THE IMPACT OF ACCESS TO ELECTRICITY ON EDUCATION AND OTHER ESSAYS IN SPATIAL ECONOMICS

BY

TIMOTHY L. SQUIRES

B.A., QUINNIPIAC UNIVERSITY, 2009

M.A., BROWN UNIVERSITY, 2010

A DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY IN THE DEPARTMENT OF ECONOMICS AT BROWN UNIVERSITY

PROVIDENCE, RHODE ISLAND
MAY 2015

This dissertation by Timothy L. Squires is accepted in its present form

by the Department of Economics as satisfying the

dissertation requirements for the degree of Doctor of Philosophy.

Date ______________ ______________________________

Doctor David Weil, Adviser

Recommended to the Graduate Council

Date ______________ ______________________________

Doctor Andrew Foster, Reader

Date ______________ ______________________________

Doctor Adam Storeygard, Reader

Approved by the Graduate Council

Date ______________ ______________________________

Pete M. Weber, Dean of the Graduate School

VITA

Tim Squires was born May 30th, 1987 in Newton, Massachusetts and attended Newton North High School. He graduated with honors from Quinnipiac University in 2009 with a Bachelor of Science in Economics and a minor in math, finance and computer science. He enrolled in Brown University's Masters program in 2009 and obtained his masters in 2010. In 2010 he was enrolled in the doctoral program at Brown University. At Brown University he was a fellow in the Population Studies and Training Center and the Spatial Structures in the Social Sciences. He was a Graduate Merit Dissertation Fellowship and in 2013 was named fellow of the year in the Spatial Structures in the Social Sciences.

ACKNOWLEDGEMENTS

I first would like to thank Professors David Weil and Andrew Foster for their help and support over the years. David's support and patience was vital to not only my development but my success in graduate school. He helped me through the process, guiding me along the way and when needed pushing me along. Andrew was the first person to take a chance on me in graduate school and created the master's program of which I was the the first enrollee. Without that masters program and his support I would not have been at Brown or been as successful. Lastly, I appreciate the hard work of the economics office staff, especially Angelica Vargas. Her help in guiding me through many of the processes was irreplaceable and her unending support made my entire time here much more enjoyable.

I want to thank the Population Studies and Training Center and the Spatial Structures in the Social Sciences for their financial support and for giving me a place to interact with students of all different backgrounds.

I want to thank all of my teachers at Cabot Elementary School, Day Middle School and Newton North High School for their guidance. I would especially like to thank Ms. McFarland, Mr. Demarest, Mr. Ford, and Mr. Brown for their belief in me. I would also like to thank my Professors at Quinnipiac University especially Donn Johnson, Chris Ball, Eric Brunner, Erik Johnson and Kathy Cooke who pushed me to be more than I thought I could.

My friends and family also provided continuous support to me over the years. I want to especially thank my parents and sister for putting up with all the highs and lows of graduate school and always being there when I needed to talk. I would like to thank my friends Sam Redlinger, Paul Jones and Jen Ellsworth who have been my rock through all of this. To create an exhaustive list would be impossible but I would like to thank those people who have supported me during my time in graduate school including Mariá José Boccardi Chalela, Morgan Hardy, David Glancy, Daniela Scidá, Michelle Marcus, Angélica Meinhofer Santiago, Desislava Byanova, and Phil DaCosta.

Contents

List of Tables

List of Figures

Abstract of "The Impact of Access to Electricity on Education and Other Essays in Spatial Economics" by Timothy L. Squires, Ph.D., Brown University, May 2015

The first chapter examines how district size affects the bargaining power of teachers unions and the allocation of school resources. Our identification strategy exploits the fact that 33 states mandate collective bargaining while 5 states prohibit it. In states that mandate collective bargaining, we find that beginning salaries and the premium paid to experienced teachers increase with district size while the teacher-pupil ratio declines with district size. In contrast, in states that prohibit collective bargaining we find a negative relationship between district size and the premium paid to experienced teachers. The second chapter examines malaria's impact on economics activity. Previous research estimating the impact of malaria on economic activity has been inconclusive and has struggled with certain estimation problems. Specifically, there are concerns that previous results have been driven by country specific characteristics, poor measures of malaria, or the presence of other diseases. This paper addresses each of these concerns. I introduce a novel and more accurate measure of malaria at the grid square level to reduce the measurement error problem. By using this geographically finer measure of malaria, combined with nighttime lights and population data as a proxy for economic activity, I am able to control for country fixed effects. The results suggest that reducing malaria prevalence by 20% would increase GDP by 20%. The third chapter estimates the effect of access to electricity on school attendance and educational attainment. I take advantage of individual level data matched with community level electrification dates, allowing me to study the effect of access to electricity on completed education as well as the age-specific hazard of dropping out of school. Contrary to expectations, I find that access to electricity *reduces* educational attainment. The reduction in education was accompanied by an increase in childhood employment, suggesting that improved labor market opportunities, due to electricity access, led to the increased dropout rates.

Chapter 1

The Bargaining Power of Teachers' Unions and the Allocation of School Resources

1.1 Introduction

"One of the wonderful things about being a teacher is that you get to help elect your own bosses. Illinois AFL-CIO president, 2010."

How does the bargaining power of teachers unions affect teacher compensation and the allocation of school resources? The answer to that question remains controversial. One view is that teacher unions use their bargaining power to promote the self-interests of their members, namely teachers. Those that espouse this view argue that greater bargaining power leads to higher salaries, better working conditions (e.g. smaller class sizes) and more generous returns to seniority. The alternative view is that bargaining power allows teacher unions to distribute resources more efficiently and increase the productivity of teachers. Those that espouse this view argue that even if teacher unions impact wages and other working conditions they do so simply to make teachers more productive, attract higher quality teachers and increase student performance. Despite the controversial role teacher unions play in the provision of education, surprisingly little empirical evidence is available on the determinants of teacher union bargaining power or how union bargaining power affects the allocation of school resources.

