salary)	-0.010	0.004	-0.033**	-0.040***	-0.049***
	(0.022)	(0.015)	(0.015)	(0.014)	(0.017)
ln(Teacher Benefits)	-0.130***	-0.124***	-0.148***	-0.166***	-0.190***
	(0.047)	(0.047)	(0.049)	(0.051)	(0.057)
N	4,180	4,180	4,180	4,180	4,180
Years	2006-15	2006-15	2006-15	2006-15	2006-15
Year FE	x	x	x	x	x
District FE	x	x	x	x	x
District-trends	x	x	x	x	x
District characteristi	x	x	x	x	x
Reweighted	x	x	x	x	x

Table shows the effect of Act 10 by predicted compensation quintile. Each column limits the sample to a different predicted wage quintile, reporting the coefficient on *PostAct10* from equation 1.3. Predicted compensation is determined by out of sample prediction from 2003 payroll data. The prediction equation regresses total compensation on teacher experience and education, and their interaction. Results are shown for teacher-weighted regressions, though results are robust to regressing on district-level averages within each predicted quintile. Standard errors are clustered at the district level. All regressions include district and year fixed effects, district trends, and reweight observations to match off-cycle districts. Each cell is a seperate regression, using different outcome variables by row.
*** p<.01, ** p<.05, * p<.10

Table 1.8: Balance in observable characteristics by fiscal shock

	Quartile of 2012 aid cut			
District Characteristic	(1)	(2)	(3)	(4)
Share of rural districts	.489	.404	.467	.587
	(.5)	(.493)	(.501)	(.495)
Share of black and hispanic students	.0716	.0733	.0691	.064
	(.101)	(.0747)	(.0776)	(.0736)
Share of economically distressed students	.324	.305	.347	.35
	(.133)	(.138)	(.15)	(.178)
Federal grants (per student)	1,232	1,114	1,174	1,244
	(439)	(309)	(362)	(1,153)
State grants (per student)	6,283	5,853	5,592	4,202
	(1,339)	(1,473)	(1,939)	(2,562)
Change in state+federal grants, 2011-2012	-1,042	-610	-364	-28.9
	(1,116)	(65.5)	(77.3)	(163)
Property value (per student)	561,316	602,611	806,987	1,373,755
	(403,590)	(322,008)	(1,103,519)	(1,622,145)
Total enrollment	2,483	2,367	2,201	1,124
	(8,588)	(3,142)	(3,566)	(1,282)
Revenue Limit (per student)	10,299	10,078	10,129	11,014
	(990)	(738)	(1,124)	(1,841)

Table shows the balance in observable characteristics after sorting districts by the size of their change in external (Federal+state) grants between 2011 and 2012. Quartile 1 had the largest (negative) change in external grants, which quartile 4 had the smallest (in some cases, positive) change. The patterns suggest there is substantial selection in these categories, and that state grants were cut least for the wealthiest districts. These districts are the least reliant on state grants, and fund operations primarily through property tax revenues.

Table 1.9: Treatment in teacher compensation by fiscal shock

	Quartile of Change in External Grants			
Outcome	(1)	(2)	(3)	(4)
ln(Teacher Comp)	-0.079***	-0.077***	-0.076***	-0.071***
	(0.023)	(0.022)	(0.023)	(0.022)
ln(Teacher Salary)	-0.032**	-0.032**	-0.026*	-0.024
	(0.016)	(0.015)	(0.015)	(0.014)
ln(Teacher Benefits)	-0.180***	-0.172***	-0.179***	-0.168***
	(0.046)	(0.045)	(0.046)	(0.045)
N	1180	1180	1180	1180
Years	2006-15	2006-15	2006-15	2006-15
Year FE	x	x	x	x
District FE	x	x	x	x
District-trends	x	x	x	x
District characteristics	x	x	x	x
Reweighted	x	x	x	x

Table shows heterogeneity in the effect of Act 10 after sorting districts based on the change in external (Federal+state) grants, between 2011 and 2012. Each cell is a separate regression. Rows describe different outcomes and columns describe different sampling groups. Quartile 1 had the largest (negative) change in external grants, while quartile 4 had the smallest (in some cases, positive) change. Each regression compares the quartile of initially treated districts to the universe of off-cycle districts, reporting the coefficient on $PostAct10$ from equation 1.2. Standard errors in parentheses, clustered by district. Table 1.9 shows that wage effects are remarkably consistent across each category, despite differences in initial pay levels and district wealth.
*** p<.01, ** p<.05, * p<.10

Table 1.10: Weightings for synthetic Wisconsin

State	Weight
California	0.106
District of Columbia	0.012
Delaware	0.193
Indiana	0.09
Nebraska	0.291
New Hampshire	0.286
Rhode Island	0.022

Tables summarizes the weights assigned to other states used to construct a synthetic Wisconsin. Matching properties are based on teacher compensation in odd years from 2005-2011, the state trend in teacher compensation, special needs enrollment, race and ethnicity of the state student population, the share of students in the state attending rural school districts, and the share of students in the state attending urban districts.

Table 1.11: Characteristics of Wisconsin and Synthetic Wisconsin

	Wisconsin	All other states	Synthetic WI
Teacher compensation (Average pre)	49,671	41,188	50,145
Teacher compensation (2005)	47,763	39,281	47,826
Teacher compensation (2007)	52,106	42,552	51,998
Teacher compensation (2009)	55,428	45,876	55,478
Teacher compensation (2011)	61,436	47,809	61,312
Share of special needs students	0.144	0.136	0.145
Share of black students	0.102	0.155	0.118
Share of Hispanic Students	0.072	0.122	0.125
Share of local-sourced Revenues	0.428	0.406	0.458

Table summarizes the treatment and covariates in Wisconsin, the unconditional average for other states, and synthetic Wisconsin. Summary demonstrates that Wisconsin teachers are paid more than teachers in other states, and Wisconsin has a greater share of special needs students. Compensation figures are taken from the Common Core of There are multiple measures of teacher compensation in the Common Core of Data. I show here a measure which is consistently lower than the Wisconsin administrative data, but is available for all states from 2005-2013. Results show in figures 1.10 and 1.11 are robust to other measures which more closely track administrative data but are unavailable for all states over this period.

CHAPTER II

Does the Stimulus Stick? The Role of Federal Spending on Local Hiring

2.1 Introduction / Motivation

In this paper, I investigate the effects of a significant, but temporary, expansion of of Federal grants on local government spending. Specifically, I identify the causal effects of Federal grants from the American Recovery and Reinvestment Act ("Recovery Act", "stimulus", or "ARRA") of 2009, which saw the Federal government increase its role in education spending by 50% over a four year period. I do this by exploiting non-linearities in the formula allocation of Federal grants, which generates cross-sectional variation in grants across local school districts. Because most stimulus grants are distributed by pre-existing formulas, they are unlikely to be endogenous to local economic conditions school districts.

This expansion of Federal education grants was part of a larger anti-recessionary measure. Proponents of the Recovery Act argued that it had two main goals. First, it sought to increase total employment by increasing government spending to offset reduced private demand. Second, to the extent the Federal government was increasing total spending, it sought to do so in a manner consistent with long term growth, by investing in public goods.

The setting of this analysis speaks to two important empirical questions. First, is it possible for the Federal government to "stimulate" local economic conditions? Second, to what extent do intergovernmental grants influence spending patterns?

The first question is related to the "flypaper effect", an empirical puzzle first identified by Gramlich et al. (1973), whereby intergovernmental grants tend to increase spending by more than an equivalent amount of private income. That is, because government funds are fungible, we might expect that intergovernmental grants would simply crowd out other funds but not dramatically shift the overall distribution of spending. Instead, we tend to see these grants "stick where they hit."

A useful overview of the Flypaper Effect can be found in Hines and Thaler (1995). Estimates can vary widely - a dollar in grants can increase sub-government spending by anywhere from $.25 to $1. Across a variety of settings, however, the literature has consistently found that intergovernmental grants increase spending by amounts greater than would have been found by increasing private income. Understanding the magnitude of response to short term education grants is thus key to understanding the effectiveness of stimulus education grants.

A variety of papers have looked at the flypaper effect of education grants. These are typically rooted in one of two settings. The first is expansions of Federal programs, typically related to Title I. Title I is a Federal program designed to provide assistance with low income students. Gordon (2004) looked at the effects of Title I grants on school district spending by exploiting changes in the allocation of grants driven by Census measurement updates, and found that while grants lead to short term increases in spending, local governments tended to offset the changes in Federal grants within three years. Cascio, Gordon and Reber (2013) looked at the effects of Federal grants immediately following the introduction of Title I from the Elementary and Secondary Education Act of 1965. They found Federal grants increased local spending by anywhere from $.46 to $1 for each dollar or additional greats, depending on the local government's scope of offset. The differences in these estimates show that the exact setting and source of variation can be important for the response of local governments to Federal grants. This fact important in assessing the ability of the stimulus to increase local education spending, because ARRA represents the largest short term increase in Federal education spending to date, and previous estimates based on long term changes may not be informative.

The second group of papers which have estimated the education flypaper effect are based on state finance reforms. Throughout the 70s, 80s, and 90s several state governments transitioned from financing K-12 schools primarily at the local level, to placing a greater emphasis on state funding. These reforms often had the intent to equalize per student spending across local schools within a state by increasing the state's level of financial support. The literature includes papers by Card and Payne (2002), which looks at court mandated reforms across the US, and Hyman (Forthcoming), which does a deep dive exploring cutoffs in the state grant formula in Michigan. Both find that a dollar in additional state grants increase total (all-source) student spending by 50 to 60 cents. While this paper focuses on changes in Federal policy, the stimulus let states have some discretion in allocation Federal greats, which makes the state-local relationship an important part of the equation in evaluating this policy.

Even if increased Federal grants are successful in increasing student spending, it is not clear to to what extent this will succeed in increasing employment. The literature on stimulus program evaluation has found wide-ranging estimates of the cost per additional job-year. On the low end, Chodorow-Reich et al. (2012) has found that the Recovery Act was able to create an additional job-year with each \$26,000, by exploiting state level variation in additional Medicaid grants. Towards the mid-range of estimates, Wilson (2012) assessed a variety of programs and found that it took \$125,000 in Federal spending to create an additional job year. On the higher end, Feyrer and Sacerdote (2011) found a cost of \$170,000 per job-year (from all sources of stimulus funding), and a cost of \$567,000 from education-based funding.

The range of estimates in the stimulus program evaluation literature emphasize that different programs in the Recovery Act have different local average treatment effects. In particular, education grants have previously been considered to be one of the least effective programs, at least as far as increasing employment.

This paper has two key contributions. First, I evaluate the flypaper effect of a large, short term increase in Federal spending. The education flypaper literature to date has primarily followed permanent changes, which are likely to have different dynamic responses than temporary changes. Second, I exploit local variation in analyzing the effectiveness of stimulus funds. I accomplish this by exploiting three key formulas which drove most of the ARRA education grants. Most prior efforts to evaluate the Recovery Act have relied on either time series of state-level cross-sectional variation. With this approach, I find an average stimulus induced job-year cost of \$108,000, and I find that the expansion of permanent programs (Title I, Special Education) had a large flypaper effect, whereas the temporary program (the Stabilization Fund) had none.

The remainder of this paper proceeds as follows. In section 2.2 I discuss background on funding for K-12 education and the Recovery Act itself. Section 2.3 discusses the data used in this analysis. In section 2.4 I discuss my identification strategy. In section 2.5 I discuss results. Section 2.6 concludes.

2.2 Background

In this section, I provide the reader with institutional context to better understand how the Recovery Act changed the Federal, state, and local financial relationships in funding K-12 education. This background informs the identification strategy I describe in section 2.4. First, I describe the funding system for K-12 education in the US, prior to the enactment of the Recovery Act. I then describe the Recovery Act itself. I end this section with a discussion of the related Federal policy changes surrounding the Recovery Act.

2.2.1 The Federal Role in K-12 Education

Funding for K-12 education in the US is primarily at the state and local level. Prior to 2009, Federal funds made up about 8% of K-12 spending, with the remainder roughly evenly split between state and local sources. Average per pupil spending in this period was about

$11,000 per student. See Figure 2.1 for a summary of both total per pupil revenue, and share of spending, by government source.

Historically, Federal funding for K-12 education is primarily driven by two programs: Title I and Special Education. These programs make up two thirds of total Federal K-12 spending.[1] Remaining programs are a mix of competitive teacher incentive grants, aid for schools impacted by military bases, ESL programs, and a variety of other small programs. The Title I and Special Education programs are important to this analysis not only due to their size and scope, but also because these grants are determined by formula, and the amounts received by a given district are outside of its ability to control.

Title I grants are targeted toward school districts serving low income students. This program was created as a part of the Elementary and Secondary Education Act of 1965. In practice, Title I is actually four distinct grant programs, which distribute funds to a district based on their poverty, enrollment, and the average per pupil expenditures of the state.[2] In my sample, 85% of the districts in 2007 received Federal Title I grants. The top 10% of school districts, ranked by the size of Title I grants received in 2007, make up about 35% of the program.

Special Education grants are funds provided to support equal education access for all individuals with disability in the United States. These grants are provided from the Federal Government to the state level, and states must in turn determine a process for allocating these to school districts. While there is variation in how states allocate these grants to districts, typically it is a function of the share of students with special needs attending that district [see Cullen (2003)]. In my sample 70% of school districts received Special Education grants. The top 10% of school districts, ranked by the size of Title I grants received in 2007, make up about 45% of the program.

[1] See US Department of Education (2005)

[2] The full details surrounding the determination of Title I funds are extensive, but are detailed in National Center for Education Statistics (2016). The formula has multiple kink points for different grant types, as well as a hold harmless provision which prevents a district from losing too much in the way of funding from year to year.

2.2.2 The American Recovery and Reinvestment Act

The American Recovery and Reinvestment Act was a policy response to the Great Recession. Between February 2008 and February 2010, the US economy shank by 5.1% of GDP, and lost approximately 8.7 million jobs. In the middle of this contraction, the US government transitioned presidencies, from the Administration of George W. Bush to that of Barack Obama. The Recovery Act, signed in to law in February 2009, in the midst of the great recession, was the first major policy initiative of the Obama Administration.

The primary goal of the Recovery Act was to increase government spending as counter-cyclical fiscal policy. Proponents of the law, however, pursued spending not just for the sake of spending, but targeted programs which were meant to be consistent with long term growth. For example, as described by Barack Obama, "We will act not only to create new jobs, but to lay a new foundation for growth", or then-chair of the National Economic Counsel, Larry Summers, "The Obama plan represents not new public works, but, rather, investments that will work for the American public. Investments to build the classrooms, laboratories, and libraries our children need to meet 21st-century education challenges."

The Recovery Act thus had a twin set of goals: to spend money quickly, while also targeting programs which can be seen as long term investments. This mix of priorities is reflected in the actual breakdown of spending in the program, shown in Table 2.1. Of the roughly $800 billion in spending, slightly more than one third targeted tax benefits. 13% (or 20% of the remainder) was spent on education, which is the largest category of spending outside of tax cuts, outpacing spending on Medicaid grants to states, infrastructure spending, and increased spending on supplemental income programs.

Of the $104 billion in new education spending, I bring particular focus on three main programs. The first two are existing programs, discussed in section 2.2.1 above: Title I and Special Education. The third was a new program created entirely within the Recovery Act: the State Fiscal Stabilization Fund (or "SFSF"). The relative sizes of these programs can be seen in Table 2.1. Taken together, these three programs make up $74 billion of the $104 billion in total education grants in the Recovery Act. The Title I and Special Education increases were roughly the size of a full year of appropriations for these two programs. The Stabilization Fund was the largest of the three, slightly more than the Title I and Special Education expansions combined.

The new program created by the Recovery Act, the State Fiscal Stabilization Fund, was intended as way to get money to state and local governments as quickly as possible. Of the $51 million included in Table 2.1, roughly $40 million was meant to be for higher and K-12 education, whereas the remainder could be spent either on education or public safety. These funds were allocated to states based on the share of the state's total population aged between 5 and 24.

All three major education programs came with a certain degree of flexibility afforded to the state governments. First, and perhaps most importantly, the state governments had discretion with regards to when to release these funds to local school districts. Special Education and Title I funds were released to states on April 1 and September 30, 2009, spanning the 2009 and 2010 fiscal years (at the state and local level, which typically ends on June 30). However all funds had until September 30, 2011 to be obligated to local districts–effectively, the end of the 2011 fiscal year.

The Stabilization Fund had even more state-level flexibility than the other programs. As with Title I and Special Education, it had an obligation window which lasted until September 30, 2011. It also gave states the flexibility to spend all funds on higher education, provided the state meets its Maintenance of Effort requirements.[3] Further complicating the matter, the Stabilization Fund has a twin set of mandates. First, to spend funds in order to avoid layoffs. Second, to spend funds strategically so as to avoid a "cliff" when they were no longer available. These two goals were in direct conflict with one another. A state such as

[3] Because one of the policy goals of the Recovery Act was to engage in counter-cyclical spending, the Federal government was somewhat lax in terms of actually releasing these funds to states. In concept, there was a Maintenance of Effort provision which required that districts maintain spending at 2006 levels for both K-12 and higher education. In practice, most states met this standard easily, or received a waiver from this requirement. Thus, or this total, almost all was received by states and spend by school districts.

Wisconsin, for example, used its Stabilization Funds to avoid layoffs and to continue to pay teachers at current levels. This lead to a significant shortfall when the funds dried up after the 2011 school year.

While the programs I describe above are the largest portion of the Recovery Act dedicated to K-12 education, they are not the only ones. Smaller competitive grants were made available for school improvements, investments in information technology, and perhaps most importantly, the Race to the Top program. Race to the Top was a $4 billion competitive grant program, smaller than the other programs I discuss, but often credited with inducing a large amount of change across school districts. It evaluated state-level programs to improve teacher effectiveness, adopt common standards, turn around low performing schools, and build and use data systems.

2.2.3 Related Policy Changes

As the window of my analysis covers several years, it is important to keep in mind changes in Federal education policy operating outside of the Recovery Act. In particular, 2006-2012 saw several changes in government control, which had a direct impact on Federal K-12 funding. Beginning in 2007, Democrats took control of the House and the Senate, following a period of united Republican control. The first full budget following this, for the 2008 fiscal year, saw Title I funding jump by 8%. Beginning in 2009, Democrats had larger control of both houses, as well as the presidency. In this time frame, the normal appropriation increased by 5%. After Republicans took back control of the House in 2011, Federal K-12 appropriations saw no nominal growth, declining slightly in real terms. The 2013 school year[4], in particular, saw a 9% nominal decline.

In August 2010, Congress passed a short extension of the Recovery Act, called the Education Jobs Fund. This was a $10 billion program allocated along the same overall methodology

[4] Throughout this paper, I follow the fiscal year standard of referring to fiscal years by the end year. So, school year 2013 refers to the period July 1, 2012 to June 30, 2013. My analysis will follow this fiscal period, as it is the operating time frame for the local school districts which form the basic unit of analysis of the data used in this analysis.

as the State Fiscal Stabilization Fund. The primary difference between these programs was that the Education Jobs Fund had a shorter time horizon and required that funds be spent on primary and secondary education.

The sharp decline in Federal spending in 2013 was driven by the Budget Control Act of 2011, the agreement which resolved the debt limit crisis of 2011. The declines in nominal Federal education grants is unique to 2013, as future years saw partial offsets to many of the budget cuts passed in this legislation. The combined effect of the Education Jobs Fund, the Budget Control Act of 2011, and the Recovery Act means that I expect to see increased Federal K-12 spending through the 2011 fiscal year, followed by a sharp decline in 2012 and particularly 2013.

2.3 Data

In this section, I describe the main data sources used in the analysis. I begin by describing the original data used and the construction of the analysis database. I then discuss sample selection. Finally, I summarize descriptive statistics used in the analysis.

2.3.1 Data sources and sample selection

The main data sources used in this analysis are the Common Core of Data district level finance and population files. The Common Core is a database maintained by the National Center of Education Statistics at the Department of Education. Is a comprehensive survey taken of all school districts in the United States, including charter and other special purpose school districts. These data sets are originally populated by the Census of Governments and ultimately updated by each state individually before being made available online. Data include both school district finances (such as total spending, or Federally sourced revenue) and school district population characteristics (such as enrollment, the share of students with special needs, or the share of student on poverty).

In addition to the Common Core of Data, I use two additional sources of data to measure additional district characteristics. First is the Census's Small Area Income Population Estimates. This is an annual summary of poverty and population by school district in the United States, and is the main source of data the Department of Education uses in allocating Title I grants to each school district. I also use the Stanford Education Data Archive ("SEDA"), which maps school district boundaries to Census regions for years 2009 and later.[5]

The last piece of data I add is budgetary state-level Recovery Act funding published by the Department of Education. Including this data allows me to construct district-level stimulus fund measures as described in section 2.2. This includes state-level summaries on grants related to Title I, Special Ed, and the State Fiscal Stabilization Fund, attributable to the Recovery Act. For all data, I normalize them to 2009 dollars, using the Bureau of Economic Analysis's Employer Cost Index.

The full analysis database covers the years 2005-2014 at the school district unit of analysis, although I focus primarily on the period 2008-2012, as this is the window in which I feel most confident interpreting my results to be a direct result of the Recovery Act. I include all types of school districts in the database, subject to three main screens. First, I drop top and bottom percentile of school districts, in terms of per pupil spending, to limit sensitivity to outliers. Next, I focus on school districts with more than 100 students. There are several advantages to this approach. Larger school districts are the most policy relevant, serving 99.8% of students and 99.6% of ARRA grants. Because I normalize all data by district enrollment, larger districts are less sensitive to swings in school size. Also important is that this is more comparable to the population served by previous efforts, such as Gordon (2004), so limiting my analysis to this sample allows for clearer comparisons with the existing literature.

[5] As the SEDA data does not cover 2008, I extend the 2009 values back one year to include this year in my analysis window. The main results are not sensitive to either including these measures or dropping 2008.

2.3.2 Descriptive Statistics

I include descriptive data in Table 2.2. Table 2 groups school districts by the quintile of the size of the formula-driven grants each district received in the Recovery Act. Those districts in the first quintile received $532 in Recovery grants, per pupil, and those in the top quintile received $2,214 per pupil. I find that higher-grant districts tend to have lower enrollment, and tend to receive more in the way of Federal and state revenues. As a result, high-Recovery fund districts are typically less reliant on local revenues. In other ways, these districts are quite similar across quintiles. I find that there is no pattern in total per pupil spending across the quintiles. I also find that districts lost a similar amount of local revenue between 2007 and 2009, suggesting there is not a significant relationship between local shocks and ARRA funds.

Figure 2.1 describes the general pattern in funding by revenue source over the analysis period. Average Federal per pupil spending increased by just over $400 from 2008 to 2010, before returning to baseline levels. Local revenue was more steady than state revenue. This is consistent with the results of Seegert (2012). Local governments typically rely on property taxes, which Seegert showed to be less volatile during a recession. State government spending increased steadily from 2004 to 2008, and then sharply dropped following the recession. It then remained low and did not show signs of real growth until 2014. This pattern has several likely explanations – changing control of state governments, recession, and, possibly, a fiscal response to the Recovery Act.

2.4 Methods

In this section I discuss the main identification strategy in my analysis, and underlying assumptions. I start by defining the main parameter of interest, and the challenges in identifying it causally in this setting. I then describe the identifying variation I use in this analysis and the benefits of this approach. I next describe preferred causal model. The main approach I use in this analysis is a two-stage least squares approach which exploits non-linearities in the funding formulas for Title I, Special Education, and the State Fiscal Stabilization Fund.