What evidence is available comes primarily from studies that compare the allocation of resources in unionized and non-unionized school districts. For example, a number of studies compare wage differentials in unionized and non-unionized districts or teacher salaries before and after they were represented by a union. The results of those studies are mixed. Baugh and Stone (1982) and Hoxby

(1996) find that unionization increases teacher pay by approximately 5 to 12 percent. In contrast, Kleiner and Petree (1988) and Lovenheim (2009) find that unionization has little impact on teacher pay. Studies that examine other outcomes of unionization such as class size and student performance also tend to find mixed results.[1]

The purpose of this paper is to provide new evidence on the determinants of teacher union bargaining power and how bargaining power affects the allocation of school resources. While most of the previous literature has focused on comparing educational outcomes in districts with and without unions, we focus on the idea that even among unionized districts union power may be quite heterogeneous. For example, Brown and Medoff (1988) find that that the wages paid by unionized local government employers are substantially higher in larger districts than in smaller districts, a finding that suggests unions in larger districts may have greater bargaining power. Similarly, using data from California, a state where nearly all districts engage in collective bargaining, Rose and Sonstelie (2010) find that beginning teacher salaries and the salary premium paid to experienced teachers both increase with district size. Rose and Sonstelie (2010) and Hoyt (1999) develop public choice models that provide explanations for why union bargaining power may increase with district size. Following this literature we focus specifically on how district size affects union bargaining power and the subsequent allocation of school resources.

To identify the causal impact of district size on union bargaining power and outcomes of the collective bargaining process, we estimate models that include state and local labor market fixed effects and exploit the fact that in 33 states collective bargaining is mandatory while in five states collective bargaining is explicitly prohibited.[2] We are therefore able to compare the relationship between district size and outcomes of the collective bargaining process in states that mandate and states that explicitly prohibit collective bargaining while controlling for state and labor market unobservables. Our analysis is based on data from the restricted use version of the 2007-2008 Schools and Staffing Survey (SASS), a large and nationally representative sample of school districts. The survey contains information on teacher salaries, including salaries at various points in the teacher salary schedule and the number of teachers employed in a district. We use those variables to construct our dependent variables of interest which include beginning teacher salaries, the salary premium paid to

[1] For detailed summaries of this line of research see Stone (2002) and Eberts (2007) and Lovenheim (2009).

[2] There are actually 35 states with mandatory collective bargaining laws. We omit two of those states from our analysis, namely Hawaii and Alaska. The five states that prohibit collective bargaining are: Georgia, North Carolina, South Carolina, Texas and Virginia. There are also 11 states that neither mandate nor prohibit collective bargaining. Because the decision to bargain in those 11 states is likely endogenous, we exclude them from our analysis.

experienced teachers and the teacher-pupil ratio, all of which can potentially be influenced by the collective bargaining process. The restricted use version of the SASS also contains school district identification codes that allow us to merge the SASS survey data with a rich set of school district demographic controls from the 2005-2009 American Community Survey and the National Center for Educational Statistics. We find that in states that mandate collective bargaining, base salaries and the salary premium paid to experienced teachers increase with district size while the teacher-pupil ratio declines with district size, findings consistent with the results of Rose and Sonstelie (2010). In contrast, in states that prohibit collective bargaining we find a negative relationship between district size and the salary premium paid to experienced teachers. Furthermore, we find that the positive relationship between district size and base salaries is stronger while the negative relationship between district size and the teacher-pupil ratio is weaker in states that prohibit collective bargaining.

A series of robustness checks provides evidence that these effects are highly robust. For example, our results persist across specifications with and without demographic controls, reducing the concern that our estimates are being driven by unobservables. They are also robust to dropping large districts with high concentrations of potentially difficult to teach students, suggesting our results are not being driven by larger districts paying higher salaries to compensate for a more difficult working environment. The estimated relationship between district size and outcomes of the collective bargaining process also persists, and is found to be monotonic, when we replace our continuous measure of district size with a set of indicator variables that correspond to district size quartiles.

In terms of magnitude, our results suggest that in states that mandate collective bargaining, moving from a district one standard deviation below the mean to one standard deviation above the mean in district size would increase the base salary of teachers by 4.5% and the salary premium paid to experienced teachers by 2.4% for teachers with a BA and by 3.6% for teachers with an MA. These latter estimates imply that a teacher with 10 years of teaching experience working in a district one standard deviation above the mean of district size would have a salary that was between 7% to 8.3% (depending on degree) higher than a teacher with the same years of teaching experience working in a district one standard deviation below the mean of district size. To put these estimates into perspective, Hoxby (1996) estimates that unionization increases average teacher salaries by approximately 5% while West and Mykerezi (2010) find that unionization increase beginning teacher salaries by approximately 4%. Thus, our estimates of the effect of a two standard deviation increase in district size on salaries are comparable to previous estimates of the effect of unionization on salaries.

1.2 Conceptual Framework

The effect teachers unions have on the allocation of school resources and school productivity depends on two factors: (1) the objective function of teachers unions and (2) the ability of unions to influence the allocation of resources through their collective bargaining power. There are two alternative views concerning the objective function of teachers unions. The first assumes teachers unions are primarily interested in promoting the self-interests of their members. Under this rent-seeking view, the primary objective of a union is to bargain for higher salaries, better working conditions and other items that benefit mainly teachers. As such, teacher unions may lead to an inefficient allocation of resources that favors salaries over other inputs to school quality. In support of this view a number of studies have found that teacher salaries are higher and student-teacher ratios lower in unionized versus non-unionized districts. In the most comprehensive of these studies, Hoxby (1996) finds that unionization increases teacher pay by approximately 5%, decreases the student-teacher ratio by 1.7 students per teacher and leads to an increase in high school dropout rates.