2.4.1 The key relationship of interest

I seek to identify the causal effect that additional Federal education grants have on school district outcomes such as state revenues, local revenues, and school district employment. The key identifying equation is:

$$\Delta_{2007} Y_{dt} = \gamma \Delta_{2007} FedGrant_{dt} + \beta X_{dt} + \delta_{st} + \phi_d + \varepsilon_{dt} \tag{2.1}$$

In this equation, the outcome variable, Y_{dt}, tracks district revenues, from state or local sources, or district-level employment. These outcome variables, along with the treatment variable, $FedGrant_{dt}$, are measured on a per-student basis, and in terms of their difference from 2007 levels.[6] Measuring these variables as differences from 2007 allows me to isolate the policy-induced differences, while also avoiding grant timing autocorreleation which would be inducing by measuring pure year-over-year differences. This is a combination of the approaches used by Gordon (2004) and Cascio, Gordon and Reber (2013).

This specification includes district fixed effects, ϕ_t, which in the case of a differenced model, allows each district to have a distinct linear trend. I include state-by-year level fixed effects in δ_{st}. At the district level, I include in X_{dt} the share of the students living in poverty, the share of the students with special needs, district enrollment, the district share of black and Hispanic students, and employment within the district. Finally, ε_{dt} is a district specific error term.

[6] While 2008 is the last year prior to the creation of the Recovery Act, I prefer differencing based on 2007 because it is the latest "normal" year in the analysis. 2008 was a full year of recession, combined with a somewhat higher Federal K-12 budget, which makes it an imperfect baseline. Differencing based on 2007 also lets the analysis period include an additional "baseline" year with no ARRA grants. This helps to identify the difference between the effect of, for example, an additional special needs student in any given year, versus the surge in Special Education grants coming from ARRA.

The key parameter of interest is γ, which estimates the causal effect of additional Federal grants in a given district-year observation on our outcomes of interest. In order to interpret γ as causal, however, it must be the case that changes in Federal grants from 2007 are uncorrelated with other unobserved factors which may affect changes in the outcome over this same time horizon (conditional on the other covariates included in equation 2.1. Given the discussion of the circumstances surrounding the Recovery Act in section 2.2, there are several possible threats which would result in $E[\varepsilon_{dt}|FedGrant_{dt}, X_{dt}] \neq 0$. Many of these threats I attempt to address by using a differenced specification, fixed effects, and district observables. For example, by differencing I implicitly control for district fixed effects, and avoid drawing conclusions off baseline differences in districts. Including the share of students with special needs or in poverty will help control for changes in a district's student composition or financial needs, beyond the scope of the Recovery Act.

There are, unfortunately, some threats to a causal interpretation of γ which I may not be able to solve through a selection on observables approach. For example, some funds in the Recovery Act were based on competitive grants, which may go towards districts more organized and capable of receiving grants of all types. This would attenuate any flypaper effect. In addition, state discretion in the use of Recovery Act funds may result in those districts struggling in the recession receiving more in the way of state and Federal funds. I attempt to address these types of threats using a instrumental variables strategy, described below.

2.4.2 Variation in Federal Formulas

As Federal K-12 grants may be endogenous (to district competitiveness, to district economic circumstances, or with other non-educational Recovery Act program funds), I employ an instrumental variables strategy to estimate the effects of Federal grants on district spending. I construct three instruments, one for each of the major K-12 programs in the Recovery Act.

The goal of these instruments is to rely on formula-induced policy variation, which is unlikely to be a tool manipulated by policymakers at at any level of government. Each of the major K-12 Recovery programs relied on a formula which has no obvious relationship the major threats I identify above. In particular, Title I and Special Education grant formulas were developed decades ahead of the recession, and rely on nonlinear components which limit the ability of lawmakers to employ these programs strategically at the district level.

Title I is a Federal program which targets K-12 funds towards economically distressed students. The formula for Title I is calculated at the school district level. Funds are allocated as a function of the average state spending of the district's state, the count and share of economically distressed students within a district, and that district's historical share of Title I funds. In addition, the formula is broken down into 4 distinct programs, each with various kinks for different poverty levels. This means two districts may be extremely similar in terms of the number and share of students in poverty served, but receive different Title I grants. As part of the Recovery Act, the Federal government roughly doubled the size of the Title I program, based on existing district level allocations. I calculate a simulated district Title I allocation as

$$Z_{TitleI} = \frac{TitleI_{d,2008}}{TitleI_{2008}} * \$13,000,000,000 \tag{2.2}$$

That is, each individual district's share of national Title I funding in 2008, times the total additional allocation from the Recovery Act. A summary of the variation in Title I grants can be seen in Figure 2.2.

Special Education funds are Federal grants designed to assist schools in complying with the Individuals with Disabilities in Education Act. They are calculated by the Federal government at the state level. States then have a responsibility to distribute these funds to school districts in a systematic basis. In 2009, similar to Title I, the Federal government roughly doubled its annual appropriation for Special Education. Because state governments have discretion in how these funds are distributed, I use the 2008 district share of state

Special Education funds to proxy for the share going to a given district.[7] The total Special Education funding received by a state is a function of the initial count of special needs students in the state in 1994, along with current state measures of population and poverty. This quirk, combined with differences in how states allocate Special Education grants, creates the variation seen in Figure 2.3.

$$Z_{SpecialEd} = \frac{SpecialEd_{s,d,2008}}{SpecialEd_{s,2008}} * SpecialEd_{s,ARRA} \tag{2.3}$$

Equation 2.3 estimates each district's Special Education funds by using the 2008 share of the enrollment with special needs, multiplied by that state's budgetary allocation for Recovery Act Special Education grants.

With the State Fiscal Stabilization Fund, each state received a grant as a linear function of its population distribution - that is, the share of the school-aged (5-24) population to total population. In this program, the Federal government offered education grants to states, which the state could use for either K-12 or higher education. In order to receive these funds, districts must commit to 1) keeping K-12 and higher education spending at 2006 levels per student or higher, and 2) allocate the K-12 funds in accordance with the state's established aid system. I estimate each district's SFSF treatment by taking the 2008 share of general state aid received by the district, multiplied by the total age-determined SFSF appropriation. The factor .0789 is the share of the SFSF which was ultimately spent on K-12 education, rather than higher education.[8]

$$Z_{SFSF} = \frac{StateGeneralRev_{s,d,2008}}{StateGeneralRev_{s,2008}} * SFSF_{s,ARRA} * 0.789 \tag{2.4}$$

[7] While each state must devise a formula in terms of how these grants reach districts, there is a wide array in formula detail across the US, in terms of the types of disabilities measured and how students with disabilities factor in to the allocation formula. Because the annual flow of changes in students with disability should be small relative to changes in the stock, for any given year, I argue this state share measure to be a close substitute for each state-specific formula.

[8] See Cohen (2011).

Refer to Figure 2.4 for a summary of the variation in the sum of instruments across districts, relative to the size of the fiscal shock, the share of a district's reliance on Federal resources, and total enrollment. Taken together, these policy-induced measures of variation are conceptually and empirically a valuable source of variation for this analysis. By using the formula-induced policy variation from these three measures, I argue that Federal grants received via the Recovery Act are plausibly exogenous to other district level threats which to interpreting the effects of new Federal funding as causal.

2.4.3 Two-stage least squares approach

Below I use the instrumental variables defined in section 2.4.2, along with the main equation in section 2.4.1 to define the 2SLS strategy as follows:

$$\Delta_{2007} FedGrant_{dt} = \phi_{t,d} + \alpha Z_{dt} + \beta_z X_{dt} + \delta_{z,st} + \varepsilon_{z,dt} \tag{2.5}$$

$$\Delta_{2007} Y_{dt} = \phi_t + \gamma \Delta_{2007} \widehat{FedG} rant_{dt} + \beta X_{dt} + \delta_{st} + \varepsilon_{dt} \tag{2.6}$$

The above equations are similar to equation 2.1, in terms of the fixed effect and covariates included. The key differences is using the formula driven variation of the Recovery Act to guide the analysis.

Z_{dt}, here, is a vector of 12 instruments, one for each major K-12 program in the Recovery Act, interacted over each year the program was active. In Figures 2.6, 2.7, and 2.8, I run the reduced form of this equation separately by year. This allows me to flexibly review the timing differences afforded to states in releasing Recovery funds. It also allows an implicit falsification test for the period preceding the analysis, and to see any persistence of crowd-out.

I run the main analysis for the period 2008-2012. This time frame allows me to include the peak of the stimulus and have a pre-period baseline, while avoiding much of the coinciding policy changed which potentially contaminate later years of the analysis. Additionally, this is the period for which I have demographic variables such as race and district unemployment. All fiscal variables are normalized by student enrollment and standard errors are clustered at the district level.

2.5 Results

2.5.1 Discussion of first stage

Refer to Figure 2.5 for a residualized plot of the ARRA induced spending variation and Federal grants. Because I use multiple instruments, it can be difficult to visualize the direct effect of the instrument on the treatment variable. To simplify visualizing the relationship, I add each program instrument together and show the effect for a single year (2010, the peak stimulus year). There is a clear positive relationship between the Recovery Act funds and the change in Federal spending, as would be predicted by the institutional background.

Refer to Figure 2.6 for a summary of the first stage effects of these instruments by year. In this case we see there is essentially no effect of the instruments until 2010 and 2011, at which point they combined explain nearly the total sum of Recovery Act grants. These are the only two years the effect is statistically significant in isolation.

Table 2.3 summarizes the effect of each instrument on the first stage regression. Each column includes one program instrument at a time, until column four, which includes all three. Both SFSF and Special Education Instruments produce a combined effect of about one dollar of additional Federal Grants per dollar of instrumented spending.[9] The Title I instrument is less powerful. This is primarily because the Title allocation process is direct to districts; it lacks the additional state-induced variation of the SFSF and Special Education instruments.[10] The effects are statistically powerful, with an F stat over 25.

[9] I report the linear combination of the sum of coefficients for each instrument from 2009-2012. The SFSF combined coefficient is slightly more than 1, as a result of the Education Jobs Fund extending this program.

[10] Another way of seeing this is to remove district poverty as a covariate. When I do this, the combined coefficient for Title I increases to 1. This does not change my results in any meaningful way, but I leave poverty as a district observed covariate because it may affect spending patterns. Implicitly, this means that my results rely more on variation driven by SFSF and Special Education than Title I.

2.5.2 Financial Effects of the Stimulus

I describe the financial effects of the stimulus in Table 2.4. This includes the OLS and 2SLS results of the main estimating equations 2.5 and 2.6 over the period 2008-2012. I show three specifications: including district fixed effects, including district covariates (demographics), and including both. All three regressions also include state by year fixed effects. In all cases, standard errors are clustered by district.

My main outcomes are state revenue, local revenue, and total revenue. In general, the third will equal state+local+1 (for the additional dollar of Federal revenue). Each of these variables has two main comparisons of note. In general I expect the coefficients on state and local Revenue to be bounded between -1 and 0. A coefficient of -1 implies "full crowd out" or "no flypaper effect" because Federal revenues are being offset elsewhere. Alternatively, a coefficient on state or local revenue of 0 would mean "no crowd out" or a "full flypaper effect." The coefficient on Total Revenues is likewise bounded between 0 and 1, where 1 represents the no crowd out / full flypaper scenario.

While I expect the coefficient to be so bounded, it is not strictly necessary. Other works such as Delaney (2011) have found crowd-in effects with a coefficient greater than 1. While there are certain models which might lead to this result, I general I do not interpret any of my results in this way. Some OLS results are larger than 1, which I interpret as bias relating to possible concurrent state policy shifts or local grant-seeking. Some 2SLS point estimates are larger than 1, but none are bigger in a statistically significant sense.

The first three columns show the OLS effects of and additional dollar of Federal revenue on other revenue sources. I focus my discussion primarily on column 3, which is my preferred specification. I find no evidence of crowd out, and perhaps evidence of crowd-in, when relying on the OLS results. This suggests that an additional dollar of Federal grants might induce more state and local spending. I interpret this as a specification error and focus on the 2SLS results.

Columns 4-6 highlight the 2SLS results. In general, these results are more credible, in the sense of being more consistent with priors about the likely range. I note that columns 5 and 6 are roughly identical, suggesting that district fixed effects add little once I have already differenced the outcome and treatment variables, and included demographic covariates. My point estimates for both state and local revenue responses are positive, but low and not significantly different from zero, suggesting little in the way of crowd out. When I look at the effects on total revenues, I can reject any point estimate lower than .4, or 40 cents on the dollar. Thus, this shows a significant flypaper effect, but casts a wide range in possible crowd out. Below, I will decompose this into the effect driven by the State Fiscal Stabilization Fund, and that driven by other existing programs such as Title I and Special Education.

2.5.3 Discussion of Employment Effects

Table 2.5 shows the employment effects of my main specification. I note that the identifying equation picks up the number of teachers hired per every \$100,000 dollars of Federal grants. In this discussion I will also refer to job-year costs, which I create by converting this using the factor $\frac{1}{\gamma}$.

I show three main categories of response – teacher cost, "other" employees cost, and total employment cost. "Other" employees include all nonteacher types of employment, for example, teaching aides, guidance counselors, or librarians. I consolidate these into a single category for two main reasons. First, most sub-groups are too small to be analytically interesting on their own. Second, data reliability for sub-groups is limited over the full analysis period, so I instead rely on the more stable aggregate. The cost of any-source employment is typically the lowest of the three, by construction. Including either just teachers or just other employees only gives a limited picture of the total employment response. On average, school districts have about .071 teachers per student and about .062 other employees per student. At baseline, the two groups are similar in size.

The point estimates for the OLS results are included in columns 1-3. I again note that results are insensitive to adding district fixed effects once district demographics are included.

I find that teachers are more expensive than other employees, by nearly a factor of three. We see that $100,000 in Federal grants creates .311 new teachers, .835 new other employees, and 1.146 employees in total Taken together, this means that every additional $87,260 in stimulus grants leads to an additional job-year in direct hiring.

Columns 4-6 show the 2SLS effects of additional Recovery Act grants on district employment. All point estimates are broadly lower than in job-year cost that the OLS estimates. This is consistent with my flypaper analysis, which found OLS spending effects to be larger (which would subsequently amplify the perceived hiring). Again as in OLS, I find that teacher hiring is much less responsive than other employee hiring. In my preferred specification, I find that school districts are able to provide an additional job-year of employment for every $108,000 in Federal grants (or .925 jobs per $100,000)

This point estimate is noteworthy in at least two ways. First, it is precise enough to reject many of the higher estimates for Recovery Act job costs, such as those found in Feyrer and Sacerdote (2011) (over $500,000 for an education job) or the low end, such as in Chodorow-Reich et al. (2012) ($26,000). Second, my estimate is unique from other efforts to study this problem because I rely on direct employment results only. This means that in some ways, my results could be considered a lower bound, with larger downstream employment effects. Thus, these estimates bring the effects of stimulus education spending in line with other areas such as highway and construction spending[11].

2.5.4 Decomposition of Results - Stabilization Funds

The Stabilization Fund is distinct from the other two program instruments (Title I and Special Education) in several key ways. It is more flexible, as states have the choice to spend it on higher education. Even if it is spent entirely at the K-12 level, additional flexibility remains, as districts can then chose to spend either directly on employment or on fixed costs. In addition, the SFSF will go away after a few years, whereas states will continue to use Title

[11] For example, both Feyrer and Sacerdote (2011) and Wilson (2012) found these job costs to be in the range of $100,000 to $125,000

I and Special Education funds for years after the stimulus ends. It stands to reason that these two programs would not produce the same effects on employment and total spending.

Tables 2.6 and 2.7 summarize the effects of the SFSF on district fiscal responses and employment. Looking at the preferred specification in column 6 of Table 2.6, we see a point estimate of -.181 for local revenues, which is precise enough to reject the possibility of no crowd-out. The crowd-out effect on state revenues is even higher. A dollar of SFSF grants causes state governments to send 72 cents less to a given district. This suggests state governments saw the SFSF as a near perfect substitute for their own general revenues. The total effect shows nearly complete crowd-out. I can reject any crowd-out less than 70 cents on the dollar, with a point estimate of a flypaper effect at roughly 10 cents per dollar.

As a result of the fiscal effects described above, the employment effects of the SFSF in Table 2.7 show that the SFSF did not effectively lead to additional employment. For both teachers and nonteachers, the point estimates suggest a job year cost in excess of $1 million. If SFSF did not in fact induce additional spending at the district level - which it appears it did not - this conclusion is natural.

2.5.5 Decomposition of Results - Program Analysis

Finally, I look at the effects of additional grants, specifically for Title I and Special Education, without including the Stabilization Fund. The results are included in Table 2.8 and are roughly similar to those in Table 2.4, only with less power. This is intuitive; if the SFSF did not have a significant flypaper effect, then the rest of the Recovery Act grants should have a similar effect to the total. I can reject a crowd-out effect any greater than 50 cents on the dollar in this specification.

Finally, in Table 2.9, I look at the employment effects of the Title I and Special Education grants. The results here follow a familiar pattern; teachers are the least responsive to additional grants, with only .281 new jobs per $100,000 in additional grants. For both teachers and nonteachers, the average job year cost is $98,000. The point estimate here is about as precise as that on the total employment effect in Table 2.5. The 2-sigma range is roughly \$35,000 to \$160,000.

2.6 Conclusion

In this paper, I assess the effects of a significant, but temporary expansion of the Federal role in funding K-12 education. This situation provides for a useful natural experiment which can help shed insight in to long-standing empirical phenomena such as the flypaper effect. It can also address important policy questions debated during the crafting of the Recovery Act, such as whether fiscal stimulus for the Federal government is an efficient way to meet the goal of increasing employment.

This approach has two major contributions. First, I assess a temporary grant expansion, which is distinct from previous efforts such as Cascio, Gordon and Reber (2013) and Gordon (2004), which dealt with permanent changes in Federal grant-making. Second, I am able to exploit district level, formula-driven variation which is plausibly exogenous and powerful enough to deliver precise estimates of the employment effects of the stimulus.

With respect to the effect of temporary grants on state and local finances, I find that the Recovery Act had a significant flypaper effect, ranging between 40 cents on the dollar to a full flypaper effect. The Federal government was successful in increasing its spending at the state and local level. With respect to employment costs, I found that school districts employed an additional person for one year for every \$108,000 spent in the Recovery Act. Running back of the envelope calculations from this estimate, the K-12 education section of the Recovery Act alone lead to 563,000 additional jobs (spread over four years). My analysis is fairly precise within the realm of the Recovery Act literature, allowing me to reject some of the more extreme estimates in the literature, on both the high and low end.

Comparing the type of grant suggests that the there was a much larger flypaper effect associated with Title I and Special Education funds than the State Fiscal Stabilization Fund. There are multiple possible interpretations of this pattern. It may simply be that Title I and

Special Education had a more strict set of requirements than the SFSF, which made them more difficult for state and local governments to offset. Similarly, it may be the case that as a new program, the SFSF required new compliance rules and was not as well understood as the existing programs, which ultimately made it more difficult for districts to implement into their overall budget.

The final key finding of this paper is that I find school district hiring of non-teachers is more responsive to additional grants than the hiring of teachers. I interpret this as being related to labor protections. If teachers are difficult to fire, and fiscal pressure is forcing a district to make layoffs, they will start with non-teachers first. If those districts receive additional Federal grants, they will have less need to make these kinds of layoffs. Further work could be done to look at the differences in employment effects in areas with strong vs. weak labor protections.

These results may inform policymakers looking to respond to recessions or other fiscal crises in the future. It seems likely that the lack of flexibility in times of distress is an important component of evaluating labor regulations. Also shown in this paper, the flypaper effect is sensitive to the structure of the program. Piggy-backing off of preexisting frameworks may lessen the state and local capacity to respond.

2.7 Figures

Figure 2.1: Time series share of K-12 funds sourced from the Federal government

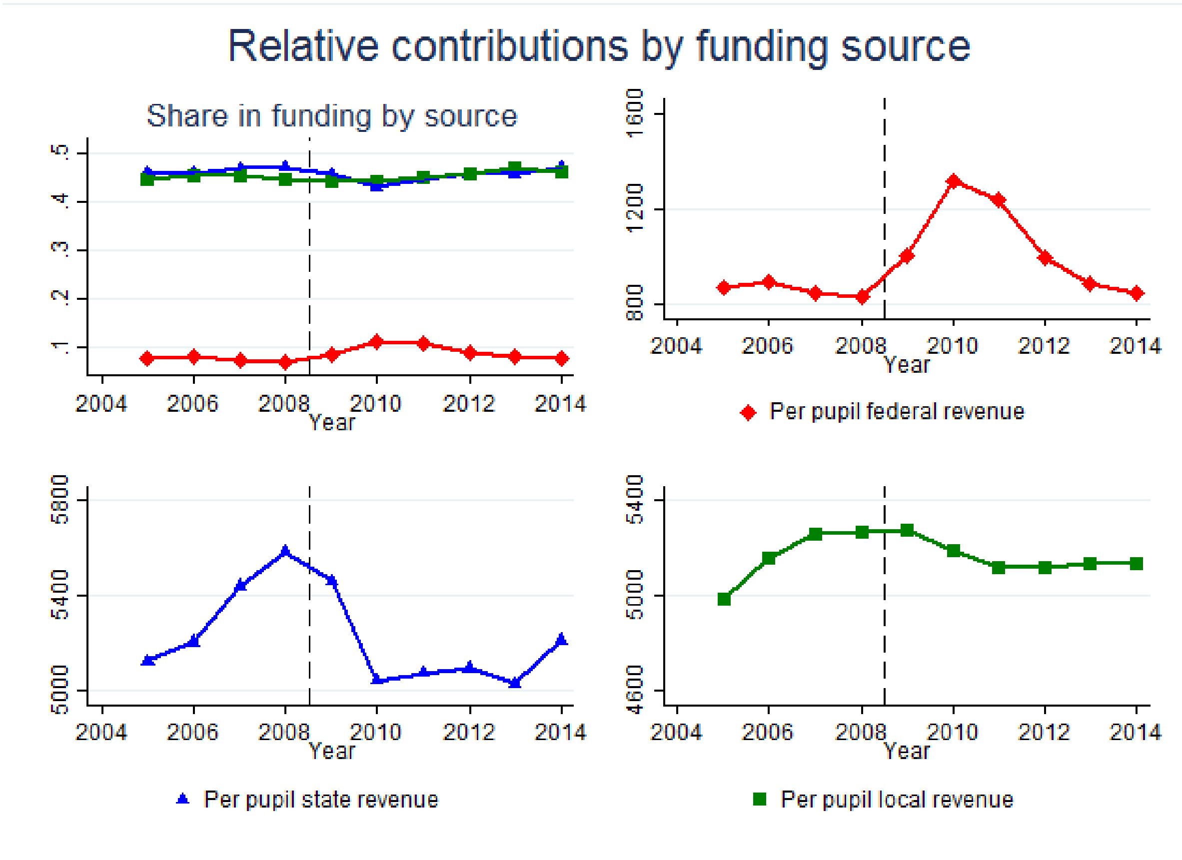

Figure shows the trend in the share of total school district resources in the US coming from Federal funds (calculated by the aggregate Federal K-12 spending in the US over aggregate total revenues across all districts in the US, from the common core of data). The dashed lines emphasize the beginning of the Recovery Act. In this period the total Federal share jumps dramatically, before returning to a more typical level. State revenues decline following the recession in a pattern which could plausibly be driven by the Recovery Act. Local spending remains relatively constant.