The second view assumes that teacher unions have the same objective function as parents and local voters, namely to maximize school quality. Under this view, unions may have little impact on the allocation of school resources and school productivity. In support of this view, Lovenheim (2009) finds that unions have no effect on teacher pay, class size or student achievement as measured by high school dropout rates. However, even if the objective function of teacher unions is to maximize school quality, they may still influence the allocation of school resources. As Hoby (1996) notes, given their direct contact with students, teachers may have better information than administrators and local voters on how resources could be used most efficiently. As a result, teachers unions may bargain on behalf of higher teacher salaries, larger returns to experience and smaller class sizes not because they benefit mainly teachers but because they help attract higher quality teachers, retain the most productive teachers and increase school productivity.

The second factor that determines the degree to which unions affect the allocation of school resources is the ability of unions to influence outcomes of the collective bargaining process. As noted by Moe (2006), Rose and Sonstelie (2010) and Strunk (2010), the primary source of union bargaining power comes from a unions ability to influence local school board elections and hence the composition of the school board with whom it must bargain. Are some unions likely to be more successful at influencing school board elections than others? Rose and Sonstelie (2010) develop a public choice model that addresses that question. In their model, teachers unions and homeowners compete for bargaining power during local school board elections. In this political competition, both homeowners and the union attempt to elect school board candidates that are sympathetic to their interests. While both the union and homeowners can contribute time and money to the candidate of their choice, homeowners are at a particular disadvantage. Political contributions and other efforts taken on behalf of a candidate are a public good and therefore subject to the free rider problem. Teachers unions on the other hand can overcome the free rider problem by taxing members in the form of union dues. Because the free riding problem increases with group size, homeowners in larger districts are at more of a disadvantage than homeowners in smaller districts. As a result, the bargaining power of teachers unions tends to increase with district size.

Hoyt (1999) develops an alternative model that also predicts unions should be more powerful in larger communities. In his model, bargaining power increases with district size because the cost of inefficiencies and lower quality public services is not fully capitalized into housing values in larger districts. Specifically, Hoyt (1999) demonstrates that while government inefficiency and tax increases are fully capitalized into property values in small communities they are not fully capitalized in large communities.[3] Consequently, since monitoring government activity requires time and effort, voters in larger communities have less incentive to monitor union activities and the potential misallocation of resources than voters in smaller communities.

Regardless of the mechanism, if the relative power of teacher unions in school board elections increases with district size, teacher unions should be relatively more active and influential in larger districts. Moe (2005) presents evidence consistent with that hypothesis. Based on survey responses from 526 school board candidates in California, he finds that union activity and influence tends to increase with district size. For example, when asked whether the union supported specific candidates, only 24% of respondents in districts with less than 1,000 students responded yes while 92% of respondents in districts with more than 25,000 students responded yes. Similarly, when asked if the union contributed money to school board candidate campaigns or mobilized its members on behalf

[3]Brasington (2001) provides evidence consistent with the prediction that government service inefficiencies and lower quality public services are not fully capitalized into housing values in larger districts. Using housing transaction data from the six largest metropolitan areas in Ohio, he finds that the rate of capitalization of school quality into housing values is significantly weaker in large school districts than in smaller school districts. A result that is robust to a number of specification checks.

of candidates, less than 30% of respondents in districts with less than 1,000 students responded yes while nearly 95% of respondents in districts with more than 25,000 students responded yes. Finally, when directly asked whether the teacher union played an important role in elections, 43% of respondents in districts with less than 1,000 students responded yes while 82% of respondents in districts with more than 25,000 students responded yes.[4]

While the survey results of Moe (2005) suggest that teacher unions tend to be more active and influential in larger districts, we are particularly interested in whether union influence in school board elections translates into more favorable collective bargaining outcomes. In particular, if union power increases with district size, the contracts negotiated by unions in larger districts should reflect that power in the form of greater weight placed on collective bargaining outcomes that are of most interest to union members. In the empirical work that follows, we focus on three outcomes of the collective bargaining process: beginning teacher salaries, the premium paid to experienced teachers and the teacher-pupil ratio. Of these three outcomes, we believe the premium paid to experienced teachers is perhaps the best measure of union bargaining power and rent-seeking behavior on the part of unions for several reasons. First, most studies have found that beyond the first three or four years, experience contributes very little to teacher performance.[5] Thus it seems unlikely that large returns to experience can be justified on the grounds of productivity. Second, as noted by Lankford and Wyckoff (1997) and Bifulco (2010), increases in experienced teacher salaries are a less effective method of attracting and retaining high-quality teachers then increases in beginning teacher salaries. For example, Murnane et al. (1991) find that salary increases have a much larger effect on the career decisions of new teachers than the career decisions of experienced teachers.[6] Third, based on a sample of unionized school districts in Pennsylvania, Babcock and Engberg (1999) find that the premium paid to experienced teachers is higher in districts with more senior teachers, and that this relationship is most pronounced in districts where support for union activity is high. That evidence is consistent with the notion that high returns to experience represent a rent to senior teachers. Finally, in their analysis of the returns to teacher seniority, Ballou and Podgursky (2002) find that

[4]Hess and Leal (2005) provide further evidence that teacher unions tend to be more active and influential in larger districts. Based on a nationally representative survey of 827 school board members, they find that the activity level of teacher unions in school board elections tends to increase with district size and that in larger districts teacher unions are more active in school board elections than any other group including business and parent organizations.

[5]See for example, Hanushek, Kain, OBrien and Rivkin (2005), Clotfelter, Ladd, and Vigdor (2006), Gordon, Kane and Staiger (2008) and Rockoff (2004). An exception is Wiswall (2011) who finds significant returns to later teaching experience on 5th grade math scores in North Carolina but not on 5th grade reading scores.