Figure 2.2: Variation in Title I Instrument by District Poverty

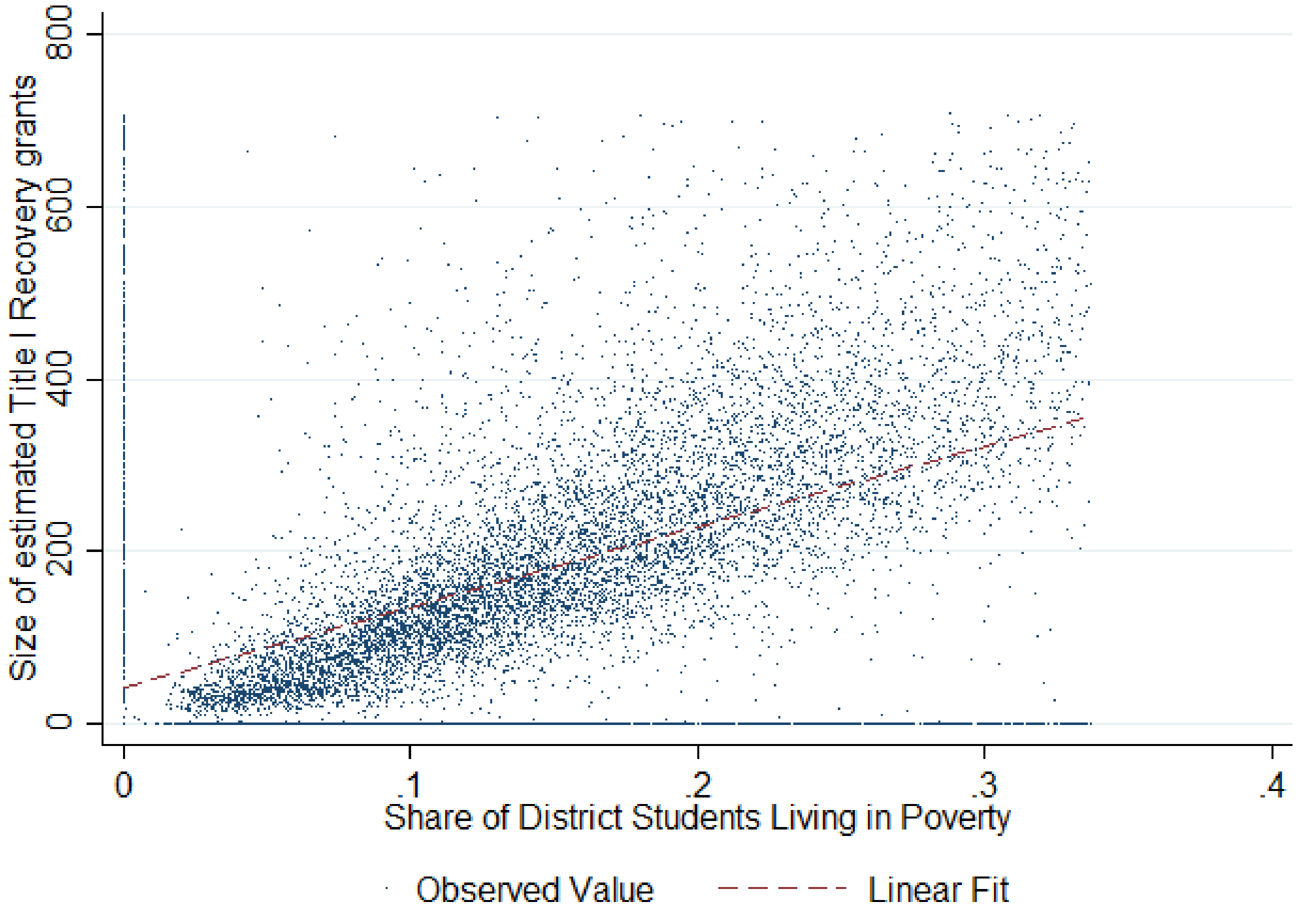

Figure shows the per-student instrument in Title I funds, relative to district poverty as of 2009. The Title I revenue instrument is the district level statutory amount computed per district as described in section 2.4. While the instrument is increasing in poverty, the off-fit variation is what identifies my analysis. For any level of poverty, there is significant formula-driven variation between districts in Title I funding.

Figure 2.3: Variation in scaled Special education instrument by special education enrollment

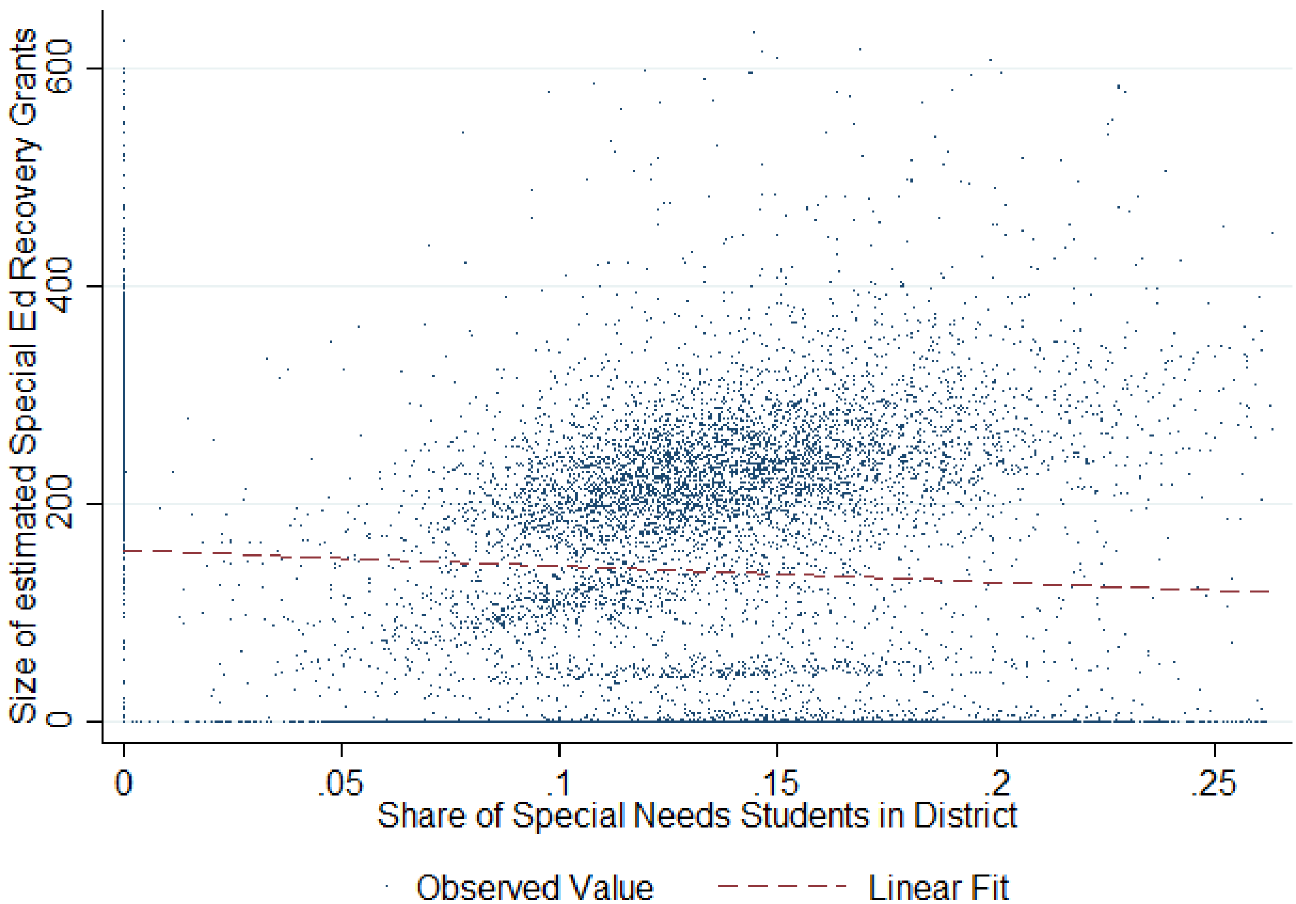

Figure shows the per-student instrument in Special Education funds, relative to district special needs enrollment as of 2009. The Special Education revenue instrument is the district's share of total Special Education funding with the state in 2008, multiplied by the Federal-to-state special education grant boost in ARRA. While the instrument is slightly decreasing in special needs enrollment, the off-fit variation is what identifies my analysis. The negative slope in this case signifies that a large share of the variation is driven by interactions of state and Federal policy that operate beyond a simple measure of special needs enrollment.

Figure 2.4: Variation Total ARRA Funding

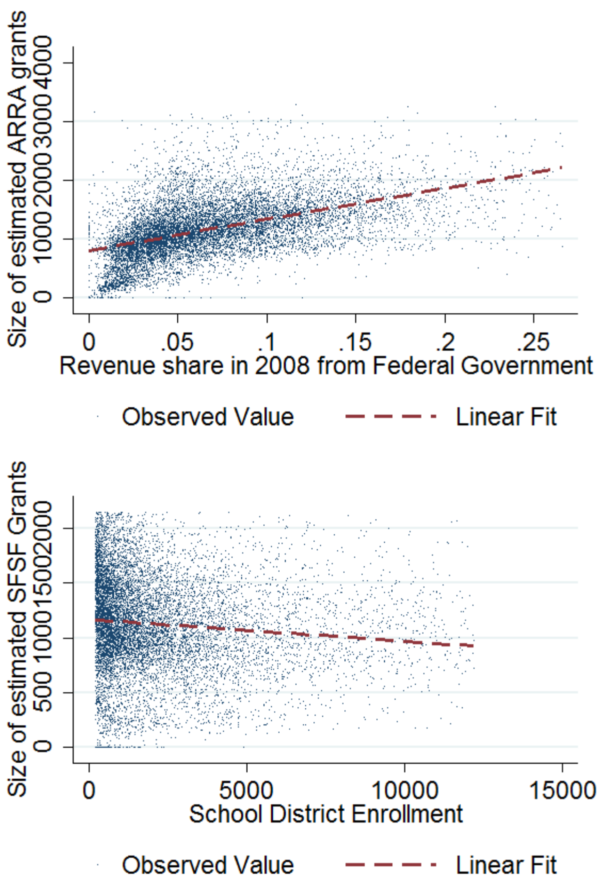

Figure 2.4 shows the off-fit variation for the Stabilization Fund as well as for the total grants. I find in the first panel that a district's share of Federal revenue in 2008 strongly predicts its total Recovery Act funding. In the upper right panel, I show that a district's fiscal shock (change in 2007 to 2009 local revenue) is unrelated to the total grants it receives. Finally, as Stabilization Funds are driven by state population, I show that total district enrollment only weakly predicts total district enrollment.

Figure 2.5: First Stage Residualized Relationship

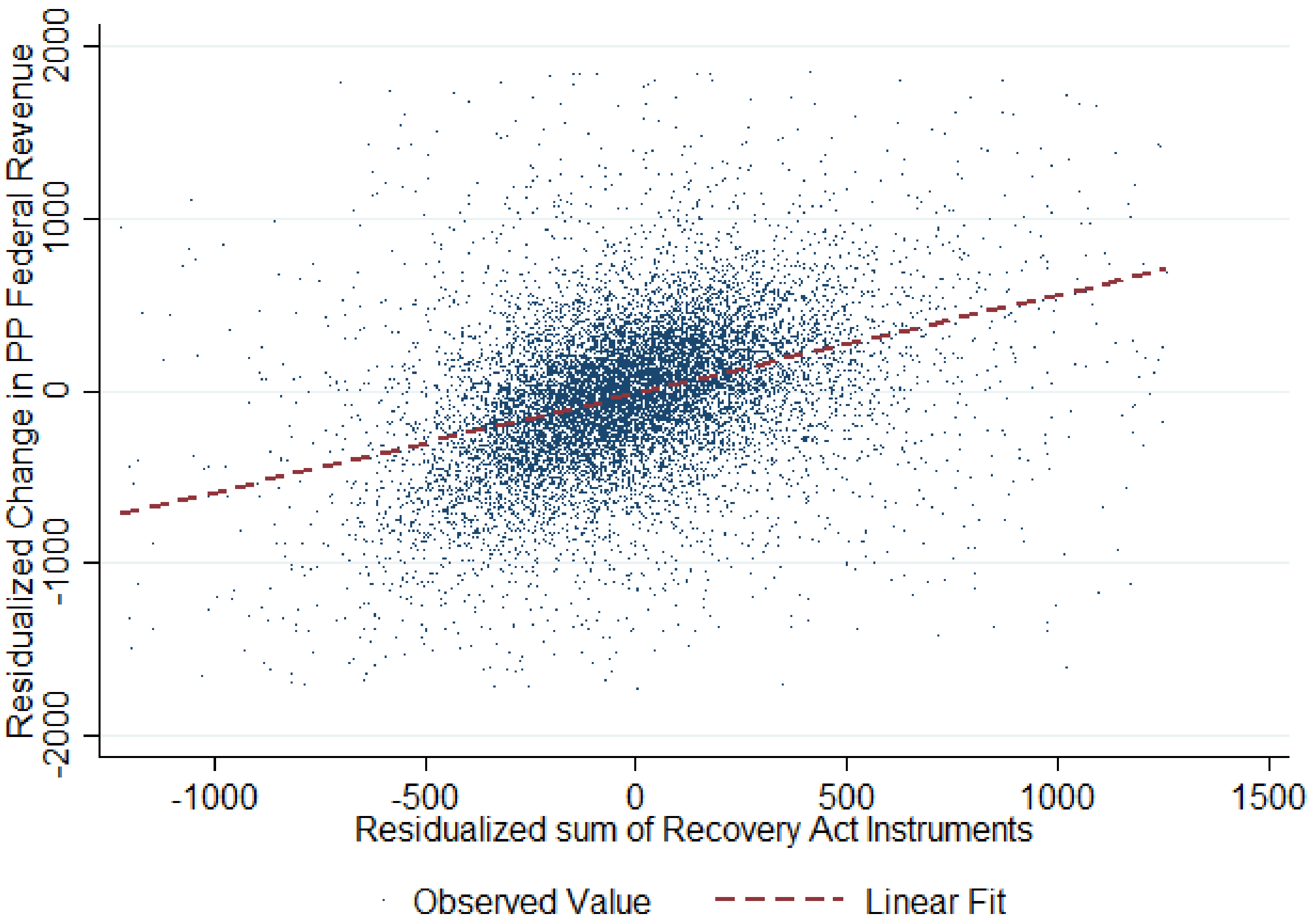

Here I show the first stage relationship between the growth in district-level per pupil Federal revenue, from 2007 to 2010, and total Recovery Act grants, as defined by the sum of my instruments. I residualize my instrument and per pupil Federal revenue against state fixed effects and district demographics. Year and local fixed effects are implicit in differencing and running this within a single year. 2011 shows a similar, albeit weaker pattern.

Figure 2.6: Reduced Form Effect of Recovery Grants on District-level Federal Spending by Year

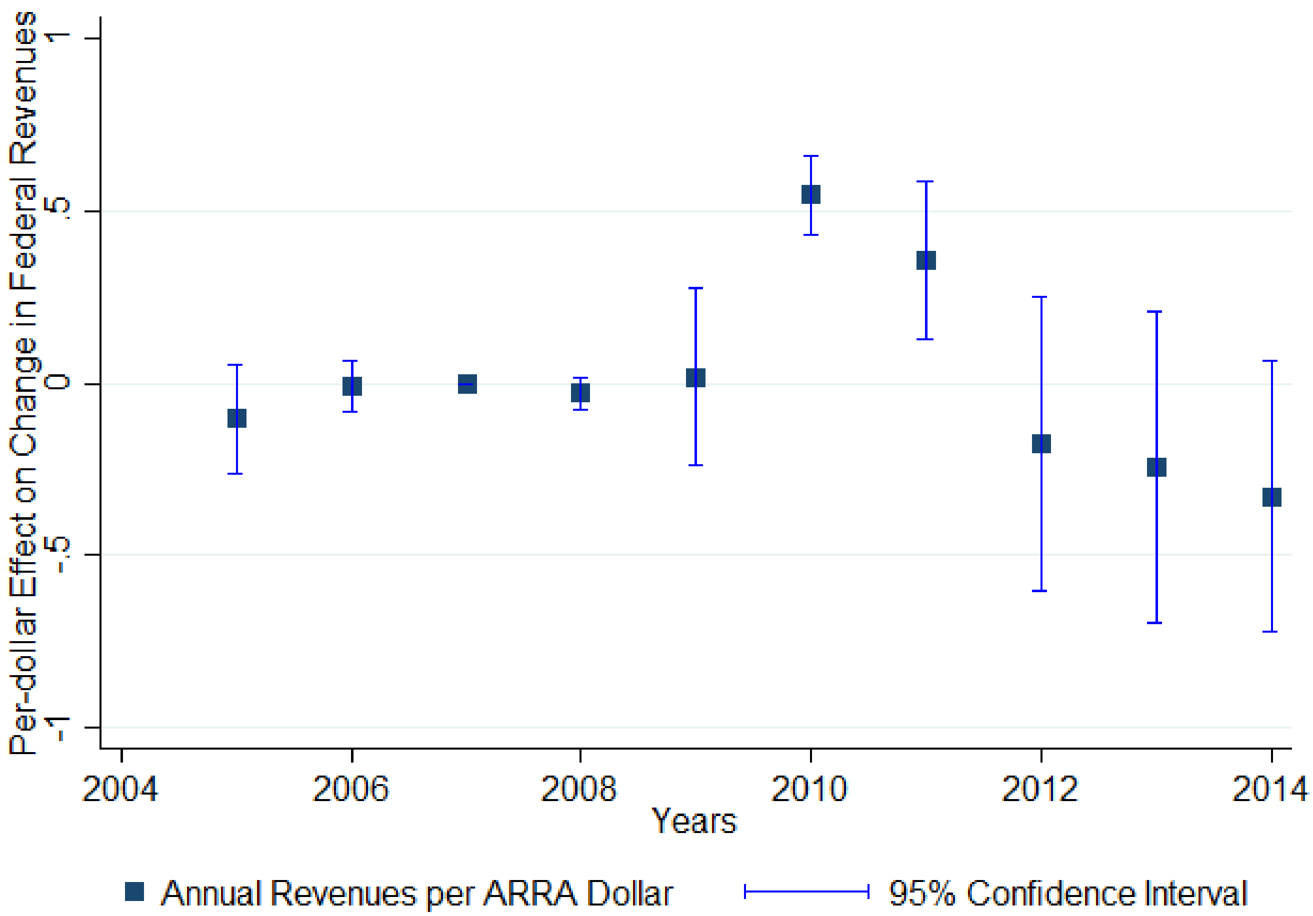

Figure shows the annual effect of my total Recovery Act grant instrument on the change in Federal per pupil revenues, by year. 2007 is zero by construction. Prior to the Recovery Act, the instrument has no explanatory power on district Federal revenues, and is close to zero. In 2010 and 2011, the size jumps to about 50 cents per dollar granted. The regressions shown here are done on ARRA dollar grants, by each year, with district demographics (race, enrollment size, free and reduced lunch eligibility, and employment) and and state fixed effects included as covariates. Including other district covariates which are likely to be correlated with Federal spending, such as poverty or special needs students, will show a similar pattern but lower the explanatory power of the ARRA grants due to multicolinearity.

Figure 2.7: Reduced Form effect of State Revenues

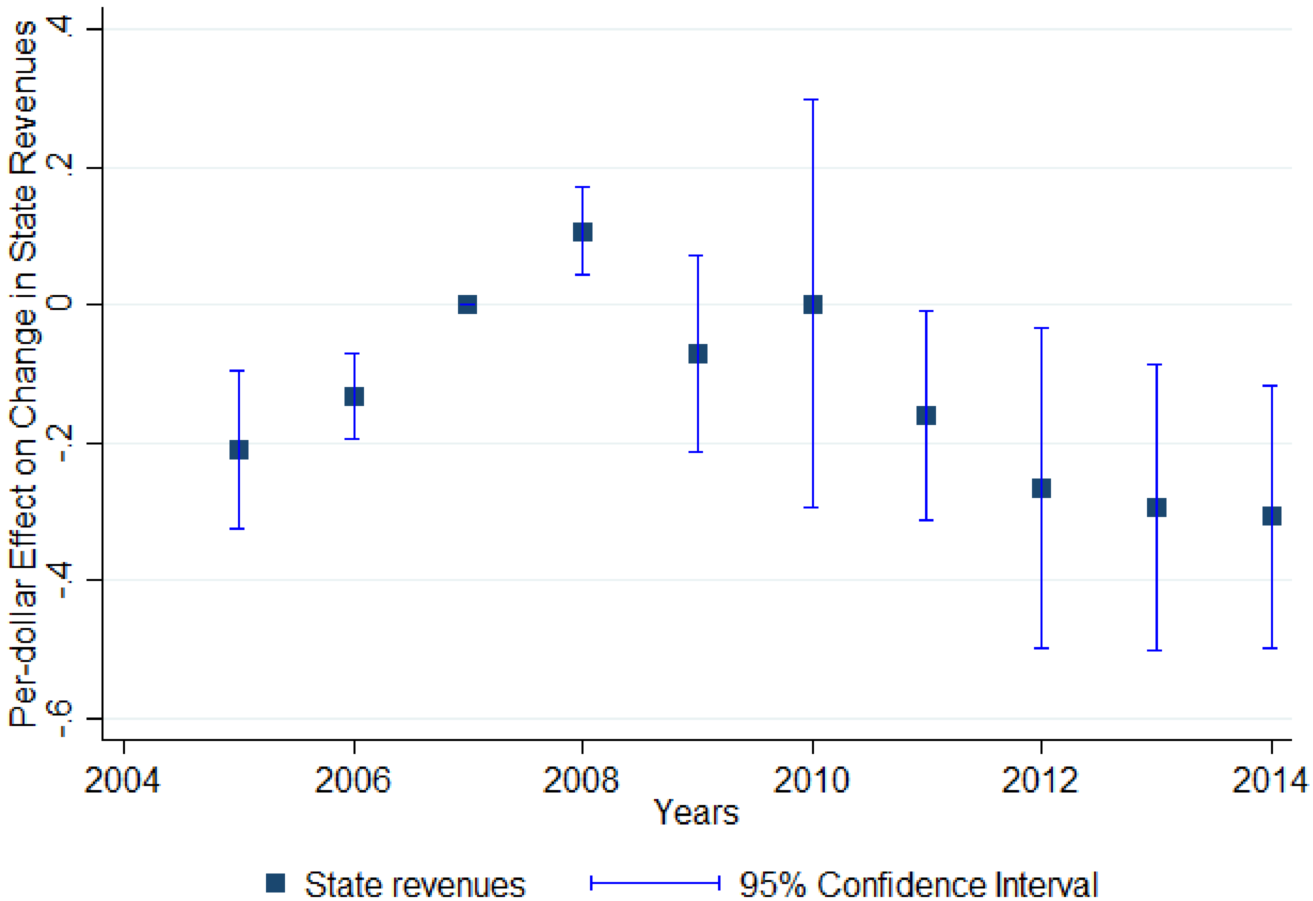

Figure shows the effect of ARRA grants on the change in per pupil state revenue at the district level (again, differences relative to 2007, making 2007 zero by construction). Unlike the Figure relating to Federal revenues, the state revenues have a secular trend prior to my analysis period, complicating the analysis. The bias prior to the implementation of the Recovery Act is bounded between -.2 (2005) and .2 (2009). The post period is consistent with a bias of -.2. This time series approach suggests a fairly limited crowd-out effect, consistent with my main results in Table 2.4.