[6]Specifically, Murnane et al. (1991) find that while higher salaries increase the probability that a new teachers will choose to remain in the classroom, the relationship between salary and the probability of leaving a teaching job declines steadily with years of teaching experience. After six to eight years of teaching, salary has almost no effect on the decision to exit the teaching profession.

the observed returns to teacher experience are inconsistent with three prominent theories for an upward sloping wage-experience profile, namely human capital accumulation, monitoring costs, and turnover costs. Having found little evidence consistent with standard rationales for high returns to experience, the authors conclude that the structure of teacher pay is more consistent with rent seeking than efficient contracting.

Our focus on the salary premium paid to experienced teachers as the primary measure of union bargaining power leads to a broader issue. As discussed previously, whether teachers unions are purely rent seeking or seek simply to maximize school quality, the objective function of a teachers union is likely to be an increasing function of beginning teacher salaries, the premium paid to experienced teachers and the teacher-pupil ratio since all of these factors may enhance both school quality and the well-being of teachers. Thus, none of our outcome variables represent absolute indicators of union bargaining power or rent seeking behavior on the part of unions. We have argued that perhaps the best measure of union bargaining power is the premium paid to experienced teachers. As noted by Rose and Sonstelie (2010), the experience premium represents a rent for senior teachers who have tenure in their current positions and are unlikely to move to another district. If that claim is true, then under a rent seeking model of union behavior, we should expect teachers unions to care more about the premium paid to experienced teachers than administrators and the residents of a school district. The natural way to test that hypothesis is to examine whether the premium paid to experienced teachers is larger in districts that are likely to have more bargaining power, namely larger districts and districts located in states that mandate collective bargaining.

Our analysis allows such a test. Because the SASS contains a nationally representative sample of school districts, we are able to exploit variation in state collective bargaining laws to further isolate the causal impact of district size on teacher union bargaining power. We make use of the fact that in 33 states collective bargaining is mandatory while in five states collective bargaining is explicitly prohibited. If district size has a causal impact on the bargaining power of teacher unions we should expect that impact to be more pronounced in states with mandatory collective bargaining laws. More specifically, if our assumption that teachers care more about the premium paid to experienced teachers than other voters is correct, we should expect to find a stronger relationship between district size and the premium paid to experienced teachers among districts located in states that mandate collective bargaining.

1.3 Data

Our primary source of data is the restricted-use version of the 2007-2008 Schools and Staffing Survey (SASS), a nationally representative sample of schools and teachers that contains information on individual teachers, schools and districts. Since teacher salary schedules are typically set at the district level, our primary unit of analysis is the school district. As noted in section II, we focus on three outcomes of the bargaining process: beginning (base) teacher salaries, the salary premium paid to experienced teachers and the teacher-pupil ratio. District-level data on teacher salaries and teacher-pupil ratios comes from the survey of school district administrators, a subcomponent of the SASS. We measure the teacher-pupil ratio as the ratio of the number of full time equivalent teachers to the number of students attending a districts schools. We measure base salaries as the salary of a teacher with a bachelors degree and no teaching experience. We constructed two measures of the salary premium paid to experienced teachers. The first measure is the difference between the salary of a teacher with a bachelors degree and no teaching experience and the salary of a teacher with a bachelors degree and 10 years of experience. The second is the difference between the salary of a teacher with a masters degree and no teaching experience and the salary of a teacher with a masters degree and 10 years of experience. We constructed both measures by taking the natural log of the ratio of the two points in the salary schedule.[7]

We are primarily interested in how district size affects the bargaining power of teacher unions. Using district-level data from the 2005-2009 American Community Survey (ACS), we measure district size as the number of residents in a district who are citizens aged 18 and over. We also include an additional variable designed to measure the bargaining power of teachers unions. Babcock and Engberg (1999) focus on the idea that the bargaining power of unions will depend in part on the degree of community support for union activities. As they note, school board members are elected by local voters and thus voter sympathy for union objectives may influence bargaining outcomes. Since Democrats tend to be more sympathetic towards unions than Republicans, we proxy for the degree of community support for union activities using the fraction of voters within the county that

[7] We chose these two measures of the experience premium based on data availability. Specifically, the 2007-08 SASS provides only limited information on a districts salary schedule thus limiting our ability to construct other measures based on alternative differences in seniority. To examine whether our results were robust to alternative measures of the experience premium we used the 1999-2000 wave of the SASS, which contains alternative information of a districts salary schedule, to construct two other measures of the experience premium: 1) the difference between the salary of a teacher with a bachelors degree and no teaching experience and the salary of a teacher with a masters degree and 20 years of experience and 2) the difference between the salary of a teacher with a masters degree and no teaching experience and the salary of a teacher with a masters degree and 20 years of experience. Using these alternative measures of the experience premium yielded results that were qualitatively and quantitatively similar to those reported in Table 2. Results are available upon request.

a school district is located that voted for the Democratic presidential candidate during the 2004 presidential election. County-level data on the share of voters voting Democratic was obtained from POLIDATA, a company that compiles election results from official sources.

We use data from the 2007-2008 SASS and school district-level data from the 2005-2009 American Community Survey to create a number of variables that capture the economic and demographic characteristics of a districts students and residents. Those variables are: (1) median household income, (2) the fraction of the population age 25 or older with a college degree, (3) the fraction of homeowners, (4) the fraction of students eligible for free or reduced price lunch, (5) the fraction of students that are non-white, (6) a gini index of income inequality and (7) an index of racial heterogeneity. All of these variables are designed to capture district characteristics that might influence outcomes of the collective bargaining process. For example, since the capitalization of school quality into housing values directly affects the welfare of homeowners, we might expect to see greater monitoring and higher levels of engagement in school board politics in districts with more homeowners (Fischel, 2001; Grosskopf et al, 2001; Davis and Hayes, 1993). Similarly, teacher unions may attempt to bargain for higher wages or lower class sizes when the population of children they serve is more disadvantaged. Finally, the ability of teacher unions to influence school board elections is likely to be greater in communities with lower levels of civic engagement. Alesina and LaFerrara (2000) find that civic engagement is significantly lower in communities that are heterogeneous both in terms of race and income. Thus, we include an index of racial heterogeneity and a gini index of income inequality.[8] Finally, we include two additional variables in our analysis to account for the grade structure of school districts: an indicator variable that takes the value of unity if a district is an elementary district and an indicator variable that takes the value of unity if the district is a secondary district. The omitted group is unified districts.