Figure 2.8: Reduced Form effect on Local Revenues

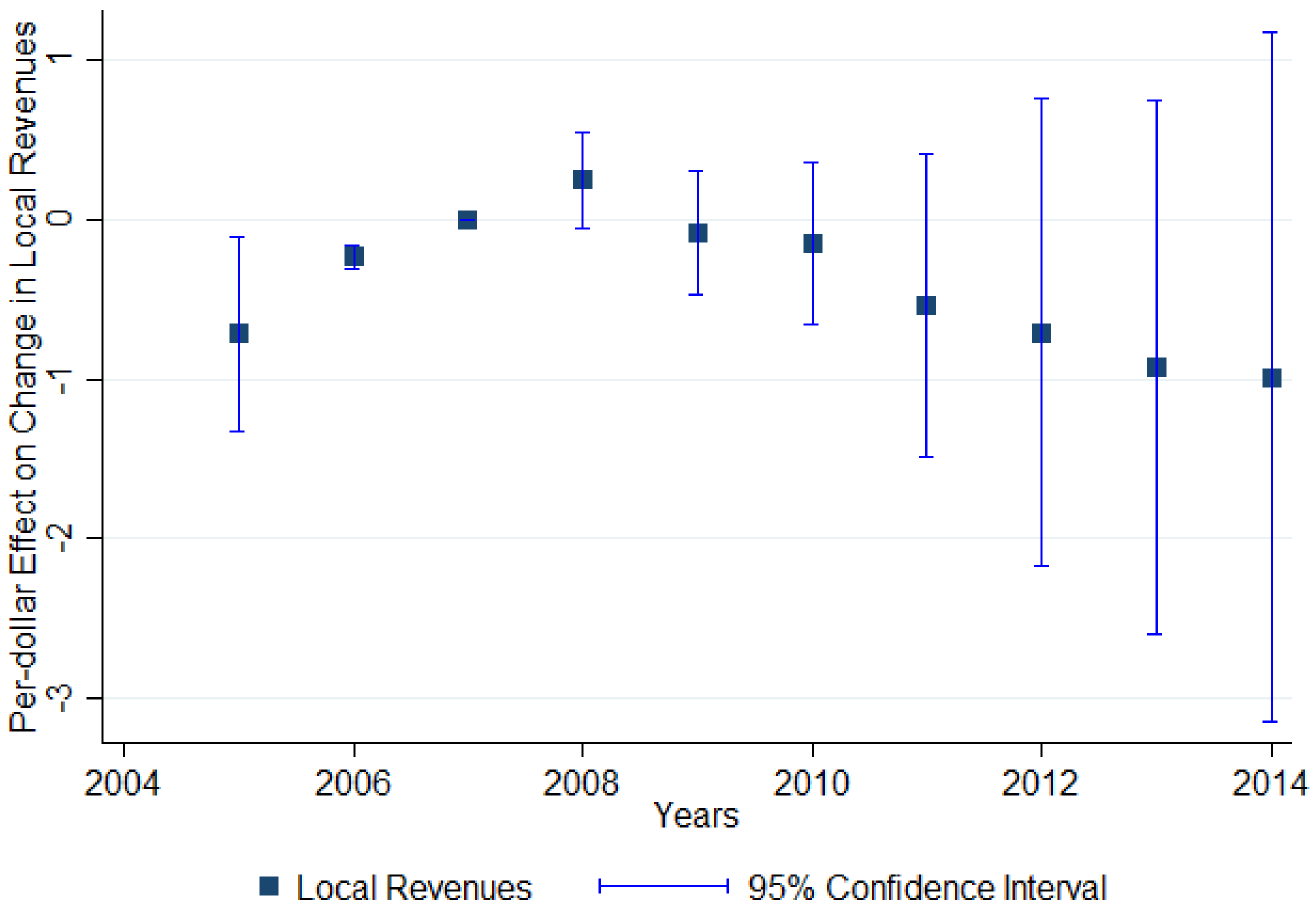

Figure shows the effect of ARRA grants on the change in per pupil local revenue at the district level (again, differences relative to 2007, making 2007 zero by construction). As in the previous Figure relating to state revenues, here I see that there is some bias in the prior periods. Due to the size of the confidence intervals I cannot draw strong conclusions about the possibility of local revenue crowd out. It is noteworthy, however, that the pattern of the coefficient increasing over time and relatively low in 2010 is inconsistent a scenario that involved a large amount of local crowd-out. As above, this Figure is broadly consistent with Table 2.4.

2.8 Tables

Table 2.1: Summary of ARRA spending by program

Composition of the Stimulus (in billions)

Total Stimulus Spending	
Contracts, Grants and Loans	261.2
Tax Benefits	290.7
Entitlements	264.4
Total	816.3

Total Grants, Contracts, and Loans	
Education	94
Transportation	39.2
Infastructure	33.8
Energy/Environment	30.2
Research	15.9
Housing	14
Health	12.2
Misc Other	21.9
Total	261.2

Total Education	
SFSF	51.4
Pell	16.5
Title I	12.4
Special Ed	12.2
All other	1.5
Total ARRA	94
Education Jobs Act (2011)	10
Education Stimulus	104

Table 2.1 summarizes the total funding in Recovery Act across programs. Of the $816 million total, it was split in roughly three ways between tax cuts, entitlements, and operating programs ("Contracts, Grants, and Loans"). Of the total that went into operating programs, Education was by far the largest share, making up over 1/3 of the total. Studying education is then, in a very practical way, central to understanding the consequences of the Recovery Act. The final column shows that Education funds in the Recovery Act were predominantly K-12. The instruments I use in this analysis capture over 80% of the total education funds.

Table 2.2: Summary Statistics of School Districts by Recovery Funds Received

	Quintile				
	(1)	(2)	(3)	(4)	(5)
ARRA Instrument	532	929	1,145	1,397	2,214
	(224)	(64)	(64)	(86)	(1,969)
Enrollment	3,560	3,395	3,315	2,977	2,504
	(4,990)	(5,216)	(5,357)	(4,990)	(4,681)
Change in local revenue (2007-2009)	-344	-144	-90	-110	-296
	(1,743)	(973)	(1,726)	(1,279)	(3,974)
Federal Revenue (PP)	470	607	730	954	1,936
	(482)	(496)	(521)	(598)	(3,624)
Title I (PP)	71	93	153	239	475
	(99)	(104)	(120)	(137)	(764)
Special Ed (PP)	117	90	130	171	448
	(113)	(105)	(110)	(127)	(1984)
Local Revenue (PP)	9,931	5,006	4,332	3,987	4,944
	(5,185)	(3,197)	(3,100)	(3,246)	(9,051)
State Revenue (PP)	3,361	5,294	5,724	6,108	8,356
	(1,845)	(1,880)	(1,937)	(2,400)	(6,294)
Total Expenditure (PP)	13,648	10,948	10,712	10,921	14,895
	(5,347)	(3,401)	(3,344)	(3,798)	(15,144)

Here I summarize the finances and size of school districts, by each quintile of total Recovery grants. High and low grant districts are different in some key fundamental ways. All numbers are in real 2009 dollars. High grant districts are smaller (lower enrollment) and rely more on state and Federal revenues. As a result they are likely lower income as well. Low ARRA grant districts are less reliant on state and Federal resources, and as a result likely to be more resilient against a recession (local governments tend to rely on low volatility property taxes, see Seegert (2012). Crucially, there is no clear pattern across quintiles in either total per pupil expenditure, or in the size of the revenue shock as a result of the recession.

Table 2.3: Power of First Stage Instrument

	Model			
Stat	(1)	(2)	(3)	(4)
Stabilization Fund	1.643***			1.350***
	(0.203)			(0.182)
Title I		0.721***		0.069
		(0.328)		(0.274)
Special Ed			1.181***	1.148***
			(0.390)	(0.392)
F	19.088	19.893	28.434	25.975
N	63,884	63,884	63,884	63,884
Years	2008-2012	2008-2012	2008-2012	2008-2012
Demographic Controls	x	x	x	x
District Fixed Effects	x	x	x	x
State by Year FE	x	x	x	x

Table includes first stage coefficient on the sum of four years of instruments, for each program instrument. That is, in column 1, I show the combined effect of SFSF on per pupil Federal grants, for 2009, 2010, 2011, and 2012, when the SFSF is the only instrument used. Each column also includes standard errors in parentheses and the first stage F statistic. I include the programs as instrumental variables one at a time in the first three columns, and then all together in the final column. Instruments appear to be powerful, exceeding the necessary F statistic. While Title I appears to be weakly correlated with Federal per pupil spending, this is primarily a result of including district poverty in the main regression. Removing this poverty measure would change the Title I instrument to have a coefficient of 1 in column four, and does not materially change the estimates of this paper.

Table 2.4: Effect of Federal Grants on District Finance

	Model					
	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Local Revenues	0.176	0.734***	0.655***	-0.108	0.445	0.350
	(0.118)	(0.227)	(0.228)	(0.271)	(0.326)	(0.383)
State Revenues	0.031	0.260***	0.269***	-0.122**	-0.072	0.008
	(0.034)	(0.066)	(0.066)	(0.059)	(0.125)	(0.127)
Total Revenues	0.446***	1.994***	1.924***	-0.107	1.373***	1.359***
	(0.166)	(0.284)	(0.286)	(0.375)	(0.434)	(0.495)
N	63,884	63,884	63,884	63,884	63,884	63,884
Years	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012
State by Year FE	x	x	x	x	x	x
Demographic Controls		x	x		x	x
District FE	x		x	x		x

Table shows the effect of additional Federal grants on total district expenditures. Each cell is a distinct regression. Rows are different outcomes and columns are different modeling assumptions. Regression results from equations 2.5 and 2.6. Preferred specification in column 6. OLS results suffer from endogeneity and show consistent crowd-in. IV results show no identifiable state or local response. In the final row, I can reject a flypaper effect any lower than 40 cents on the dollar, but cannot reject a full flypaper effect with no state or local crowd out. Standard errors clustered by district.
*** p<.01, ** p<.05, * p<.10

Table 2.5: Effect of Federal Grants on District Employment

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
New teachers per $100,000	0.298***	0.345***	0.311***	0.163	0.281***	0.255***
	(0.066)	(0.070)	(0.064)	(0.135)	(0.077)	(0.086)
Other jobs per $100,000	0.774***	0.888***	0.835***	0.616**	0.718***	0.670***
	(0.181)	(0.190)	(0.187)	(0.284)	(0.229)	(0.246)
Total jobs per $100,000	1.071***	1.233***	1.146***	0.778**	0.999***	0.925***
	(0.238)	(0.254)	(0.245)	(0.377)	(0.286)	(0.307)
N	63,884	63,884	63,884	63,884	63,884	63,884
Years	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012
State by Year FE	x	x	x	x	x	x
Demographic Controls		x	x		x	x
District FE	x		x	x		x

Table shows the effect of and additional $100,000 in Federal grants on total district employment. Each cell is a distinct regression. Rows are different outcomes and columns are different modeling assumptions. Regression results from equations 2.5 and 2.6. Preferred specification in column 6. The 2SLS estimate of .925 shows that it costs just over $100,000 (approximately $108,000) for each new job created by ARRA. I find that stimulus funds have a larger effect on non-teacher employment than teacher employment, creating .67 new jobs per $100,000. I can reject job year results larger than $178,000, and in fact argue that this is a upper bound as I cannot account for spillovers outside of direct education hiring.
*** p<.01, ** p<.05, * p<.10

Table 2.6: Effect of Stabilization Fund Grants on District Finances

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Local Revenues	0.176	0.734***	0.655***	-0.419**	-0.071	-0.185***
	(0.118)	(0.227)	(0.228)	(0.201)	(0.059)	(0.071)
State Revenues	0.031	0.260***	0.269***	-0.880***	-0.859***	-0.718***
	(0.034)	(0.066)	(0.066)	(0.157)	(0.050)	(0.056)
Total Revenues	0.446***	1.994***	1.924***	-0.489	0.071	0.097
	(0.166)	(0.284)	(0.286)	(0.325)	(0.093)	(0.107)
N	63,884	63,884	63,884	63,884	63,884	63,884
Years	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012
State by Year FE	x	x	x	x	x	x
Demographic Controls		x	x		x	x
District FE	x		x	x		x

Table shows the effect of additional Federal grants on total district finances, when limiting the variation to that induced from the Stabilization Fund. Each cell is a distinct regression. Rows are different outcomes and columns are different modeling assumptions. Regression results from equations 2.5 and 2.6. Preferred specification in column 6. Here, I find evidence of a large crowd-out and almost no evidence of a flypaper effect. Both local and state revenues are crowded out by SFSF grants, although there is a much larger effect in state revenues. The effect on total revenues is almost zero, suggesting a very low or zero flypaper effect. I can reject a flypaper effect larger than 20 cents on the dollar.
*** p<.01, ** p<.05, * p<.10

Table 2.7: Effect of Stabilization Fund Grants on District Employment

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
New teachers per \$100,000	0.298***	0.345***	0.311***	-0.082	0.059	0.004
	(0.066)	(0.070)	(0.064)	(0.105)	(0.046)	(0.057)
Other jobs per \$100,000	0.774***	0.888***	0.835***	-0.098	0.144**	0.089
	(0.181)	(0.190)	(0.187)	(0.198)	(0.067)	(0.083)
Total jobs per \$100,000	1.071***	1.233***	1.146***	-0.179	0.203**	0.093
	(0.238)	(0.254)	(0.245)	(0.253)	(0.102)	(0.123)
N	63,884	63,884	63,884	63,884	63,884	63,884
Years	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012
State by Year FE	x	x	x	x	x	x
Demographic Controls		x	x		x	x
District FE	x		x	x		x

Table shows the effect of an additional \$100,000 in Federal grants on total district employment. Each cell is a distinct regression. Rows are different outcomes and columns are different modeling assumptions. Regression results from equations 2.5 and 2.6. Preferred specification in column 6. Here I find that as the SFSF induced very little additional education spending. The coefficient of .093 corresponds to a job year cost of roughly \$1 million.
*** p<.01, ** p<.05, * p<.10

Table 2.8: Effect of Program Grants on District Finances

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Local Revenues	0.176	0.734***	0.655***	-0.106	0.525	0.409
	(0.118)	(0.227)	(0.228)	(0.274)	(0.369)	(0.431)
State Revenues	0.031	0.260***	0.269***	-0.116*	0.068	0.128
	(0.034)	(0.066)	(0.066)	(0.060)	(0.126)	(0.131)
Total Revenues	0.446***	1.994***	1.924***	-0.105	1.593***	1.537***
	(0.166)	(0.284)	(0.286)	(0.378)	(0.480)	(0.548)
N	63,884	63,884	63,884	63,884	63,884	63,884
Years	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012
State by Year FE	x	x	x	x	x	x
Demographic Controls		x	x		x	x
District FE	x		x	x		x

Table shows the effect of additional Federal grants on total district finances, when limiting the variation to that induced from the existing programs (Title I and Special Education). Each cell is a distinct regression. Rows are different outcomes and columns are different modeling assumptions. Regression results from equations 2.5 and 2.6. Preferred specification in column 6. Here, I find no evidence of crowd out at the state and local level, and find a large flypaper effect. I can reject a flypaper effect lower than 50 cents on the dollar. Thus, it seems that the grants related to existing programs, rather than the SFSF, drove the increase in district spending.

*** p<.01, ** p<.05, * p<.10

Table 2.9: Effect of Program Grants on District Employment

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
New teachers per $100,000	0.298***	0.345***	0.311***	0.214*	0.296***	0.281***
	(0.066)	(0.070)	(0.064)	(0.129)	(0.084)	(0.090)
Other jobs per $100,000	0.774***	0.888***	0.835***	0.723**	0.813***	0.738***
	(0.181)	(0.190)	(0.187)	(0.295)	(0.249)	(0.267)
Total jobs per $100,000	1.071***	1.233***	1.146***	0.937**	1.109***	1.020***
	(0.238)	(0.254)	(0.245)	(0.384)	(0.310)	(0.331)
N	63,884	63,884	63,884	63,884	63,884	63,884
Years	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012	2008-2012
State by Year FE	x	x	x	x	x	x
Demographic Controls		x	x		x	x
District FE	x		x	x		x

Table shows the effect of an additional $100,000 in Federal grants on total district employment, when limiting the variation to that induced from the existing programs (Title I and Special Education). Each cell is a distinct regression. Rows are different outcomes and columns are different modeling assumptions. Regression results from equations 2.5 and 2.6. Preferred specification in column 6. Here I find that these programs induced additional employment at an even lower costs than estimated in Table 2.5. Thus, existing programs met this particular goal more effectively than the SFSF.
*** p<.01, ** p<.05, * p<.10

CHAPTER III

Substituting Higher Education for Medicaid: A Study on the Growth of Entitlements

with Morris Hamilton

3.1 Introduction

In this paper, we evaluate the causal relationship between state Medicaid obligations and state subsidies towards higher education. Medicaid and state higher education programs both constitute a large share of state spending. Moreover, these shares have been trending in opposite directions for years (see Figure 3.1 and Table 3.1). On average, Medicaid spending has grown to almost 20% of state budgets in 2010, relative to the 7% of state budgets at the program's inception in 1977. Over the same time period, higher education subsidies have fallen to about 10% of state budgets (a decline of 5 percentage points).

There is an intuitive explanation for why this relationship is likely to be causal. Growth in the Medicaid program is driven by technology and Federal policy, much of which is exogenous to state choices. Funding for higher education, on the other hand, is inherently a state decision. As a result, increases in Medicaid obligations, driven by external factors, will require states to find a way to offset the new spending requirements elsewhere. States are likely to respond to this pressure through reduced higher education subsidies. Unlike other major categories of state spending, such as K-12 education or law enforcement, higher education the state typically directly charges the users of higher education, via tuition. As a result, states can cut funding to higher education expecting that institutions will seek higher tuition, private donations, or Federal research grants to mitigate any adverse effects. Thus, if states face an increase in Medicaid spending beyond their control, they will likely reduce higher education subsidies.

We find evidence consistent with this framework. In our preferred specification, we find that a $1 shock to Medicaid spending reduces higher education subsidies by 20 to 37 cents. The effect on other state outlays appears to be less strong. The effects of Medicaid are robust to different model specifications, and as we argue below, we interpret these estimates as causal. In addition, we show that the response in higher education spending is larger than other categories of spending through which the state might respond, such as K-12 education, transportation (roads, airports, trains), and justice (prisons and police).

One challenge in estimating the main effect in this paper is that there are many possible mechanisms through which Medicaid spending and higher education subsidies may be endogenous at the state level. Crucially, both categories of spending depend on overall state generosity. If more generous states are interested in spending more on both health care and education, the effect that Medicaid costs have on higher education spending will be understated in magnitude. Further complicating the matter is that states may become more or less generous over time.

To overcome these limitations in our analysis, we exploit long-term *Federal* variation in Medicaid policy. In particular, we use changes in Federally administered Supplemental Security Income (SSI) policy to instrument for Medicaid spending. As an instrumental variable, changes to SSI policy have several attractive features. First, Federal statute dictates that SSI enrollment automatically triggers Medicaid enrollment in most states. Second, SSI enrollees are disproportionately more expensive than other Medicaid enrollees, which implies a strong link between SSI enrollment and Medicaid spending. Lastly, there have been numerous changes to Federal SSI policy over the past four decades, which give us the variation we need to identify changes in Medicaid spending.

The estimates above are similar to but smaller in magnitude than previous ones found in Kane and Orszag (2003). In their work, the authors find that a dollar of Medicaid spending crowds out 30 to 50 cents of higher education spending. Kane and Orszag rely on two instruments to predict state Medicaid spending—the share of the population over age 65 and the share of the population below the Federal poverty line (FPL). While these two variables reliably predict total Medicaid spending, it is not obvious that they meet the exclusion restriction assumption necessary for an instrumental variable. That is, there are many other ways age and poverty might affect state budgets other than through Medicaid.

In sum, our paper offers several important contributions to the existing literature. First, we argue that our use of SSI enrollment as an instrumental variable solves many of the co-determination problems with previous instrumental variables approaches to this problem. We are also, to our knowledge, the first paper to digitize state-level SSI and Medicaid enrollment and expenditure records by subgroup since data were first reported in 1975. Next, our approach has numerous methodological improvements, particularly the use of Bartik interactions with SSI enrollment data, which allows us to estimate more precise results. Finally, we analyze heterogeneity in state responsiveness to new Medicaid obligations.

3.2 Background

In this section, we discuss the institutional background of Medicaid, higher education, and state budgets. In particular, we describe trends in Medicaid and noteworthy Federal policy changes that influence these trends. We then discuss the role for state flexibility in Medicaid generosity and differences in the costs of the Medicaid program at the state level. Finally, we discuss state fiscal issues, particularly those relating to the interplay between Medicaid and higher education spending. We discuss the institutional backgrounds of these programs with ultimate goals of motivating the causal relationship and the instrumental variable used to identify it.

3.2.1 Trends in Medicaid Enrollment and Expenditures

Medicaid began in 1965 with a simple directive: to replace two Federally administered medical care programs for the poor, one for welfare recipients (low-income families, blind, and disabled) and one for low-income elderly. Eligibility was restricted to these four groups, and the covered benefits were limited in scope. By 1975, every state but Arizona had adopted Medicaid. Also in 1975, the program enrolled 22 million beneficiaries across the nation, and national expenditures were $49.0 billion in 2010 dollars, or less than 1% of GDP.

Since that time, both enrollment and expenditures have grown considerably. By 2010, Medicaid enrollment tripled to 66 million beneficiaries, and expenditures rose even more to $339.0 billion in 2010 dollars, or 2.3% of GDP. At this pace, Medicaid easily qualifies as the fastest growing of the three largest mandatory entitlement programs, a list which also includes Social Security and Medicare. In addition, it is poised to become the second largest entitlement program behind Social Security. By 2010, Medicaid expenditures have grown to be 76.0% of those for Medicare, and after the Affordable Care Act was implemented, Medicaid expenditures grew to 85.4% of Medicare with only partial state expansions.

The Medicaid and CHIP Payment and Access Commission (MACPAC) describes this trend, along with its root causes, in their June 2016 report. It determined that 70.7% of the growth in real Medicaid benefit spending is due to new enrollment, while the remaining 29.3% is due to growth in spending per enrollee. Moreover, MACPAC decomposed the trend in spending by the four major eligibility classifications: aged, disabled, children, and adults. From this decomposition, we learn that approximately half of the growth in spending is attributed to trends among disabled enrollees. The remaining growth is attributed roughly evenly across the other three eligibility groups.

There are at least two characteristics about Medicaid that led to this type of growth in spending. First, Medicaid is uniquely structured as a joint Federal and state entitlement program. Second, because growth in spending is driven by new enrollment, much of the change in Medicaid spending is the result of new eligibility requirements dictated at the Federal level. We discuss the role these characteristics played in greater detail below.