In several of our specifications, we also include two additional variables, namely total revenue per pupil and average teacher experience. Both these variables may potentially affect outcomes of the collective bargaining process. Districts with greater revenue generating capacity may be able to pay higher salaries and achieve higher teacher-pupil ratios. Similarly, Babcock and Engberg (1999) find that districts with more experienced teachers tend to bargain for higher returns to experience. Of course, both revenue per pupil and teacher experience are likely endogenous. For example, if higher

[8] The racial heterogeneity index we employ is defined as: $H_j = 1 - \sum_{r=1}^{R} S_{ij}^2$, where S_{ij} is racial group rs share of the population in district j. Greater values of this index are associated with greater racial heterogeneity.

experience premiums induce experienced teachers to remain in a district, teacher experience and the premium paid to experienced teachers will be simultaneously determined. In light of that fact we follow Hoxby (1996) among others and control only for the plausibly more exogenous characteristics of a district in our baseline models. We then add revenue per pupil and teacher experience into specifications designed to examine the robustness of our results. Data on total revenue per pupil in each district in 2007-2008 comes from the National Center for Education Statistics, Common Core of Data. Since no comprehensive dataset on teacher experience is available, data on average teacher experience was obtained by contacting individual state education agencies. Through that process we were able to obtain administrative data on the average years of teaching experience in 2007-2008 for approximately 92% of the school districts in our sample.[9]

We restrict our sample in a number of ways. First, we drop charter schools, state-operated institutions and other non-traditional districts. Second, we drop districts that do not utilize a teacher salary schedule and therefore do not report information on the salary earned by teachers at various points in the salary schedule. We also exclude a small number of districts with missing data on district demographics. Table 1 provides the mean and standard deviation of the variables used in our analysis. We present these summary statistics for the sample of districts located in states that mandate collective bargaining and, for comparison purposes, the sample of school districts located in states that prohibit collective bargaining.

1.4 Empirical Framework

To examine the relationship between district size and outcomes of the collective bargaining process, we estimate a model of the following form:

$$Y_{jms} = \beta_0 + \beta_1 N_{jms} + \beta_2 N_{jms} * CBP_s + \beta_3 X_{jms} + \delta_m + \lambda_s + \varepsilon_{jms} \tag{1.1}$$

where Y_{jms} denotes one of the four outcomes of the collective bargaining process (base salaries, the two experience premium measures, and the teacher-pupil ratio) for school district j located in labor

[9] In specification where we use average teacher experience, the estimates reported in the paper are based on the sample for which we have administrative data on average teacher experience. However, we also constructed an estimate of the average years of teaching experience for the districts with missing data using the SASS teacher sample and estimated models where we replaced the missing administrative values for average teacher experience with the estimates from the SASS sample. The results based on this expanded sample where nearly identical to those reported in the paper.

market m and state s, N_{jms} is the number of eligible voters (district size), CBP_s is an indicator variable that takes the value of unity if a school district is located in a state that prohibits collective bargaining, X_{jms} is a vector of district-level control variables, δ_m and λ_s are teacher labor market and state fixed effects respectively and ε_{jms} is a random disturbance term.[10]

The coefficients of primary interest are β_1 and β_2, the coefficients on the number of eligible voters and the interaction term between the number of eligible voters and the indicator variable for districts located in states that prohibit collective bargaining, respectively. The estimated coefficient on the number of eligible voters represents the effect of district size on bargaining outcomes in states that mandate collective bargaining while the coefficient on the interaction term represents the difference in that effect for districts located in states that prohibit collective bargaining. The theoretical models of Rose and Sonstelie (2010) and Hoyt (1999) both predict that union bargaining power should increase with district size implying that β_1 should be positive in specifications where the dependent variable is either base salaries or the premium paid to experienced teachers. Furthermore, if our assumption that teachers care more about the premium paid to experienced teachers than other voters is correct, we should expect to find a stronger relationship between district size and the premium paid to experienced teachers in states that mandate collective bargaining, implying that β_2 should be negative in specifications where the dependent variable is the experience premium.

We proxy for teacher labor markets using the year 2000 commuting zones developed by the Economic Research Service of the United States Department of Agriculture. Similar to metropolitan statistical areas (MSAs) and core based statistical areas (CBSAs), commuting zones are designed to be spatial measures of local labor markets — geographical areas composed of counties with strong commuting ties. The main benefit of using commuting zones rather than CBSAs is that commuting zones are defined for the entire United States, not just for metropolitan areas. Thus, they allow us to include rural areas that lie outside of traditional metropolitan areas.[11] The inclusion of labor market (commuting zone) fixed effects implies that we are identifying the effect district size has on bargaining outcomes based solely on within labor market variation. Thus, our specification controls for any market-level unobservables that may be correlated with district size and/or outcomes of the collective bargaining process. We also include state fixed effects to control for the fact that some

[10]Note that because our model includes commuting zone and state fixed effects, we do not include the level effect of CBP_s in the model.