3.2.1.1 Medicaid's Federal-State Partnership

Unlike Social Security and Medicare, Medicaid is a joint Federal-state program. While the Federal Government establishes mandated minimum guidelines for eligibility, benefits, and cost-sharing (insurance premiums and co-pays), each state administers its own Medicaid program. In addition, states can seek exemptions to certain guidelines through the Federal waiver program. In practice, these waivers tend to expand eligibility and benefits. Moreover, as coverage for particular groups gains critical mass, several coverage expansions are eventually adopted by the Federal government, leading to increases in enrollment and expenditures for all states [Currie and Gruber (1996), Cutler and Gruber (1996), Klemm (2000)].

Funds for Medicaid come from general state and Federal revenue, and Federal funds are delivered as matching grants to the states without a Federal spending cap. The size of the grant is determined using a formula called the Federal Medicaid Assistance Percentage (FMAP), which is described in Section 3.3.

The nature of funding through matching grants has the potential to distort state spending priorities [see Baicker (2005)]. As an example, consider a state with an FMAP of 60%. For each dollar that the state spends, it will receive 60 cents. As such, state policymakers can promise \$1.60 worth of health care at the cost of \$1.00, or they can promise \$1.00 worth of something else—like higher education subsidies—at the cost of \$1.00. Unsurprisingly, this tradeoff leads to more spending on Medicaid over higher education.

3.2.1.2 Federal Eligibility and Benefits Changes to Medicaid and Supplemental Security Income (SSI)

Most Federal expansions in Medicaid eligibility grew coverage for children and adults. By 1988, all pregnant women and young children with incomes below the Federal poverty line (FPL) were eligible for Medicaid. By 1989, children under age 6 at or below 133% of FPL were eligible, and by 1990, children ages 6 to 18 at or below 100% of FPL were eligible. Lastly, the Balanced Budget Act of 1997 (BBA97) created the Children's Health Insurance Program (CHIP), which mandated coverage assistance for children under 200% of FPL.

Because of the number and breadth of these expansions, children and adults represent a growing proportion of Medicaid enrollees—from 64.2% in 1975 to 70.5% in 2010; however, spending per enrollee for these two groups is consistently low relative to the aged and disabled. As such, new enrollment among these groups has contributed only 30.4% of long-run spending growth. By contrast, new enrollment among the disabled comprises 37.6% of long-run spending growth even though its share of enrollment grew by half that of children and adults [see MACPAC (2016)].

Among the pathways to Medicaid enrollment for the disabled, Supplemental Security Income (SSI) is by far the most common. SSI is a Federally administered cash assistance program for people too disabled to earn more than the Substantial Gainful Activity amount ($1,170 per month in 2017). It was designed in 1972 as a basic income program for the aged and disabled, where enrollees would receive a cash stipend from the Social Security Administration (SSA) and health coverage from Medicaid. For the District of Columbia and 34 other states, the SSA automatically enrolls SSI recipients into Medicaid on behalf of states. Seven states require separate enrollment into Medicaid but use the SSI enrollment criteria for eligibility. The other nine states are classified as 209(b) states, and they use a more restrictive eligibility criteria based on pre-1972 levels. Although these states are more restrictive, they are also required to adopt medically needy spend-down programs, which yield similar eligibility results to the other non-209(b) states.

Given the dominant role of the Federal Government in determining Medicaid eligibility among the disabled, changes to Federal SSI policy will also alter enrollment and expenditures for state Medicaid programs with only limited influence by the states. There are numerous instances of SSI changes since 1972, and they have been both contractionary and expansionary [Autor and Duggan (2003), Rupp and Scott (1998), Stapleton et al. (1998)].

The first such change came via the Social Security Amendments of 1980, which were designed to instill programmatic consistency across all states. This legislation strengthened Federal responsibility for disability determinations and state oversight. In addition, the amendments required that the SSA conduct regular reviews of the disability enrollees to see if they remain eligible for benefits. Coupled with a continued decline in elderly enrollment, this added authority led to a decline in SSI enrollment from 1974 to 1982.

Shortly afterwards, however, it became clear that the SSA was not prepared to conduct systematic reviews. This fact was especially true for beneficiaries with imprecisely categorized diagnoses like mental illness or multiple low-level disabilities. After much public outcry and a few lawsuits, the Federal government enacted the Social Security Disability Reform Act of 1984. Among other things, this act required that the SSA establish new standards for qualifying disabilities, especially for mental illness and multiple low-level disabilities. In 1985, the SSA published new rules, which were much more inclusive then previous guidelines, particularly for mental illness. In addition, AIDS was listed as an impairment for the first time in February 1985. As a result of these revised guidelines, enrollment accelerated, especially for non-elderly, mentally ill adults whose numbers grew almost 50.0% from 605,900 in 1986 to 886,400 in 1990 and continued to increase to 1.5 million in 1995 (see Figure B1).

Because the revised guidelines did not apply to them, SSI rolls for children remained stable throughout the 1980s, but in February 1990, the landmark *Sullivan v. Zebley* case required that the SSA evaluate children using similar criteria used for adults. As such, SSI enrollment for children under age 18 soared from 299,000 in 1989 to 970,000 in 1995 (see Figure B2). Most of this increase came from growth in enrollees with mental disorders at

333%, but there was also sizable growth among other enrollees at 144%.

In response to concerns about this dramatic growth, the Federal Government implemented several reforms in the mid- to late-1990s to slow SSI enrollment. First, the Social Security Independence and Program Improvements Act of 1994 required the SSA to perform 100,000 continuing disability reviews (CDRs) for SSI recipients each year for 1996, 1997, and 1998. These CDRs were performed to remove unqualified recipients from the rolls. The law also made it more difficult for SSI recipients whose impairment was either drug addiction or alcoholism, two impairments that had been a part of SSI listings since 1972.

Second, the Contract with American Advancement Act of 1996 permanently removed drug addiction and alcoholism as qualified disabilities beginning January 1, 1997. This reform was somewhat successful. SSI recipients with a diagnosis of "Other," which includes drug addiction and alcoholism, fell by 46.6% from 1996 to 1997. On the other hand, SSI recipients with a diagnosis of "Mental disorder (other than mental retardation)" and "Diseases of the Musculoskeletal System" increased by 11.3% and 18.3%, respectively. This dramatic increase in two categories that are closely related drug addiction suggests that only a subset of the drug addicts was purged from SSI rolls.

Also in 1996, the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA and also known as welfare reform), tightened restrictions for child SSI eligibility. As result, child SSI enrollees fell from 969,900 in 1995 to 879,827 by 1997. Lastly, the Foster Care Independence Act of 1999 implemented several aggressive anti-fraud measures, which reduced child SSI enrollment to a 5-year nadir of 847,000 in 1999.

Without any contractionary measures after 1999 and with the continued use of flexible qualifications for diagnosing disability, especially for mental illness, SSI enrollment continued to increase. From 2000 to 2010, SSI enrollment increased from 5.9 million to 7.1 million (see Figure 3.3). Moreover, this increase occurred despite the continued decline in SSI enrollment among the elderly. As of 2010, they represent only 300,000 of the 7.1 million enrollees, whereas enrollees with mental disorders total 3.4 million, or nearly half of all enrollees.

3.2.2 Medicaid at the State Level

As long as they meet the Federal core requirements for Medicaid, states can choose the generosity of their Medicaid program. To begin with, state participation is voluntary, and it was not until 1982 that Arizona joined the Medicaid program. More recently, 19 states have refused to expand Medicaid as part of the Affordable Care Act. Even if the states decide to participate, they can still provide differing levels of care through the medically needy program and Medicaid waivers.

As part of the original 1965 legislation, states hold the option to create medically needy programs. These programs allow individuals with significant medical expenses to qualify for Medicaid. In addition to being able to choose whether the state has a program, states can also choose which individuals can qualify as medically needy. For instance, they can cover elderly under this program but not the disabled. As of 2017, 32 states and the District of Columbia have medically needy programs in place.

The 1965 legislation also created Section 1115 Demonstration waivers. These waivers are broad in scope and can be used to expand coverage, add benefits, change delivery systems, or reduce costs. They must be approved by the Department of Health and Human Services and must still guarantee the Federal core requirements. They are typically approved for five years with a three-year extension. They must also be budget neutral for the Federal government. As of 2017, there are 42 Section 1115 Demonstration waivers over 29 states and the District of Columbia.

The two other major waiver programs are Section 1915(b) and Section 1915(c), which were created as part of the Omnibus Budget and Reconciliation Act of 1981. The Section 1915(b) waiver program allows states to mandate managed care enrollment for Medicaid recipients. The Section 1915(c) waiver program allows states to offer long-term care services through home- and community-based services (HCBS) rather than through more traditional institutional care. As of 2017, there are 64 Section 1915(b) waivers across 34 states, and there are 302 Section 1915(c) waivers across 47 states and the District of Columbia.

Because of the flexibility that the medically needy program and waivers provide, Medicaid spending per enrollee varies considerably from state to state. In addition, as waivers become available and more popular, the variation in spending per enrollee has increased. In 1977, the standard deviation of spending per enrollee was just $68. By 2010, it had increased to $5,644.

3.2.3 State Finances and Higher Education

State governments typically spend their resources on a small handful of major expenditure categories. A summary of state and local fiscal priorities is included in Table 3.1 below. State fiscal priorities are primarily in education (30%), welfare, and other expenditures such as highways and prisons. Within education, 15% is spent on higher education, versus 85% spent on K-12. Medicaid represents 77% of welfare spending.

According to the National Conference of State Legislatures, every state but Vermont has either a constitutional provision or a statutory requirement requiring a balanced budget (see NCSL (2010)). Given this fact, an increase in mandated benefits mechanically requires an offsetting increase in revenues or decrease in another type of expense. Because state revenues likely affect a broader base and are sensitive to economic conditions, we expect policymakers to prefer reductions in spending over increases in revenue.

We focus on higher education subsidies as the primary target for reductions in spending. Institutionally, only a portion of higher education is funded from state budgets. Higher education institutions collect revenue from the Federal Government in the form of research grants, from students in the form of tuition, and from alumni in the form of donations. In addition, institutions of higher education have some flexible control over their spending, by, for example, taking fewer students in a given year.

Previous empirical works highlight this institutional responsiveness to financial shocks. Kane and Orszag (2003) and Bell (2008) both describe the empirical relationship between higher education spending and short term fluctuations in the business cycle. Bound et al.

(2016) show how higher education programs respond to short term cuts in state subsidies by raising revenue from out of state students, especially international students.

Our paper adopts a slightly different approach. We assume state legislatures understand the flexibility that higher education institutions possess, and as such, they choose to cut higher education subsidies when faced with long-term growth in Medicaid spending. This assumption seems relatively benign when we compare higher education to other line-items in the state budget (see Table 3.1). Prison expenses, for example, are determined by sentencing guidelines and a judiciary which has independence from the legislature. Transportation spending is often directly earmarked from specific revenue sources such as state gas taxes and, as such, cannot be changed easily. Lastly, K-12 education spending is calculated by a pre-determined formula, and K-12 schools serve a broader base of students than higher education. Moreover, because K-12 schools rely more heavily on state budgets as their primary source of funds, cuts will feel more severe and affect more families than cuts to higher education.

We also consider the timing of state policymakers' decision to cut higher education. Most states follow a fiscal year ending on June 30. A total of 16 states practice biennial budgeting, as opposed to the more traditional annual. In practice, however, this distinction makes little difference.

States routinely have to deal with volatility based on the business cycle, and predicting the exact revenue and spending needs over a long period of time is impractical. In addition, state balanced budget requirements mandate that legislators be responsive to this volatility. As such, all states–including the biennial states–rely on supplemental budget amendments to smooth out differences. In his discussion of tax revenue volatility, Seegert (2012) argues that revenue volatility implies expenditure volatility, and shows that compensating responses to shocks are often contemporaneous. In section 3.6.3, we explore the lag structure of the relationship between increased Medicaid obligations and other expenses in greater detail.

3.3 Data

In this section we describe the data used in the final analysis. The majority of our analysis studies the period 1977 to 2010, which represents the intersection of time for available Medicaid, SSI, and state budget data. All data are collected at the state level for each fiscal year. First, we describe the Medicaid data, which are used as the main treatment variables. Then, we describe the SSI data which are used to construct the main instrumental variable. Next, we describe the FMAP formula and data. Lastly, we describe the state higher education, demographic, and financial data.

3.3.1 Medicaid Enrollment and Expenditures

Medicaid enrollment and expenditure data were collected at the state level for each fiscal year from 1975 to 2010. These data came from two sources: the Medicaid Statistical Information System (MSIS) and its predecessor report known as the Health Care Financing Administration-2082 (HCFA-2082). MSIS data were collected for FY1999 to FY2010 and are available online from the Center for Medicare and Medicaid Services. HCFA-2082 data were collected from FY1975 to FY1998. Unlike the MSIS data, these data required independent collection and digitization.

Among all of the Medicaid data collected, our analysis focuses on state-level enrollment and expenditures in total and by reason of eligibility. These reasons consistently classify into four groups: the elderly, low-income children, low-income adults (typically the parents of the eligible children), and the blind and disabled. Because a beneficiary can enroll in Medicaid for multiple reasons, the sum of enrollment and expenditures across eligibility groups will overestimate total; however, the sum of enrollment and expenditures within eligibility groups correctly totals to the national total for a particular eligibility group.

Expenditure data are reported in nominal dollars. Because our preferred specification uses Medicaid expenditures as a share of the state personal income, we seldom need to adjust for inflation. But, on the few occasions when we compare expenditures across periods, we use the Bureau of Labor Statistics's Consumer Price Index-Urban (CPI-U) to normalize expenditures to 2010 dollars for all periods.

There are a few caveats when using these data. First, following the Balanced Budget Act of 1997 (BBA97) and the creation of CHIP, Medicaid enrollment and expenditure values include enrollment and expenditures for CHIP enrollees even if the state operates a stand-alone CHIP that is separate from Medicaid. Because CHIP expenditures are reimbursed at a higher match rate, we will overstate the state share of Medicaid costs from FY1998 onward. As an empirical note, this error will be minor given the small fraction of CHIP expenditures relative to the overall size of Medicaid. Also, Arizona did not enter Medicaid until 1982, and it did not begin reporting MSIS data until 1991. As a result, we exclude Arizona from our analysis. Also, because the District of Columbia has less autonomy than regular states, we omit it from our analysis as well.

3.3.2 Supplemental Security Income (SSI)

SSI data were collected from the SSA's annual statistical supplement for the years 1975 to 2010. These reports provide state-level enrollment as of December for each year. In addition, enrollment is reported in total as well as by reason of eligibility: aged, blind, and disabled. Because the policy for aged eligibility has never changed, changes in aged enrollment over time only reflect changes in state demographics and economic conditions, which we control for in our estimation. Thus, to target the variation that is most closely tied to Federal policy variation, we use blind and disabled SSI enrollment as an instrumental variable for Medicaid expenditures.

Because the SSI data are reported on December of each year, we match each year of SSI data with the following year of Medicaid data. For example, SSI enrollment as of December 1999 is matched with FY2000 Medicaid data, which spans from October 1, 1999 to September 30, 2000.

3.3.3 Federal Medicaid Assistance Percentage (FMAP)

Because Medicaid data report the total expenditures made by both Federal and state governments, we require data on the FMAP to isolate the state-funded portion of Medicaid expenditures. The FMAP is calculated using the below formula for state i, and it is designed to direct more funds to poorer states.

$$FMAP_i = \max\left\{50, \min\left\{\left[100 - \left(\frac{(\text{State per capita income}_i)^2}{(\text{National per capita income})^2} * 45\right)\right], 83\right\}\right\} \quad (3.1)$$

Note that the FMAP is bounded such that Federal funds do not exceed 83 cents for each state dollar spent and do not fall below 50 cents for each state dollar spent. Per capita income is calculated as the three-year average used prior to announcement. Because of lags in data collection and reporting and the need to announce the FMAP rates in advance of each fiscal year, the three years chosen are the third, fourth, and fifth years before the fiscal year of interest. For example, per capita income for the FY2012 FMAP formula is calculated as the average over CY2007, CY2008, and CY2009. Calculated FMAP rates are reported each year in the Federal Register, and they have been collected and reported by the Office of the Assistant Secretary for Planning and Evaluation (ASPE), which is a part of the Department of Health and Human Services (HHS).

These rates, as reported, are the baseline rates only and have not been modified for legislative adjustments, including five major adjustments. First, as part of the Omnibus Budget and Reconciliation Act of 1981 (OBRA81), the Federal Government cut the FMAP rate across all states by 3.0ppt, 4.0ppt, and 4.5ppt for fiscal years 1982, 1983, and 1984, respectively. The one exception is that the statutory floor of 50.0% matching was retained, meaning that a state with a 50.0% FMAP and a state with a 52.0% FMAP would both receive 50.0% matching in these three years. Second, as part of BBA97, the District of Columbia's matching rate was increased from 50.0% to 70.0% beginning in FY1998. Third, BBA97 also created CHIP, and for any spending on CHIP-eligible children, the Federal

Government paid an enhanced FMAP ranging between 65.0% and 88.1%. Fourth, as part of the Bush stimulus known as the Jobs and Growth Tax Relief Reconciliation Act of 2003 (TRRA), the FMAP temporarily increased from April 2003 to June 2004. Most states' FMAPs increased by 2.95ppt over this period, and 21 states received slightly more than that. Lastly, the American Recovery and Reinvestment Act of 2009 (ARRA) temporarily increased the FMAP rate from October 2008 to June 2011 by 6.2ppt plus an additional unemployment-related increase. In practice, the average state's FMAP increased by 10ppt, and the range was between 6.2ppt and 17.87ppt.

To more accurately represent the proportion of state funding, we modify the baseline FMAP rates to reflect all these adjustments except for the CHIP adjustments. To properly adjust for CHIP enhanced rates, we need additional information on CHIP-related expenses, which are not reported separately in the Medicaid data. A detailed distribution of adjusted FMAP rates by state are shown in Figure D1. Each point represents the FMAP rate for a particular state in a given year.

3.3.4 Education, Demographics, and Finance

Education data are take from the Center for Education Policy Study at Illinois State University ("Grapevine"). Grapevine data were chosen over other data sets such as the Delta Cost database ("Delta") due to the length of availability. Whereas Delta only exists for the period following 1986, Grapevine data extend back to 1977. Data included are the direct subsidies of the higher education system coming from the state budget. That is, this excludes items such as tuition and research grants.[1]

In addition, we take state demographic information from the annual March Current Population Survey. This includes average state age, share of the population over 65, share of the population under 24, and state level race and ethnicity. Information of state per-capita

[1] We take these measures to capture the consumption value of higher education, either at the personal level (tuition) or transactionally at the state-level (research funding). Limiting the analysis to subsidies more directly captures the public investment of spending.

income and state-level program spending are taken from the Bureau of Economic Analysis.

3.4 Methods

In this section we describe the main analysis used in this paper. First, we discuss the OLS model and the challenges in causally interpreting the results from this specification. Next, we discuss a possible solution to the identification problem in the OLS model. We then describe the preferred two-staged least squares specification. In Section 3.6 below, we discuss robustness tests and extensions. Overall, we find the results to be robust to a variety of model assumptions.

3.4.1 The Key Causal Relationship

We are interested in the causal relationship between total Medicaid spending at the state level and total state subsidies for higher education. In the following equation for state s and year t:

$$Y_{st} = \gamma Medicaid_{st} + \beta X_{st} + \theta_s + \delta_{st} + \varepsilon_{st} \tag{3.2}$$

The outcome variable, Y_{st}, and the explanatory variable of interest, $Medicaid_{st}$, represent the share of state income spent on higher education and on Medicaid, respectively.[2] Because Medicaid is a Federal-state partnership, we measure only the share of total spending for which the state is responsible.[3] Higher education subsidies are measured as the total grants towards institutions of higher education, rather than costs borne by the users of higher education such as tuition.

X_{st} represents a vector of state-year characteristics likely to be correlated with education

[2] Because incomes are growing over time (in both a nominal and per capita sense), both Medicaid and higher education subsidies are also growing over time, albeit at different rates. Using income shares properly captures the displacement effect on the state's consumption bundle.

[3] Specifically, the total reimbursements for Medicaid patients in the state times one minus the state's Federal matching rate for that year.

spending. These characteristics include each state's share of the population by race and age, mean age, and business cycle environment, which is measured by the lagged state unemployment rate. We chose these covariates over other candidates, such as total income or education shares that could plausibly be interpreted as an outcome. The measures θ_s and δ_{st} capture the levels and trends in unobserved state characteristics, respectively. Finally, ε_{st} represents the remaining unexplained state-year variation in higher education spending.

Under specific conditions, the parameter γ identifies the causal relationship between the additional share of state income spent on Medicaid spending, but Medicaid spending is likely correlated with state preferences for higher education spending and thus violates the exclusion condition ($E[\varepsilon_{st}|Medicaid_{st}] = 0$). First, states with more generous attitudes towards social spending may be likely to spend more on both Medicaid and higher education, which would bias γ upward. Moreover, because states have some flexibility in Medicaid spending through waiver programs, a decision to adjust higher education subsidies could affect Medicaid spending and not the other way around. Thus, a causal interpretation of γ will require variation in Medicaid which is not attributable to state voter preferences or state policy consideration.

One straightforward solution to this threat is to include state fixed effects and state-specific trends, which should control for average state preferences over our sample period and linear growth in state spending over time. Including these measures, however, still does not address all possible concerns regarding endogeneity. Crucially, because we are looking at long-run effects over several decades, state preferences are not guaranteed to be constant for the entire observation period, nor to change in a monotonic, linear manner. For example, if a state's preference for social safety nets grows exponentially relative to Medicaid growth—possibly as a critical mass of voters enroll—we would still expect Equation 3.2 to downwardly bias the absolute magnitude of γ. Alternatively, state preferences for generosity could contract as more enrollment in Medicaid increases salience and antipathy for redistributive programs. This type of variation would increase the magnitude of γ.

Previous efforts to address this problem (see Kane and Orszag (2003)) have used demographic-driven variation among the poor and elderly. This approach is limited because it implicitly assumes that the share of poor and elderly will have similar effects throughout the 34-year sample period. This assumption cannot hold. People can migrate to states because of state benefits policy, which introduces simultaneity concerns. In addition, these concerns may compound as people's preferences for government spending and proclivity to migrate change over our sample period. Lastly, and perhaps most importantly, the composition of people who are poor and elderly has changed over time. For example, in regards to the elderly instrument, in 1977 all seniors pre-dated the Social Security Administration, but by 2010 some seniors were born following World War II. Such dramatic societal changes as a backdrop necessitate the search for more plausibly exogenous variation.

3.4.2 The Identifying Variation

As an alternative, we propose to use variation in Medicaid spending driven by per capita blind and disabled Supplemental Security Income enrollment (henceforth SSI enrollment). Unlike demographics, SSI enrollment is an ideal instrument for Medicaid spending because it is both Federally administered and directly tied to Medicaid enrollment. Because it is Federally administered, changes in state policy and state voter preference will not affect changes in SSI policy. In addition, Section 3.2 indicates that SSI enrollment is a primary driver of Medicaid spending over the past four decades. The combination of these two factors make SSI both a plausibly exogenous and strong instrument.

One concern with this instrument would be if the variation were purely Federal. If this were true, changes in SSI policy would affect all states equally and would eliminate the cross-sectional variation that we need to identify cross-sectional changes in Medicaid spending. Empirically, this issue does not appear to be a concern. Figure 3.4 illustrates the variation in SSI enrollment over each year in our sample period after adjusting for controls, state fixed effects, and linear state trends. Even after these adjustments, there remains considerable variation across states and over time. SSI enrollment varies from a range of 600 enrollees per 100,000 people in 1996 to 1,600 enrollees per 100,000 people in 1977.