[11]Commuting zones located within metropolitan areas typically contain the same counties as core based statistical areas. Thus, the only real difference between CBSAs and commuting zones is that commuting zones are more comprehensive since they also include rural counties.

metro areas cross state boundaries. These state fixed effects further insulate our estimates from any state-level unobservables that might otherwise bias our estimates.[12]

Despite our inclusion of both labor market and state fixed effects, one still might be concerned that our estimates may be biased by other unobservables. For example, larger district may contain students that are more difficult to teach and therefore pay higher base salaries and higher experience premiums in order to attract and retain teachers. We attempt to mitigate that concern by including a rich set of control variables in our analysis including controls for the fraction of free and reduced price lunch students and the fraction of students that are non-white. We also note that if our results were being driven by unobservables related to larger districts having more difficult students to teach, we would expect to find that district size has a similar effect on teacher salaries in both states that mandate and states the prohibit collective bargaining. Thus, our interaction between district size and the indicator variable for states that prohibit collective bargaining serves as a test of that hypothesis. If we find that district size has a very different impact on the premium paid to experienced teachers in states that prohibit and states that mandate collective bargaining this would suggest that our results are not being driven by unobservables related to difficult to teach students.

Another potential concern with our identification strategy involves economies of scale. A number of papers have found significant economies of scale in the production of K-12 education, implying that larger districts face lower per-unit costs.[13] Given these lower costs, larger districts might rationally choose to use those cost savings to increase teacher salaries and the premium paid to experienced teachers and/or to increase the teacher-pupil ratio. However, if our results were being driven by economies of scale, rather than union bargaining power, the effect district size has on outcomes of the collective bargaining process should be similar for districts located in states that mandate and states that prohibit collective bargaining. Thus, one again, if we find that district size has a differential effect on bargaining outcomes in states that mandate and states that prohibit collective bargaining it would cast doubt on whether economies of scale could be an explanation for our results.

[12] For example, Figlio (1997) finds that state imposed property tax limitations are associated with lower teacher-pupil ratios and lower starting teacher salaries.

[13] See for example, Kenny (1982), Chakraborty, Biswas, and Lewis (2000), and Andrews et al. (2002).

1.5 Results

Results based on the estimation of Equation (1) are presented in Table 2. The standard errors reported in Table 2 and all subsequent tables are clustered at the commuting zone level to allow for within-local labor market area autocorrelation of the disturbance term. Each column of Table 2 represents results from separate regressions that include the same set of control variables but different outcomes of the collective bargaining process. All of our dependent variables as well as the number of eligible voters and household income are measured in natural logarithms.

We begin by noting that the estimated coefficient on the number of eligible voters is statistically significant at the one percent level or better in all four specifications. Thus, in states that mandate collective bargaining, base salaries and the premium paid to experienced teachers increase with district size, while the teacher-pupil ratio declines with district size. These results are highly consistent with Rose and Sonstelie (2010) who examine the relationship between district size and bargaining outcomes in California, a state that mandates collective bargaining. Similar to us, they find that base salaries and the salary premium paid to experienced teachers increase with district size while the ratio of teachers to pupils declines with district size.

In terms of magnitude, our results suggest that moving from one standard deviation below to one standard deviation above the mean in district size would increase base salaries by 4.5%. That finding is comparable to the estimates reported in the literature on the effect of unionization on salaries. For example, West and Mykerezi (2010) find that unionization raises beginning teacher salaries by 3.9 percent. Similarly, Hoxby (1996) finds that unionization increases average teacher salaries by approximately 5%. In terms of the premium paid to experienced teachers, the estimates reported in columns 2 and 3 suggest that moving from one standard deviation below to one standard deviation above the mean in district size would increase the premium paid to experienced teachers by 2.4% and 3.6% respectively. To get some sense of the magnitude of these effects, consider a district with an average beginning salary of $29,243 and an experience premium for a BA10/BA0 of 1.22, which correspond to the average beginning salary and average experience premium of districts in the first quartile of district size (which includes a district one standard deviation below the mean in district size). The estimated coefficient on district size reported in column 2 implies that a teacher with 10 years of experience and a BA located in a district one standard deviation above the mean of district size would have a salary that was 7% higher ($2,500) than a teacher with 10 years of experience and a BA in the smaller district. Under the same scenario (but using the starting salary for a teacher with an MA and the experience premium for a MA10/MA0) a teacher with 10 years of experience and a MA located in a district one standard deviation above the mean of district size would have a salary that was 8.3% higher ($3,300) than a teacher with 10 years of experience and a MA in the smaller district.

Turning to the estimated coefficient on the interaction term between district size and the indicator variable for states that prohibit collective bargaining, we note that in three of the four specifications the estimated coefficient on the interaction term is the opposite sign of the coefficient on the number of eligible voters and is statistically significant in all four specifications. While the premium paid to experienced teachers increases with district size in states that mandate collective bargaining, it declines with district size in states that prohibit collective bargaining.[14] Furthermore, relative to states that mandate collective bargaining, the positive effect district size has on base salaries (column 1) is about twice as large and the negative effect district size has on the teacher-pupil ratio (column 4) is about one half as large in states that prohibit collective bargaining.[15]

Based on the conceptual framework of Section II, we interpret the positive and statistically significant coefficient on district size in the experience premium specifications as suggesting that the bargaining power of teacher unions increases with district size.[16] If that interpretation is correct, then the results reported in Table 2 suggest that more powerful unions in large districts tend to bargain for more generous returns to teacher seniority at the expense of staffing ratios and base salaries.[17] That interpretation is reinforced by several other results in Table 2. As noted previously,

[14] Summing the coefficient on the number of eligible voters and the coefficient on the interaction term gives the effect of district size for districts located in states that prohibit collective bargaining. For the two experience premium specifications, summing the two coefficients results in a negative point estimate which, based on a joint hypothesis test, is statistically different from zero in both specifications.

[15] We also estimated two additional specifications where the dependent variable was either total teacher compensation (salary plus benefits) as a share of total expenditures or total revenue per pupil. We found that teacher compensation as a share of total expenditures increased with district size in both states that mandate and states that prohibit collective bargaining with the effect being slightly weaker in states that prohibit collective bargaining (i.e. the coefficient on the interaction term was negative although small in magnitude and only statistically significant at the 10% level). We found no difference in the relationship between district size and total revenue per pupil in states that mandate and states that prohibit collective bargaining.

[16] One potential concern with the results reported in Table 2 is that district size may be correlated with other market-level characteristics that influence outcomes of the collective bargaining process such as the alternative salary teachers could earn in another profession or the degree of inter-district competition. For example, Hoxby (1996) finds that teacher unions tend to be more powerful in areas with less inter-district competition as measured by a Herfindahl index. Recall, however, that all the specifications reported in Table 2 include labor-market fixed effects and thus control for any labor market characteristics (such as alternative wages and the amount of competition) that may influence bargaining outcomes.

[17] That interpretation is consistent with the recent findings of Goldhaber, DeArmond and DeBurgomaster (2010). Based on a 2006 survey of teachers in Washington State, they find that 83% of respondents preferred a salary increase of $5,000 to having two fewer students in their classes, suggesting that teachers are willing to trade off lower class sizes for higher salaries.

one alternative explanation for our results is that larger districts offer teachers more generous returns to experience in an attempt to retain more experienced teachers who might otherwise leave. That explanation seems particularly relevant given that larger districts tend to have higher percentages of disadvantaged students who may be more difficult to teach. We note, however, that the estimated coefficients on the fraction of students eligible for free and reduced price lunch and the fraction of students that are nonwhite are statistically insignificant in both experience premium specifications, casting doubt on this alternative explanation. Furthermore, the fact that we find a negative relationship between district size and the salary premium paid to experienced teachers in states that prohibit collective bargaining is also at odds with this alternative explanation.

The estimates reported in Table 2 also suggest that economies of scale are not the explanation for our results. As we noted previously if our results were being driven by economies of scale, rather than union bargaining power, the effect district size has on outcomes of the collective bargaining process should be similar in states that mandate and states that prohibit collective bargaining. In fact, however, we find that the effect district size has on bargaining outcomes varies significantly across states that mandate and states that prohibit collective bargaining. Furthermore, if economies of scale were the explanation for our results, we would expect larger district to use some of their cost savings to reduce class sizes.[18] In fact, however, we find that the teacher-pupil ratio declines with district size, suggesting once again that economies of scale are not the explanation of our results.

Robustness Checks

In Table 3 we conduct a series of robustness checks designed to further investigate alternative explanations for our main finding concerning district size and union bargaining power. In the interest of brevity we report only the estimated coefficients on the number of eligible voters and the interaction term between that variable and the collective bargaining prohibited indicator.[19] In panel A of Table 3 we once again present the estimated coefficients on district size and the interaction term between district size and the collective bargaining prohibited indicator from Table 2 for comparison purposes. Panel B reports estimates based on specifications that include only the key variables of interest (district size and the interaction term between district size and collective bargaining prohibited) and locational fixed effects (commuting zone and state fixed effects). We first note that

[18] This assumes that school districts are not capacity constrained and therefore not limited in their ability to reduce class sizes. With capacity constraints districts may not be able to reduce class sizes even if they have extra revenue available.

[19] We note, however, that with the exception of the estimates presented in Panel B, all the models we estimate contain all the control variables listed in Table 1.

dropping the control variables altogether does not substantively change our basic findings: all of the estimated coefficients remain statistically significant at the one percent level and the pattern of results is quite similar to the pattern observed in our baseline specification. While the estimated coefficient on district size increases in magnitude in the base salary specification and declines in the teacher-pupil ratio specification, it remains relatively unchanged in the two experience premium specifications. Furthermore, in all four specifications the estimated coefficients on the interaction term are nearly identical to those reported in Panel A. The fact that the estimated coefficients on district size and the interaction term in the two experience premium specifications remain unchanged when we add control variables provides us with additional confidence that the observed relationship between district size and our primary measure of union bargaining power (the salary premium paid to experienced teachers), is not being driven by district unobservables.[20]

The second issue we investigate is whether our results are sensitive to the inclusion of very large districts with large fractions of disadvantaged students. Once again, our concern is that large districts with difficult to teach students may pay higher experience premiums in order to retain more experienced teachers. We address that concern in two ways. First, we drop districts where the number of eligible voters is greater than 70,000 (approximately the 90th percentile of eligible voters in the sample) and the fraction of nonwhite students is greater than 0.50 (approximately the 75th percentile of fraction nonwhite in the sample). Second, we drop districts where the number of eligible voters is greater than 70,000 and the fraction of students eligible for free or reduced price lunch is greater than 0.40 (75th percentile of free or reduced price lunch students). The results of that exercise are reported in Panels C and D. Dropping the largest districts with the most disadvantaged students has little effect on our results. In all four specifications, the estimated coefficients on district size and the interaction term are similar in magnitude to the estimates reported in Panel A and remain statistically significant at the 1 percent level or better.

The third issue we investigate is whether our results are being driven by the inclusion of districts located in mandatory bargaining states that do not engage in collective bargaining. Approximately nine percent of districts located in mandatory collective bargaining states do not engage in collective

[20]One might also be concerned about whether the interaction term between district size and the indicator for collective bargaining prohibited states is correlated with unobservables. To get a better sense of whether the potential correlation of the interaction term with unobservables is a serious threat, we examined the correlation between the interaction term and the observable district characteristics listed in Table 1. With the exception of two variables all of the correlations were less than 0.2. For the remaining two variables, fraction non-white and the racial heterogeneity index, the correlation was -0.27 and -0.30 respectively.

bargaining and these districts tend to be small in size, raising the concern that our results could be driven by the inclusion of these districts. In Panel E we report results based on specifications that drop those districts. Once again, we find that our main results are quite robust.