Even after controlling for demand drivers and state characteristics like we do above, we should expect to see this variation given the disability determination process. The first point of disability determination occurs at local Social Security Administration offices. In practice, this role means that they enroll applicants with a clear-cut disability and deny borderline cases. Even this relatively formulaic first pass results in some variation (Bound and Burkhauser (1999), Daly and Burkhauser (2003) Maestas, Mullen and Strand (2013)). The remaining applicants go through an appeals process, which can include a hearing by an administrative law judge and adds another layer of idiosyncratic enrollment (French and Song (2014)). Finally, the guidelines for determination at all levels have changed over time, and the responses to these changes by actors in the determination process will introduce more cross-sectional variation.

3.4.3 The Causal 2SLS Model

The preferred causal model for this analysis is the two-stage least squares (2SLS) instrumental variable model shown below for state s and year t:

$$Medicaid_{st} = \omega SSI_{st} + \alpha X_{st} + \rho_s + \pi_{st} + \nu_{st} \tag{3.3}$$

$$Y_{st} = \gamma \widehat{Medicaid}_{st} + \beta X_{st} + \phi_s + \delta_{st} + \varepsilon_{st} \tag{3.4}$$

In this model, the key causal parameter of interest is still γ, which now measures the change to state higher education subsidies caused by changes in SSI-related Medicaid spending. The terms ρ_s and ϕ_s are state level fixed effects, and π_{st} and δ_{st} are linear state trends. X_{st} represents the vector of state-year characteristics described for Equation 3.2 above. As with Equation 3.2 both Medicaid and higher education spending are normalized as a share of total personal income for the state, whereas SSI enrollments are normalized as a share of total population.[4]

Given the direct link between SSI and Medicaid enrollment and the relative expense of disabled enrollees, we see a strong positive relationship between the instrumental variable (SSI enrollment) and the endogenous regressor (total Medicaid spending). Figure 3.5 visualizes this relationship using the specification outlined in Equation 3.3. The effect of SSI enrollments is economically significant, with a 1ppt increase in percent of the population enrolled yielding a 0.2ppt increase in total personal income in the state spent on Medicaid. The relationship does not appear to be driven by outliers. More conventionally, Table 3.2 shows that the F-statistic ranges from 11 to 117. Our preferred specification with state fixed effects and linear state trends has an F-statistic of 43.

Note that equations 3.3 and 3.4 do not include year fixed effects. While we discuss year fixed effects in more detail in section 3.6.3, the exclusion of year fixed effects in the main analysis means that much of the variation will be driven by differences over time. That is, Medicaid is growing not just because different states have different rates of SSI growth, but also because all states experience SSI growth over time. This creates an additional threat that independent secular trends in tastes have driving these changes in higher education. We mitigate this threat with the inclusion of state-specific linear trends, and state by year observable variables. We also discuss risks related to specific Medicaid policy changes in section 3.6.3.

[4] Note that as discussed in Section 3.2, states do not universally accept Federal SSI guidelines as the basis of SSI enrollment eligibility. In order to address the potential confounding effect of these policy differences, we also run an additional specification which interacts the SSI enrollment instrument with a binary indicator for the 11 states which do not conform to these guidelines. This allows us to have two instruments, and in theory would capture any systemic difference between Medicaid costs from conforming and non-confirming states. In practice, including this additional instrument makes no difference, suggesting that non-conforming states are similar enough to conforming states, on average, as to not make this additional step necessary.

3.5 Results

We begin this section by discussing the main specification. We discuss OLS results, 2SLS results, and overall model sensitivity, for both the main outcome (higher education spending) as well as other state budget items.

3.5.1 Main Specification

The main specification is reported in Table 3.3, which report the OLS and 2SLS estimates of the effects of state Medicaid obligations for three regression specifications. The first estimate (Column 1) includes only lagged state unemployment, to control for business cycle effects on state budgets. The second model (Column 2) adds state demographic controls, including race and age. Columns 3 and 4 repeat Columns 1 and 2, but with the addition of state fixed effects. Including state fixed effects controls for differences in state higher education spending which are fixed between states. This might include state political preferences or generosity towards welfare systems. The final two specifications (Columns 5 and 6) repeat Columns 3 and 4, but add state specific linear time trends. These trends control for changes in a given state over time, such as a differential state trend in demographics or wealth.

The results across model choices generally behave as anticipated. Adding state fixed effects increases the (negative) coefficient in absolute magnitude, for both OLS and 2SLS. In addition, we find that the 2SLS estimates are consistently larger than the OLS estimates. Both of these results are consistent with our expectation that state generosity and preferences for public spending are a significant confounder. Importantly, in every model with fixed effects (Columns 3-6), the 2SLS estimates are larger in absolute magnitude than the OLS estimates, suggesting that these preferences are not simply constant over time.

The estimates range from a \$0.09 to \$0.37 reduction in higher education spending for each dollar spent on Medicaid. These estimates are somewhat lower than those used by Kane and Orszag, which range from \$0.37 to \$0.53. As part of this analysis, we also replicated the Kane and Orszag instrument using our specification and panel, and found results comparable to those reported in their paper. Overall, we argue that our lower point estimates reflect that we are better isolating the variation in Medicaid spending and not picking up other preferences driven by age or economic conditions, which may have contaminated the instrument in the Kane and Orszag paper.

Our preferred specification is the 2SLS model, including state fixed effects as well as state specific linear trends. In this model, every additional dollar of additional state Medicaid obligations lowers state grants for higher education by \$0.33. Overall, we find that state specific demographic controls and state specific linear trends are close substitutes, as the point estimates only including one or the other are comparable to those models including both. This suggests that the state trends meet our goals of controlling for unobserved state characteristics.

In Table 3.4, we repeat the analysis but allow other state expenditures to take the place of higher education subsidies. Specifically, we look at "justice" (police, courts, prisons), "transportation" (primarily roads, but also trains and airports), and "K-12 education." Looking at these other major spending categories allows us to see why higher education is a uniquely flexible component of state budgets, and thus the preferred outcome of interest.

Within each of the tested spending categories, we fail to replicate the same powerful economic and statistical relationship with Medicaid, which persists for higher education. Justice spending is a precisely estimated zero—we can say with confidence that this does not respond to higher Medicaid costs. K-12 education spending is low, increasing only 3 cents for an additional dollar of Medicaid spending, but imprecisely estimated. Transportation is the only category that produces a large point estimate. The preferred specification suggests that a dollar of Medicaid spending reduces transportation spending by 18 cents. As with K-12 spending, this estimate is unfortunately lacking the precision to draw inference.

Each of these types of spending behaves differently than higher education because of the difference in institutions and political economy. Justice spending is driven by factors such as crime rates and convictions, over which the state has limited short term power. Unlike higher education, K-12 education is universal and, thereby, more politically sensitive. The size of and large standard error for the transportation estimate is likely a result of heterogeneity in state institutions. Some states have a separate fund for transportation spending, which is determined by gas tax revenues. Other states rely more on the general fund for highways, and as a result are more likely to respond to increased Medicaid costs.

In the sections below, we will discuss the robustness of these estimates, with regards to both instrument selection and model uncertainty. In particular, we emphasize the model sensitivity to year fixed effects, which suggest that much of the variation is driven by differences over time. Finally, we discuss the heterogeneity in results by state subgroup.

3.6 Robustness and Extentions

3.6.1 Technological Endogeneity

One potential threat to the above 2SLS specification is that SSI enrollment may still not be endogenous. While it seems unlikely that states can have any meaningful control of SSI enrollment, they may be able to control its relationship with state Medicaid spending via its management of spending per enrollee. For instance, innovative use of Home and Community Based waivers (Section 1915(c)) may successfully reduce cost per disabled enrollees for a particular state. If for some reason this technology does not transfer to other states, SSI enrollment would generate differential effects on Medicaid spending, and more importantly, these differences are potentially within the purview of state policymakers.

To explore this threat, we propose adapting the Bartik instrumental variable (Bartik (1991)) to model SSI-related Medicaid spending for each state-year observation. The purpose of this approach is to use national variation in technology while excluding state-level variation.

We achieve this objective through the below specification:

$$Z_{st} = SSI_{st} * \frac{DisabilitySpend_{-s,t}}{DisabilityEnroll_{-s,t}} \quad (3.5)$$

Here, $DisabilitySpend_{-s,t}$ and $DisabilityEnroll_{-s,t}$ refer to total disability patient expenses and enrollments in states outside of state s. We take the ratio of these two values to construct an out-of-state measure of per disabled enrollee costs, and we then interact that measure with SSI disability enrollments for state s in year t. The resulting instrumental variable (called a Bartik IV) represents the predicted Medicaid spending for state s if we exclusively use out-of-state spending technology. Z_{st} replaces SSI_{st} in Equation 3.3 when using the Bartik IV.[5]

In addition to addressing the endogeneity concerns mentioned related to within state patient costs, there are additional benefits to using the Bartik IV. As it explicitly models changes in Medicaid costs, using this instrument will also improve the power of Equation 3.3. It also reduces the threat of a year-specific shock in per-patient medical costs driven by new procedures or technology.

The results from the Bartik IV are included in Table 3.5. In general, these results are both lower in absolute magnitude and more precise than the estimates described in Table 3.3, using the SSI enrollment instrument. The preferred model using the Bartik instrument (column 6) shows a reduction in higher education spending of -\$0.198 for every additional dollar in Medicaid spending (as compared to -\$0.326 in table 3.3). In this case, the estimated standard error is 0.039, considerably lower than the standard error of 0.154 in Table 3.3. While we cannot conclude that the point estimate using the Bartik IV is statistically excluded from the model using the SSI instrument, we prefer the Bartik IV estimate, both for its more credible exogeneity and for its more precise results.

Notably, the key relationships between the models persist in the Bartik specification. In particular, the point estimates are larger with fixed effects included, larger in 2SLS than

[5] We note that this is not, strictly speaking, a Bartik instrument, which measures industry driven demand shocks. It is, however, "Bartik-inspired", as an interaction of national and state trends meant to isolate out state actions.

OLS, and relatively insensitive to the inclusion of state-specific linear time trends.

3.6.2 First Differences

Next, we consider a differenced regression model, given by:

$$\Delta Medicaid_{st} = \omega \Delta SSI_{st} + \alpha \Delta X_{st} + \rho_s + \pi_{st} + \nu_{st} \tag{3.6}$$

$$\Delta Y_{st} = \gamma \Delta \widehat{Medicaid}_{st} + \beta \Delta X_{st} + \phi_s + \delta_{st} + \varepsilon_{st} \tag{3.7}$$

This is the same as equations 3.3 and 3.4, differenced by year to produce implied state fixed effects. There are two main benefits to running this model in addition to those already shown. First, this model is responsive to short-run changes in Medicaid costs. That is, while the fixed effects model compares current year Medicaid and Higher Education costs to the state average over the total time period, the differenced model is more focused on discrete, year-over-year changes. In addition, the differenced model allows for a more flexible state-specific trend controls. While the differenced model implicitly includes state fixed effects, adding state fixed effects to this framework (as in ρ_s and ϕ_s in Equations 3.6 and 3.7) is the equivalent of adding trends. Adding trends to this framework (as in π_{st} and δ_{st} in Equations 3.6 and 3.7) is thus more flexibly (non-linearly) controlling for state-specific trends.

The results of this model are included in Table 3.6. The preferred point estimate in Column 6 shows that a dollar in additional Medicaid costs reduces higher education spending by $0.343. This is comparable to the basic SSI specification in Table 3.3, and somewhat larger than the Bartik IV estimates. The delta estimates are fairly insensitive to fixed effects, trends, and year effects (see 3.6.3). In addition, these estimates suggest that the linear time trend is a reasonable approximation for more flexible options.

3.6.3 Time-Specific Variation and Lagged Effects

In this section, we discuss two final possible concerns. First, we address the concern that the increase in Medicaid costs and the decrease in higher education subsidies are both secular, unrelated trends. Then, we briefly discuss the lag structure of the treatment variable and how this might affect our estimates.

Table 3.7 shows the SSI-enrollment instrument results from Table 3.3, with three additional specifications (Columns 2, 4, and 6) added to include year fixed effects. Including year fixed effects dramatically increases the standard errors of these estimates. The point estimate of the preferred specification is near zero, but so noisy as to be essentially meaningless. In some sense, this result is to be expected. Both the increase in Medicaid costs and the decrease in higher education subsides are relatively slow-moving, long run trends. The fully saturated set of year fixed effects is too closely correlated with variation in annual increases in SSI enrollment to generate meaningful results. Nevertheless, without year fixed effects, we fail to test whether the threat of spuriously correlated long run trends is a valid concern. Given this result, we must interpret these results to be driven by time-series variation, rather than strictly by state-by-year variation in the implementation of SSI policy.

We attempt to address this threat parametrically. As discussed in Section 3.2 we identify several changes to the Medicaid and SSI programs. The two changes that probably had the largest effect on the rate of Medicaid spending were the 1985 adjustments to disability determination, which dramatically grew disabled enrollees, and the mid-1990s retrenchment from 1994 to 1996. As the main episodes of significant time series changes in Medicaid spending, we add two fixed effects identifying years that follow these policy changes, which we call "Policy Period Fixed Effects." We also add in three year-specific fixed effects, to control for the first three years following each policy change. These additional controls, in combination with our flexible state-specific trends, should allow us to address this time series threat without over-saturating the model.

The results from this specification are included in Table 3.8. This table is comparable to Table 3.7; however, it uses less flexible controls in lieu of year fixed effects. The point estimate here is somewhat larger than those without time period fixed effects (-\$0.373 versus -\$0.326, respectively), but neither estimate statistically excludes the other.

As a final note, we also consider the lag structure of the relationship between these two variables. As noted in Section 3.3, SSI data is reported in December, which would be after the end of the federal fiscal year. As a result, we match the prior calendar year SSI data to Medicaid enrollment data when running stage one of the 2SLS Model (that is, SSI enrollment in December of 1995 will match with medicaid enrollment in September of 1996). Here, we seek to ask whether a similar lag is appropriate for matching Medicaid costs to the state budget. Intuitively, a lag may make sense, as states must sign off on budgets before having full knowledge of the annual Medicaid costs. Ultimately, because states have constraints such as balanced budget requirements and flexible tools such as supplemental budget amendments, the lag timing is an empirical question. We find that including either current year or lagged Medicaid as the treatment variable produces similar results, again emphasizing the long run nature of this problem.[6]

3.6.4 Heterogeneity

As a final extension, we consider how the willingness of the state to reduce Medicaid spending may be heterogenous across states with different underlying characteristics. In order to do this, we group states according to two different criteria. The first criteria is the share of total residents in 1980 with some college experience. The measure is meant to capture the political will in the state to finance higher education, the local economy's reliance on human capital, and the existing baseline of investment. The next criteria is the share of the state's population over 65 in 1980. This is meant to measure the political

[6] While it may seem surprising that current and prior year treatments produce the same general effect, consider the timing of all three variables. Medicaid is measured at September, SSI is measured at year end, and higher education is measured as of June. As a result, any pairing will result in some degree of matching and some degree of noise. Extending the last back beyond one year noticeably reduces estimates and makes them less precise. At three years, the effect disappears entirely.

will for spending on human capital. Poterba (1996) argues that higher shares of an elderly population reduces investment in K-12 education. It is intuitive that the same dynamic may exist for higher education.

For each measure, we group states into terciles and run Equations 3.3 and 3.4 by each subgroup. We use the Bartik IV to maximize the power of these regressions, given the reduced sample size within each tercile. The results of this subgroup analysis are available in Table 3.9.

The sensitivity of higher education spending to additional Medicaid obligations is increasing in the share of the population with a college degree. For the lowest group of states, an additonal dollar of Medicaid results in a 7 cent decline in college subsidies. For the group with the most well educated baseline, an additional dollar results in a 42 cent reduction. This result could be interpreted as a form of mean reversion. The high baseline states were likely already spending a large amount of their total earnings on higher education. Thus, when pressed, they could adopt the strategies of other states and lower their overall spending. Alternatively, states which were already poorly educated could have seen preserving their spending on higher education as a higher priority.

In the bottom rows of Table 3.9, we see the sensitivity of higher education spending to additional Medicaid obligations is decreasing in the share of the population over age 65. While this result might seem counter-intuitive, it is again possible to interpret it through the lens of mean-reversion. If states with a larger shares of the elderly are already spending less on higher education, there is less for the state to cut.

3.7 Conclusion

Overall, we find that an additional dollar of Medicaid obligations reduces spending on higher education by 20 cents, in our preferred model. This number is both precisely estimated, and large enough to say that the financial obligations imposed on state governments by the Medicaid program are economically meaningful. At the same time, the estimate is smaller than those identified previously in the literature. This general finding is robust to a variety of modeling assumptions.

We argue that this paper makes a number of important contributions to the literature on this topic. First, we use state-level data relating to SSI disability enrollments covering the entire Medicaid period. This is both a data contribution, which involves digitizing older records, and a methodological improvement, as SSI is more credibly exogenous to other determinants of state higher education spending than instruments used in previous studies. This paper also integrates Bartik-style instruments into this approach, improving the power of our estimates and addressing an additional potential threat related to endogenous per-enrollee cost.

Finally, we find that the results are highly sensitive to the state's baseline level of spending for higher education. Thus, we see that in addition to lowering the overall spending on higher education, the growth of state Medicaid obligations also has the effect of compressing the distribution to total Medicaid spending across states.

Understanding the trade-offs between Federal and state fiscal policy will be an important consideration for policymakers going forward. While the Medicaid program undoubtedly helped people in both the short and long-run [Wherry et al. (2015), Cohodes et al. (2016)], it also indirectly imposed costs on human capital development and out of pocket spending for higher education. Unfortunately, the relationship between higher education and Medicaid identified in this paper is only one step in the process of understanding these relationships. Further research must show how the decline in state subsidies affected higher education at the individual level, in terms of enrollment, attainment, and out of pocket expenditures.

The literature in public benefits has previously emphasized weighing the costs of offering benefit eligibility to an imperfect candidate against the costs of denying benefits to someone who may need them [Kleven and Kopczuk (2011)]. This paper cautions that there are deeper dimensions to this type of problem both across substitutable public services and over time. Even the smallest change to a public program's rules could lead to disproportionate and persistent effects elsewhere.

3.8 Figures

Figure 3.1: Trend in State Spending on Medicaid and Higher Education, Share Tax Revenue

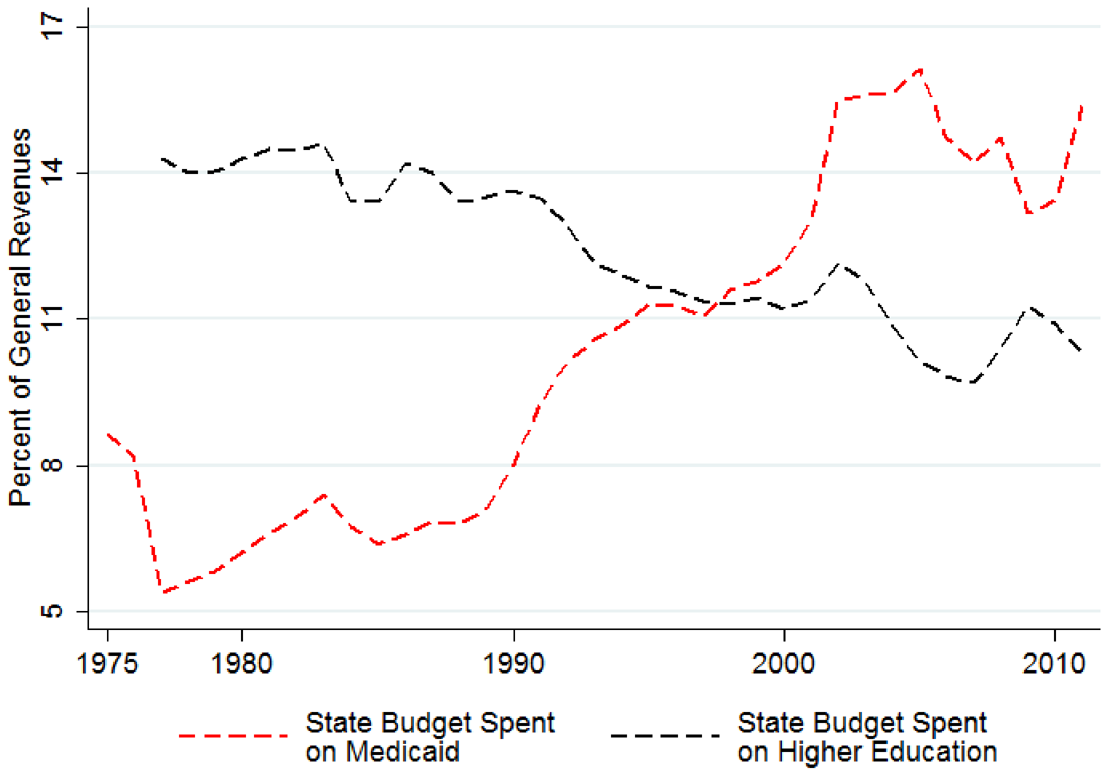

Source: Census of Governments and Bureau of Economic Analysis (1975-2011). Figure shows the comparative trend in the share of state spending directed toward higher education vs. Medicaid. Over time, an increasingly large share of general state revenues are directed towards Medicaid, while the share dedicated to funding institutions of higher education drops over time.

Figure 3.2: Trend in State Spending on Health Care and Higher Education, Per Capita

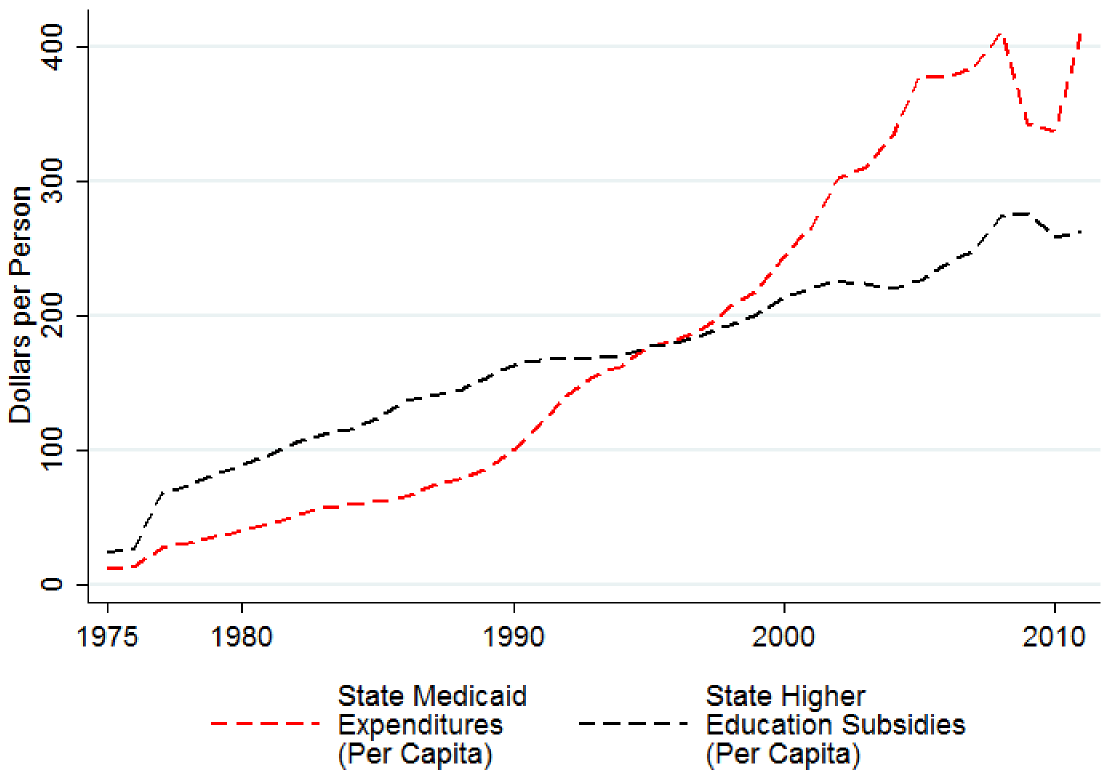

Source: Census of Governments and Current Population Survey (1975-2011). Figure shows the comparative trend in the share of state spending directed toward higher education vs. Medicaid. Over time, an increasingly large share of general state revenues are directed towards Medicaid. Higher Education increases on a per capita basis. This differs from Figure 3.1 because of the growth in the size of the economy along with overall population growth. The difference highlighted by these figures motivates using an income-based, rather than population-based normalization in this analysis.

Figure 3.3: Trend in Total SSI Recipients by Diagnosis

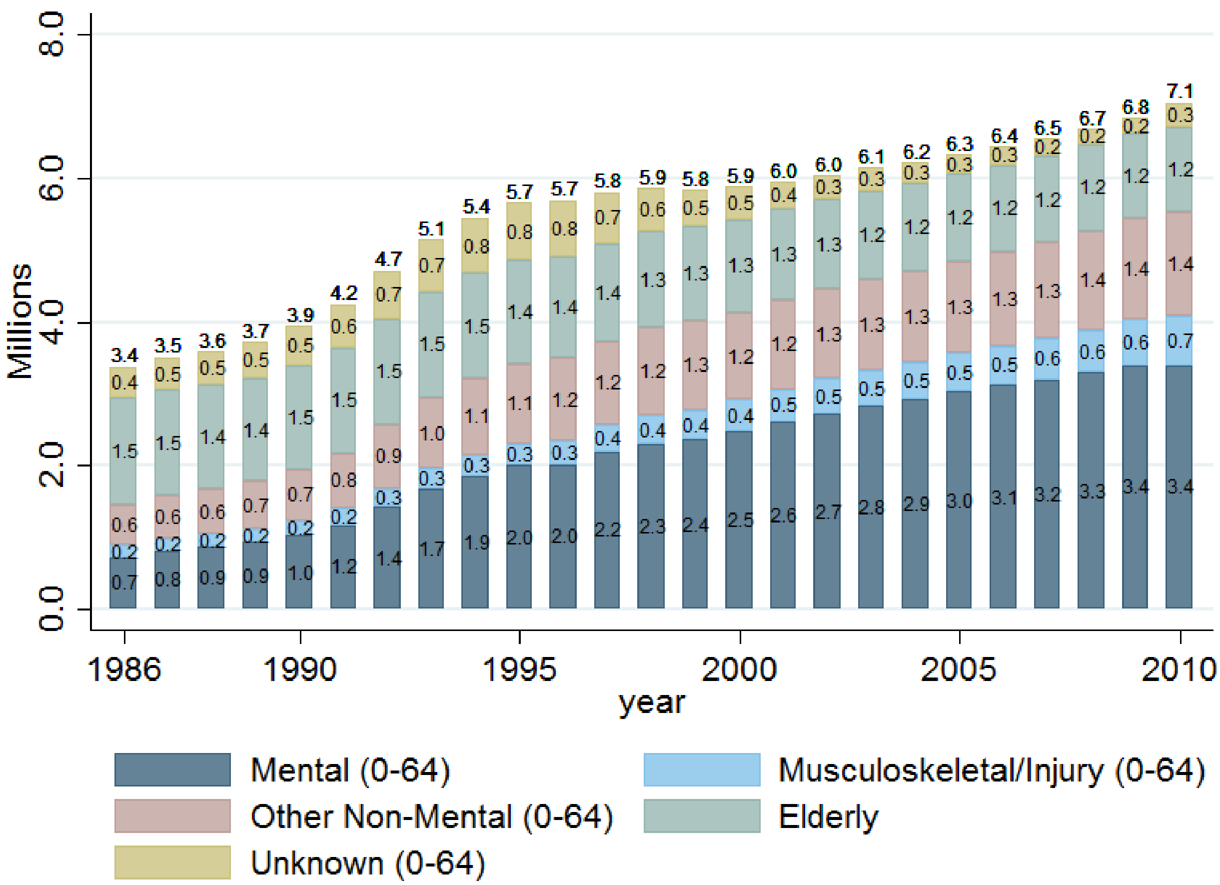

Source: SSA Annual Statistical Supplement (1986-2010). This figure demonstrates that the disability, particularly mental health disability, is the dominant source of growth in SSI recipients. Importantly, despite the aging of the baby boomer generation, the elderly actually make up a decreasing share of SSI recipients over this time frame, due to reductions in elderly poverty. Elderly recipients are exclusively ages 65 and older. They do not need a diagnosis to enroll in SSI as long as they meet the income and asset eligibility requirements. Recipients with an associated disabling diagnosis are all 64 years old and younger. The lowest group in the figure is Mental (0-64), and the second lowest group in the figure is Musculoskeletal/Injury (0-64).

Figure 3.4: Distribution of SSI Enrollment Over Time

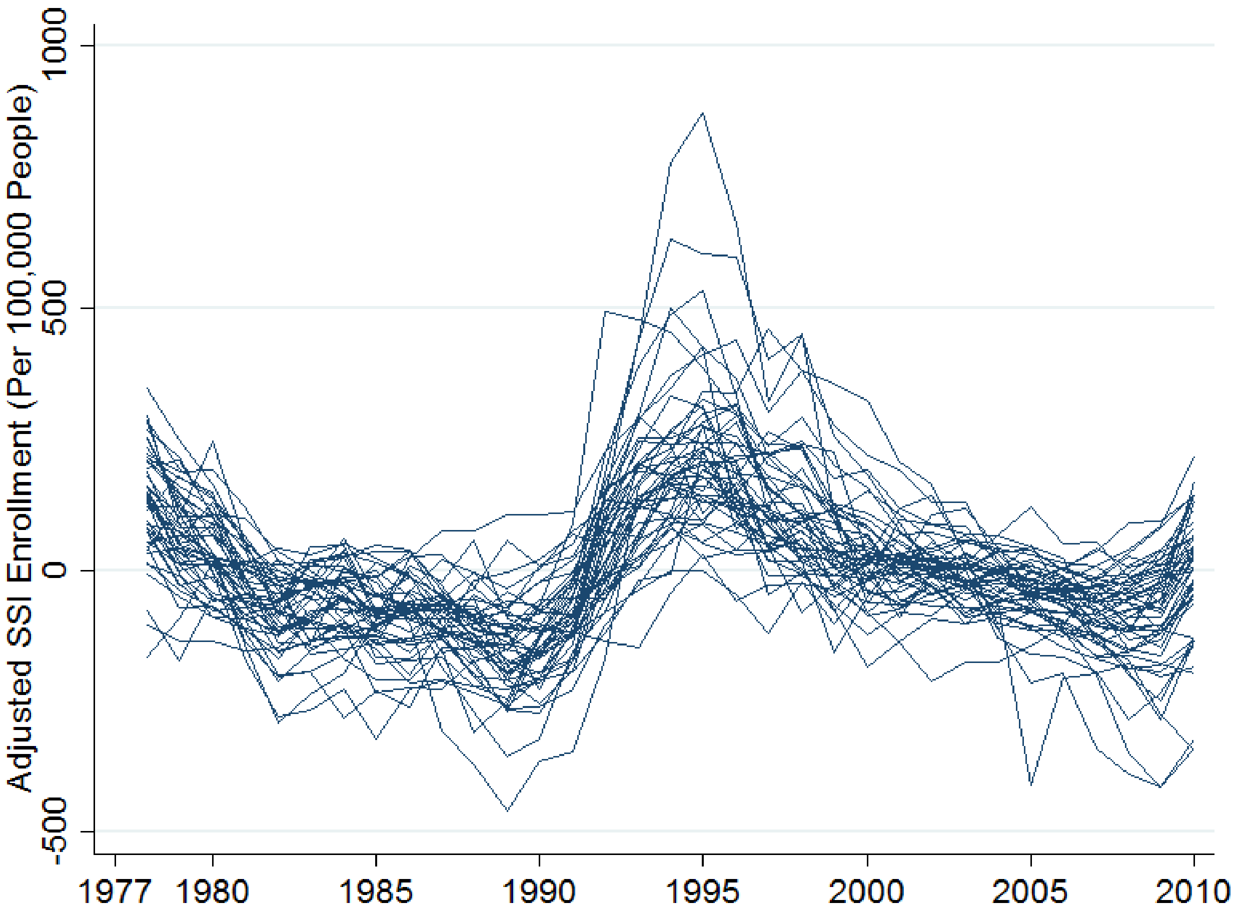

Source: Authors' analysis. The above figure shows the time series in blind and disabled SSI enrollment per 100,000 people in each state, for each state. Values have been adjusted for state fixed effects, linear state trends, demographic characteristics, and the business cycle. In any given year, enrollment ranges from 600 to 1,600 per 100,000 people. Despite SSI eligibility being driven by Federal policy, the eligibility rules can be open to administrator interpretation, which drives idiosyncratic variation at the state-year level.

Figure 3.5: Relationship Between SSI Enrollment and Medicaid Spending

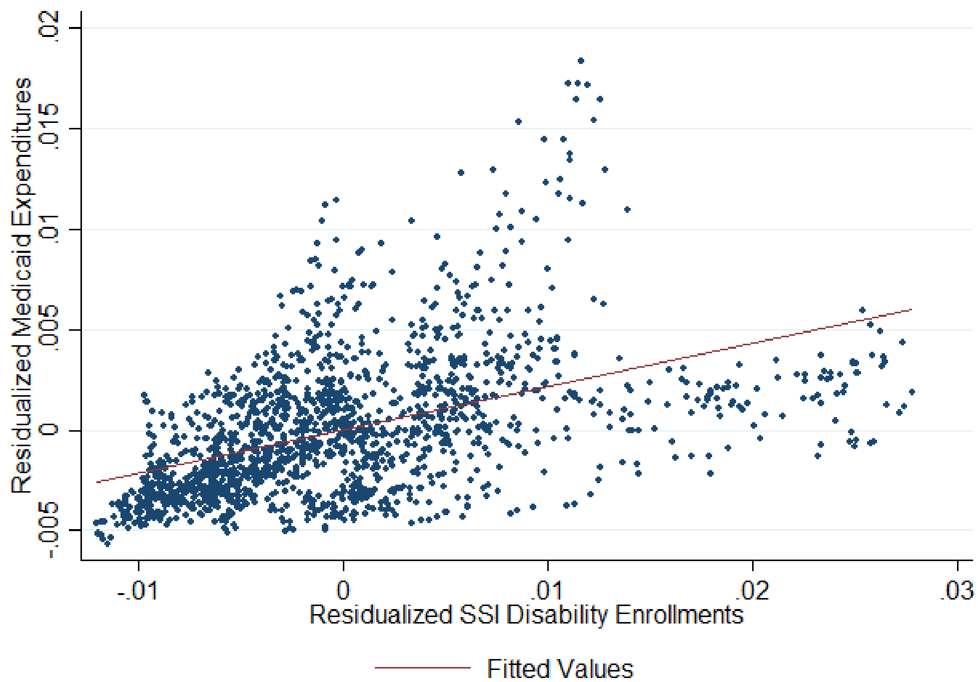

Source: Authors' analysis. Figure shows state Bartik IV (SSI enrollment times out of state average disabled patient cost) and total state-level Medicaid payments plotted together after being adjusted (residualized) for state fixed effects, linear state trends, demographic characteristics, and the business cycle. The effect of SSI enrollments is economically significant, with a single new dollar in the Bartik measure increasing state Medicaid obligations by 22 cents. This relationship does not appear to be driven by outliers.

3.9 Tables

Table 3.1: State and local expense by functional category

State and Local Functional Spending	1977	2010
General public service	17%	16%
Public order and safety	10%	13%
Economic affairs	11%	8%
Health (net)	12%	21%
Education	39%	33%
Income security	10%	7%

Source: BEA National Economic Accounts data, table 3.16. Local expenses are concentrated in categories for Public Order and Safety (Police / Fire), and Education (Particularly K-12). Table shows growth in total spending on Health, and the largest decrease in Education.

Table 3.2: First stage regressions

	Model					
Stat	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient	0.213***	0.154***	0.433***	0.270***	0.209***	0.223***
SE	(0.048)	(0.045)	(0.041)	(0.035)	(0.036)	(0.034)
F	17.972	11.075	117.044	62.468	33.666	42.716
r2	0.177	0.462	0.774	0.846	0.933	0.938
N	1666	1666	1666	1666	1666	1666
Years	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010
Lagged Unemployment	x	x	x	x	x	x
Demographic Controls		x		x		x
State FE			x	x	x	x
State Trends					x	x

Table shows the power of first stage regression results of Equation 3.3. The coefficients represent the effect of the instrument, SSI enrollment per capita, on the state Medicaid spending, measured as a share of total personal income. As an example, the specification shown in Column (6) indicates that a 1ppt increase in SSI disability enrollment per capita increases the share of total income spent on Medicaid by 0.2ppt. The F statistics are the reported first stage F statistics from the 2SLS regression. They indicate a strong relationship between the instrument and Medicaid spending. Standard errors are clustered by state.
*** p<.01, ** p<.05, * p<.10

Table 3.3: Effect of state Medicaid obligations on Higher Education Subsidies

	Model					
	(1)	(2)	(3)	(4)	(5)	(6)
Higher Ed Spending (OLS)	-0.297***	-0.081	-0.305***	-0.159***	-0.100***	-0.089***
	(0.052)	(0.066)	(0.035)	(0.042)	(0.031)	(0.029)
Higher Ed Spending (2SLS)	-0.071	0.676**	-0.351***	-0.215***	-0.370**	-0.326**
	(0.189)	(0.315)	(0.037)	(0.052)	(0.158)	(0.154)
N	1666	1666	1666	1666	1666	1666
Years	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010
Lagged Unemployment	x	x	x	x	x	x
Demographic Controls		x		x		x
State FE			x	x	x	x
State Trends					x	x

Table shows the effects of state Medicaid spending on higher education subsidies. Higher education subsidies and Medicaid obligations are both expressed in terms of the share of total state income spent on these items. Medicaid obligations are instrumented by the share of the population eligible for SSI disability benefits. Column 6 in the 2SLS specification shows that an additional dollar in Medicaid obligations lowers state higher education subsidies by about 32.6 cents.
*** p<.01, ** p<.05, * p<.10

Table 3.4: Effect of state Medicaid obligations on budget items outside of Medicaid

	Model					
	(1)	(2)	(3)	(4)	(5)	(6)
Justice Spending (2SLS)	0.089	0.028	0.055**	0.026	0.003	0.000
	(0.067)	(0.074)	(0.023)	(0.028)	(0.027)	(0.024)
Transportation (2SLS)	-0.201	-0.479	-0.216	-0.110	-0.139	-0.184
	(0.322)	(0.362)	(0.133)	(0.160)	(0.135)	(0.140)
K-12 Ed (2SLS)	-0.408	-0.135	0.034	0.099	0.002	0.031
	(0.333)	(0.415)	(0.037)	(0.067)	(0.130)	(0.128)
N	1666	1666	1666	1666	1666	1666
Years	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010
Lagged Unemployment	x	x	x	x	x	x
Demographic Controls	x	x	x	x	x	x
Policy Period FE		x		x		x
State FE			x	x	x	x
State Trends					x	x

Table shows the 2SLS effect of an additional dollar of Medicaid spending on other types of state spending, using SSI enrollment as the instrumental variable. ”Justice” includes prisons and law enforcement. ”Transportation” includes roads and other types of infrastructure such as airports or trains. Point estimates for K-12 and Justice spending are small and statistically insignificant. Point estimates for Transportation are economically significant but imprecise, which we interpret as institutional heterogeneity in the state legislature’s control of the transportation fund.
*** p<.01, ** p<.05, * p<.10

Table 3.5: Bartik Instruments

	Model					
	(1)	(2)	(3)	(4)	(5)	(6)
Higher Ed Spending (OLS)	-0.081	-0.116*	-0.159***	-0.102*	-0.089***	-0.079***
	(0.066)	(0.058)	(0.042)	(0.052)	(0.029)	(0.029)
Higher Ed Spending (2SLS)	-0.099	-0.183	-0.227***	-0.153***	-0.185***	-0.198***
	(0.127)	(0.133)	(0.042)	(0.052)	(0.043)	(0.039)
N	1666	1666	1666	1666	1666	1666
Years	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010
Lagged Unemployment	x	x	x	x	x	x
Demographic Controls	x	x	x	x	x	x
Policy Period FE		x		x		x
State FE			x	x	x	x
State Trends					x	x

Here, we redefine the instrument of the main analysis to compensate for exogenous trends in the costs of providing health care. Each state's increase in SSI enrollment is multiplied by the nationwide increase in average disability patient cost outside of that state. The result is a more powerful instrument which more accurately predicts total Medicaid costs. The main results here are slightly lower than those in the main specification, but much more precise.
*** p<.01, ** p<.05, * p<.10

	Model					
	(1)	(2)	(3)	(4)	(5)	(6)
Higher Ed Spending (OLS)	-0.218***	-0.216***	-0.220***	-0.222***	-0.225***	-0.228***
	(0.052)	(0.052)	(0.051)	(0.052)	(0.052)	(0.052)
Higher Ed Spending (2SLS)	-0.370***	-0.365***	-0.370***	-0.377***	-0.367***	-0.343***
	(0.114)	(0.115)	(0.113)	(0.114)	(0.112)	(0.108)
N	1666	1666	1666	1666	1666	1666
Years	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010

Lagged Unemployment	x	x	x	x	x	x
Demographic Controls	x	x	x	x	x	x
Policy Period FE		x		x		x
State FE			x	x	x	x
State Trends					x	x

Table 3.6: Delta model

Here we repeat the main analysis, but use a model specifying changes in enrollment and spending, rather than levels. This model emphasizes the short term changes, rather than the longer term changes implied by the fixed effects model. This model also has the additional benefit that including state effects and state trends in this model allows for a more flexible, nonlinear state trend assumption. We find that the point estimates are roughly identical to those in our main specification in Table 3.3.
*** p<.01, ** p<.05, * p<.10

Table 3.7: Sensitivity to year effects

	Model					
	(1)	(2)	(3)	(4)	(5)	(6)
Higher Ed Spending (OLS)	-0.081	-0.110*	-0.159***	-0.017	-0.089***	0.018
	(0.066)	(0.060)	(0.042)	(0.070)	(0.029)	(0.037)
Higher Ed Spending (2SLS)	0.664**	1.342	-0.212***	0.437	-0.327**	0.011
	(0.310)	(0.923)	(0.052)	(0.325)	(0.155)	(0.354)
N	1666	1666	1666	1666	1666	1666
Years	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010
Lagged Unemployment	x	x	x	x	x	x
Demographic Controls	x	x	x	x	x	x
Year FE		x		x		x
State FE			x	x	x	x
State Trends					x	x

Table shows the sensitivity of the main analysis to the inclusion of year fixed effects. The preferred specification (6) with year fixed effects shows a point estimate of 0.011. Due to the low point estimate and large standard error associated with this model, it is difficult to draw much meaningful inference from this specification. The model here can neither reject the point estimates in Table 3.3 nor a point estimate of zero. *** p<.01, ** p<.05, * p<.10

Table 3.8: Sensitivity to level policy changes

	Model					
	(1)	(2)	(3)	(4)	(5)	(6)
Higher Ed Spending (OLS)	-0.081	-0.116*	-0.159***	-0.102*	-0.089***	-0.079***
	(0.066)	(0.058)	(0.042)	(0.052)	(0.029)	(0.029)
Higher Ed Spending (2SLS)	0.664**	0.954	-0.212***	-0.040	-0.327**	-0.373**
	(0.310)	(0.599)	(0.052)	(0.109)	(0.155)	(0.177)
N	1666	1666	1666	1666	1666	1666
Years	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010	1977-2010
Lagged Unemployment	x	x	x	x	x	x
Demographic Controls	x	x	x	x	x	x
Policy Period FE		x		x		x
State FE			x	x	x	x
State Trends					x	x

Table 3.8 imposes more structure on the year fixed effects in table 3.7, and instead includes period fixed effects for major federal policy changes. These changes include a 1985 adjustment to disability determination and a 1996 effort to retrenchment effort which reduced enrollment. Including these measures does not significantly change the results from Table3.3.
*** p<.01, ** p<.05, * p<.10

Table 3.9: Heterogeneity by state subgroup

	Tercile		
	(1)	(2)	(3)
Subgroup by college share	-0.073	-0.211***	-0.420***
	(0.054)	(0.042)	(0.113)
Subgroup by elderly share	-0.310***	-0.189**	-0.147***
	(0.083)	(0.081)	(0.040)
N	1666	1666	1666
Years	1977-2010	1977-2010	1977-2010
Lagged Unemployment	x	x	x
Demographic Controls	x	x	x
Policy Period FE	x	x	x
State FE	x	x	x
State Trends	x	x	x

In this table we present regression results for states divided into terciles based on state characteristics. The first grouping is based on the share of state population having attended some college in 1980. The second is based on the share of state population over age 65. We believe higher education spending in these subgroups may exhibit different sensitivities based on state political economy. We find that states which have a higher share of the population having attended college show larger declines in spending. We also find states which had a lower share of the population over age 65 saw larger declines in higher education spending. Both of these results may be attributable to mean reversion, or higher sensitivity from places which were already overspending. In each case we present the preferred regression results, which include all control variables and use the SSI Bartik instrument.
*** $p<.01$, ** $p<.05$, * $p<.10$

APPENDIX A

SSI Enrollment by Diagnosis and by Age Group

Figure B1: Trend in Non-Elderly Adult SSI Recipients by Diagnosis

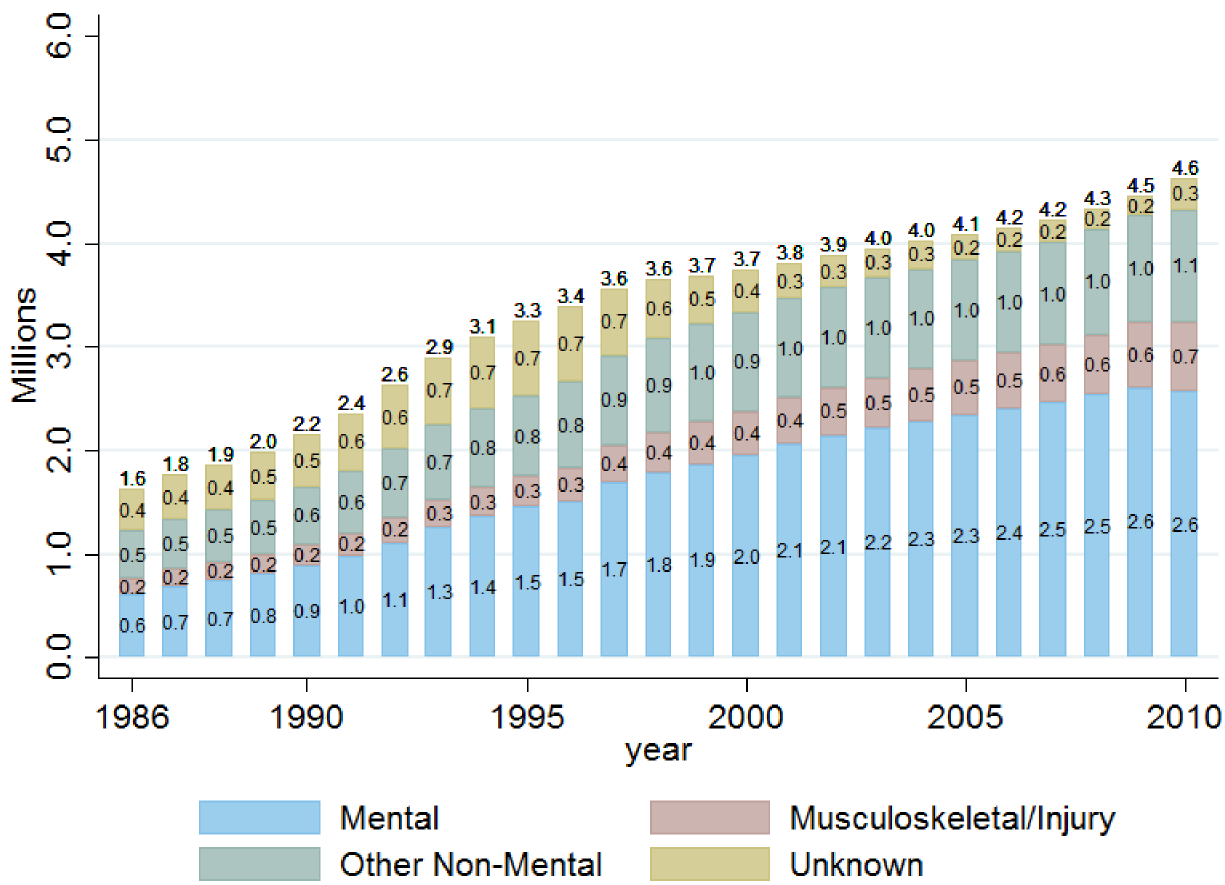

Source: SSA Annual Statistical Supplement (1986-2010). Non-elderly adults represent ages 18 to 64. The lowest group in the figure is Mental (0-64), and the second lowest group in the figure is Musculoskeletal/Injury (0-64). Similar to Figure 3.3, however this figure includes raw counts of enrollments, rather than population shares.

Figure B2: Trend in Children SSI Recipients by Diagnosis

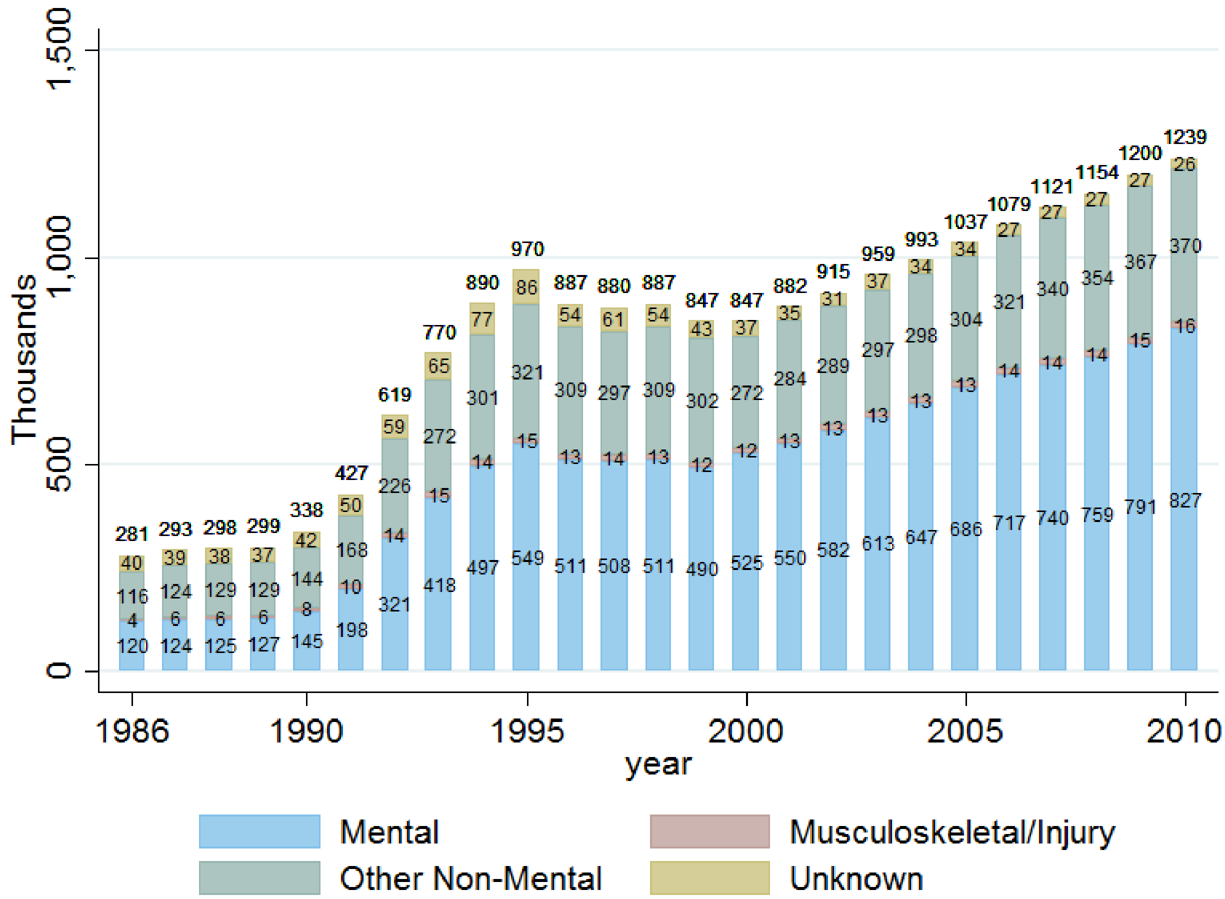

Source: SSA Annual Statistical Supplement (1986-2010). Children represent ages 0 to 17. The lowest group in the figure is Mental (0-64), and the second lowest group in the figure is Musculoskeletal/Injury (0-64). Similar to Figure 3.3. Here we see that mental health disabilities are the driving force for child SSI eligibility, as well as for adults.

APPENDIX B

Decomposition of Year-over-Year Change

Using year-over-year changes instead, Figure C1 allows us to see each step of the linkage between SSI enrollment and total Medicaid spending. Year-over-year changes in SSI enrollment are strongly correlated with Medicaid disability enrollment. In turn, year-over-year changes in Medicaid spending for disabled enrollees is also strongly correlated with total Medicaid spending.

Figure C1: Delta plots

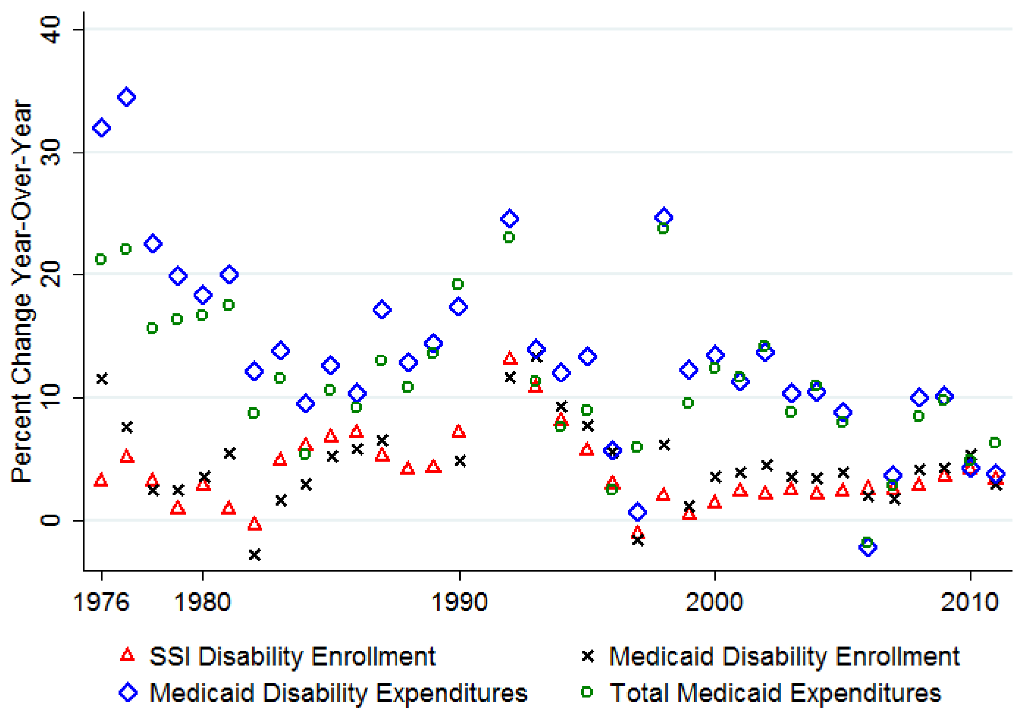

This figure shows the percent change in SSI disability enrollment, Medicaid disability enrollment, Medicaid disability payments, and Total Medicaid payments for each year, from 1976 to 2011. Overall this shows that Medicaid disability closely tracks SSI disability, and Total Medicaid costs closely track disability Mediciad costs. The relationship between Medicaid disability enrollment and Medicaid disability costs is visible but less powerful because of strong secular growth in average patient cost.

APPENDIX C

Distribution of state-level FMAPs

Figure D1: FMAP distribution

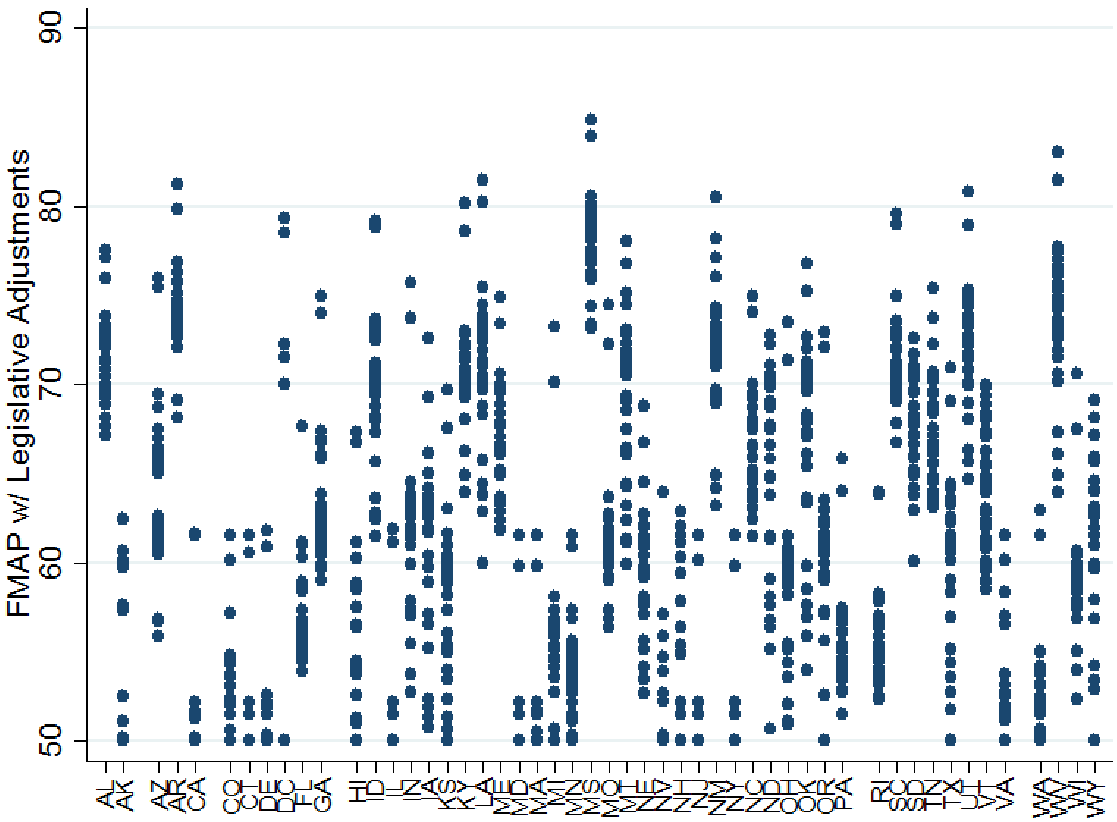

This figure shows the overall distribution of the Federal Medical Assistance Percentage within each state over 25 years. In general most states stay within a range of about 10 points. This relationship is particularly strong for states close to the lower bound of the federal matching rate.

BIBLIOGRAPHY

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2012. "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program." *Journal of the American statistical Association.* 27

Abdulkadiroğlu, Atila, Joshua D Angrist, Susan M Dynarski, Thomas J Kane, and Parag A Pathak. 2011. "Accountability and flexibility in public schools: Evidence from Boston's charters and pilots." *The Quarterly Journal of Economics*, 126(2): 699–748. 3

Autor, David H., and Mark G. Duggan. 2003. "The Rise in the Disability Rolls and the Decline in Unemployment*." *The Quarterly Journal of Economics*, 118(1): 157. 97

Baicker, Katherine. 2005. "Extensive or intensive generosity? The price and income effects of federal grants." *Review of Economics and Statistics*, 87(2): 371–384. 95

Bartik, Timothy J. 1991. "Boon or boondoggle? The debate over state and local economic development policies." 113

Bell, Julie D. 2008. *The Nuts and Bolts of the Higher Education Legislative Appropriations Process.* National Conference of State Legislatures, Denver, CO and Western Interstate Commission for Higher Education, Boulder, CO. 100

Biasi, Barbara. 2017. "Unions, salaries, and the market for teachers: Evidence from Wisconsin." *Working Paper.* 5

Bound, John, and Richard V. Burkhauser. 1999. "Economic analysis of transfer programs targeted on people with disabilities." In *Handbook of Labor Economics.* Vol. 3 of *Handbook of Labor Economics*, , ed. O. Ashenfelter and D. Card, Chapter 51, 3417–3528. Elsevier. 109

Bound, John, Breno Braga, Gaurav Khanna, and Sarah Turner. 2016. "A Passage to America: University Funding and International Students." National Bureau of Economic Research Working Paper 22981. 100

Card, David. 1996. "The effect of unions on the structure of wages: A longitudinal analysis." *Econometrica: Journal of the Econometric Society*, 957–979. 4

Card, David. 2001. "The effect of unions on wage inequality in the US labor market." *Industrial & Labor Relations Review*, 54(2): 296–315. 23

Card, David, and Abigail Payne. 2002. "School finance reform, the distribution of school spending, and the distribution of student test scores." *Journal of Public Economics*, 83(1): 49–82. 53

Cascio, Elizabeth U, Nora Gordon, and Sarah Reber. 2013. "Local responses to federal grants: evidence from the introduction of title I in the South." *American Economic Journal: Economic Policy*, 5(3): 126–159. 52, 62, 72

Chodorow-Reich, Gabriel, Laura Feiveson, Zachary Liscow, and William Gui Woolston. 2012. "Does state fiscal relief during recessions increase employment? Evidence from the American Recovery and Reinvestment Act." *American Economic Journal: Economic Policy*, 4(3): 118–145. 53, 70

Cohen, Jennifer S. 2011. "The State Fiscal Stabilization Fund and Higher Education Spending: Part 2 of 4." *New America Foundation.* 65

Cohodes, Sarah, Daniel Grossman, Samuel Kleiner, and Michael Lovenheim. 2016. "The Effect of Child Health Insurance Access on Schooling: Evidence from Public Insurance Expansions." *Journal of Human Resources*, 51. 119

Courant, Paul N, Edward M Gramlich, and Daniel L Rubinfeld. 1979. "Public employee market power and the level of government spending." *The American Economic Review*, 69(5): 806–817. 3

Cullen, Julie Berry. 2003. "The impact of fiscal incentives on student disability rates." *Journal of Public Economics*, 87(7): 1557–1589. 55

Currie, Janet, and Jonathan Gruber. 1996. "Saving Babies: The Efficacy and Cost of Recent Changes in the Medicaid Eligibility of Pregnant Women." *Journal of Political Economy*, 104(6): 1263–96. 95

Cutler, David M., and Jonathan Gruber. 1996. "The Effect of Medicaid Expansions on Public Insurance, Private Insurance, and Redistribution." *The American Economic Review*, 86(2): 378–383. 95

Daly, Mary, and Richard Burkhauser. 2003. "The Supplemental Security Income Program." In *Means-Tested Transfer Programs in the United States.* 79–140. National Bureau of Economic Research, Inc. 109

Delaney, Jennifer. 2011. "Earmarks and state appropriations for higher education." *Journal of Education Finance*, 37(1): 3–23. 68

DiNardo, John, and David S Lee. 2004. "Economic impacts of unionization on private sector employers: 1984-2001." National Bureau of Economic Research. 2, 4

DiNardo, John, Nicole M Fortin, and Thomas Lemieux. 1995. "Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach." National bureau of economic research. 23

Feyrer, James, and Bruce Sacerdote. 2011. "Did the stimulus stimulate? real time estimates of the effects of the american recovery and reinvestment act." National Bureau of Economic Research. 53, 70

Frandsen, Brigham R. 2012. "Why unions still matter: The effects of unionization on the distribution of employee earnings." 4

Frandsen, Brigham R. 2016. "The Effects of Collective Bargaining Rights on Public Employee Compensation Evidence from Teachers, Firefighters, and Police." *ILR Review*, 69(1): 84–112. 5, 7

Freeman, Richard, and James Medoff. 1984. "What Do Unions Do." 2

Freeman, Richard B. 1988. "Contraction and expansion: the divergence of private sector and public sector unionism in the United States." *The Journal of Economic Perspectives*, 2(2): 63–88. 2

Freeman, Richard B, and Morris M Kleiner. 1990. "The impact of new unionization on wages and working conditions." *Journal of Labor Economics*, S8–S25. 1

Freeman, Richard B, and Robert Valletta. 1988. "The effects of public sector labor laws on labor market institutions and outcomes." In *When public sector workers unionize.* 81–106. University of Chicago Press. 4, 5

Freeman, Richard Barry, and Eunice S Han. 2013. "Public sector unionism without collective bargaining." 4

French, Eric, and Jae Song. 2014. "The Effect of Disability Insurance Receipt on Labor Supply." *American Economic Journal: Economic Policy*, 6(2): 291–337. 109

Gordon, Nora. 2004. "Do federal grants boost school spending? Evidence from Title I." *Journal of Public Economics*, 88(9): 1771–1792. 52, 60, 62, 72

Gramlich, Edward M, Harvey Galper, Stephen Goldfeld, and Martin McGuire. 1973. "State and local fiscal behavior and federal grant policy." *Brookings Papers on Economic Activity*, 1973(1): 15–65. 52

Hines, James R, and Richard H Thaler. 1995. "Anomalies: The flypaper effect." *The Journal of Economic Perspectives*, 9(4): 217–226. 52

Hirsch, Barry T, and David Macpherson. 2016. "Union membership and coverage database from the CPS." *Available from Unionstats. com).[Accessed 22 September 2016].* 2, 7

Hoxby, Caroline Minter. 1996. "How teachers' unions affect education production." *The Quarterly Journal of Economics*, 671–718. 4, 5, 7, 14

Hyman, Joshua. Forthcoming. "Does money matter in the long run? Effects of school spending on educational attainment." *American Economic Journal: Economic Polic.* 53

Kane, Thomas J, and Peter R Orszag. 2003. *Higher education spending: The role of Medicaid and the business cycle.* Brookings Institution Washington, DC. 93, 100, 108

Klemm, John D. 2000. "Medicaid Spending: A Brief History." *Health Care Financing Review*, 22(1): 105–112. 95

Kleven, Henrik Jacobsen, and Wojciech Kopczuk. 2011. "Transfer Program Complexity and the Take-Up of Social Benefits." *American Economic Journal: Economic Policy*, 3(1): 54–90. 119

Lee, David, and Alexandre Mas. 2009. "Long-run impacts of unions on firms: New evidence from financial markets, 1961-1999." National Bureau of Economic Research. 4

Lewis, H Gregg. 1990. "Union/nonunion wage gaps in the public sector." *Journal of Labor Economics*, 8: 260–328. 2, 5

Lindy, Benjamin A. 2011. "The impact of teacher collective bargaining laws on student achievement: Evidence from a New Mexico natural experiment." *The Yale Law Journal*, 1130–1191. 7

Lovenheim, Michael F. 2009. "The effect of teachers' unions on education production: Evidence from union election certifications in three midwestern states." *Journal of Labor Economics*, 27(4): 525–587. 2, 4, 5, 14

MACPAC. 2016. "Report to Congress on Medicaid and CHIP, Chapter 1: Trends in Medicaid Spending." Medicaid and CHIP Payment and Access Commission. 96

Maestas, Nicole, Kathleen J. Mullen, and Alexander Strand. 2013. "Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt." *American Economic Review*, 103(5): 1797–1829. 109

Murphy, Marjorie. 1990. *Blackboard Unions: The AFT and the NEA, 1900-1980.* ERIC. 7

National Center for Education Statistics. 2016. "Allocating Grants for Title I." *US Department of Education, Institute for Education Science.* 55

NCSL. 2010. "NCSL Fiscal Brief: State Balanced Budget Provisions." National Conference of State Legislatures. 100

Poterba, James M. 1996. "Demographic structure and the political economy of public education." National bureau of economic research. 118

Rupp, Kalman, and Charles Scott. 1998. "Determinants of Duration on the Disability Rolls and Program Trends." In *Growth in Disability Benefits: Explanations and Policy Implications.* , ed. Kalman Rupp and David Stapleton, 31–92. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. 97

Seegert, Nathan. 2012. "Optimal taxation with volatility: A theoretical and empirical decomposition." *Working Paper.* 61, 83, 101

Sojourner, Aaron J, Brigham R Frandsen, Robert J Town, David C Grabowski, and Min M Chen. 2015. "Impacts of Unionization on Quality and Productivity Regression Discontinuity Evidence from Nursing Homes." *ILR Review*, 68(4): 771–806. 4

Stapleton, David, Kevin Coleman, Kimberly Dietrich, and Gina Livermore. 1998. "Empirical Analyses of DI and SSI Application and Award Growth." In *Growth in Disability Benefits: Explanations and Policy Implications.* , ed. Kalman Rupp and David Stapleton, 31–92. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. 97

US Department of Education. 2005. "10 Facts About K-12 Education Funding." 55

West, Martin. 2013. "Do Math and Science Teachers Earn More Outside of Education." *The Brown Center Chalkboard. April*, 17. 23

Wherry, Laura R., Sarah Miller, Robert Kaestner, and Bruce D. Meyer. 2015. "Childhood Medicaid Coverage and Later Life Health Care Utilization." National Bureau of Economic Research Working Paper 20929. 119

Wilson, Daniel J. 2012. "Fiscal spending jobs multipliers: Evidence from the 2009 American Recovery and Reinvestment Act." *American Economic Journal: Economic Policy*, 4(3): 251–282. 53, 70

Wilson, Steven F. 2009. "Success at scale in charter schooling." *Education Outlook*, 3: 1–7. 3