The final issue we investigate is whether our results are sensitive to our exclusion of revenue per pupil and average teacher experience in our baseline models. Panel F displays estimates based on specifications where we include these two potentially endogenous variables. Controlling for revenue per pupil and average teacher experience has little impact on the base salary and the two experience premium specifications, it does however, cause the estimated coefficient on district size to decline by roughly half in the teacher-pupil ratio specification. Nevertheless, even in that specification the estimated coefficients on district size and the interaction term remain statistically significant at the 1 percent level and display the same pattern as our baseline results.[21]

Nonlinearities

In Table 4 we examine if there are any important nonlinearities in the relationship between district size and outcomes of the collective bargaining process. Specifically, we replace our continuous measure of district size with a set of indicator variables that take the value of unity if the number of eligible voters within a district is in the 2nd, 3rd, or 4th quartile of district size respectively. The omitted category is districts in the 1st quartile of district size, i.e. the smallest districts. We then interacted these quartile size indicator variables with the indicator variable for collective bargaining prohibited states and re-estimated our model.[22] The first three columns of Table 4 show a consistent monotonic relationship between district size and teacher salaries. In states that mandate collective bargaining, base salaries and the premium paid to experienced teachers increase monotonically with district size. Furthermore, as the coefficients on the interaction terms reveal, in states that prohibit collective bargaining, base salaries continue to increase monotonically with district size while the

[21]Dee (2004) and Oreopoulos and Salvanes (2009) find that more educated individuals tend to be more politically active, engage in more civic activities and be more likely to volunteer, suggesting they may be less prone to free riding. Thus, to examine whether our district size results were muted in districts with larger fractions of college educated voters, we also estimated a model where we interacted district size and the interaction between district size and the collective bargaining prohibited indicator with the fraction of residents within a district that had a college degree or higher. We found that the inclusion of these additional interaction terms had no effect on our base salary or experience premium results but that they muted the negative effect of district size on the teacher-pupil ratio in states that mandate collective bargaining. Specifically, the interaction between district size and the fraction of college educated voters was positive and statistically significant in this specification. The point estimates suggest that in a district where all voters are college educated, district size has only a small negative effect on the teacher-pupil ratio.

[22]In these specifications the estimated coefficients on the non-interacted quartile size indicators represent the differential effect district size has on bargaining outcomes relative to the omitted group (small districts) in collective bargaining mandatory states. Similarly, the estimated coefficients on the interacted quartile size indicators represent the differential effect district size has on bargaining outcomes relative to the omitted group in collective bargaining prohibited states.

premium paid to experienced teachers declines monotonically with district size. These results are highly consistent with those presented in Tables 2 and 3. In column 4, the teacher pupil ratio declines monotonically with district size in states that mandate collective bargaining but increases in a slightly non-monotonic fashion in states that prohibit collective bargaining.

Compensating Differentials

A number of studies have found evidence of compensating differentials between teacher salaries and pupil-teacher ratios, suggesting that teachers may be willing to accept lower salaries in exchange for better working conditions (i.e. higher teacher-student ratios). For example, using an instrumental variables approach, Hanushek and Luque (2000) find that increasing class size by one student increases beginning teacher salaries by approximately 0.9 percent and experienced teacher salaries by approximately 1.2 percent.[23] The possibility of such compensating differentials raises an interesting question. In the empirical work conducted thus far we have shown that in states that mandate collective bargaining, base salaries and the premium paid to experienced teachers increase with district size while the teacher-pupil ratio declines with district size. Because we view both salaries and teacher-pupil ratios as jointly determined outcomes of the collective bargaining process we have interpreted our findings as suggesting that larger and more powerful unions bargain for more generous returns to teacher seniority at the expense of staffing ratios. However, it is interesting to ask whether our salary results would persist if we held teacher-pupil ratios constant. In other words, do larger and more powerful unions continue to bargain for more generous returns to experience when the teacher-pupil ratio is held constant or only when there is a tradeoff between compensation and teacher-pupil ratios?

In Table 5 we present results based on specifications designed to address that question. The table presents instrumental variable regression results from models where we add the teacher-pupil ratio as an additional control variable to the base salary and two experience premium specifications.[24] Similar to Hanushek and Luque (2000) our instruments for the teacher-pupil ratio are based on past student enrollment growth. Specifically we use the growth rate of student enrollment within a district between 1997 and 2001 and between 2002 and 2006 as instruments for the teacher-pupil ratio. As noted by Hanushek and Luque (2000), if there is a lag between student enrollment growth

[23] Antos and Rosen (1975) and Levinson (1988) also present evidence consistent with compensating differentials between teacher salaries and pupil-teacher ratios but do not instrument for pupil-teacher ratios.

[24] Ideally we would also like to estimate a specification that allowed for compensating differentials between teacher-pupil ratios and teacher salaries in a model where the dependent variable was the teacher-pupil ratio. Unfortunately, however, we do not have a good instrument for teacher salaries so we have refrained from estimating such a model.

and the hiring of new teachers, past enrollment growth will be correlated with the teacher-pupil ratio. Furthermore, it seems unlikely that past enrollment growth is correlated with the error term in our salary regressions, a necessary condition for a valid instrument.[25]

In the interest of brevity, Table 5 reports only the estimated coefficients on district size, the interaction between district size and the collective bargaining prohibited indicator and the teacher-pupil ratio. We note, however, that all specifications include the full set of controls listed Table 1. Controlling for potential compensating differentials between salaries and teacher-pupil ratios does not substantively alter our basic findings. In the two experience premium specifications the estimated coefficients on district size and the interaction term are quite similar in magnitude to those presented previously and remain statistically significant. Controlling for compensating differentials has a larger effect on our base salary results. In column 1, the estimated coefficient on district size is about half the magnitude of the estimates reported in Tables 2 and 3 while the estimated coefficient on the interaction term is slightly larger than the estimate reported in those tables. However, the pattern of results for our base salary specification remains the same as in Tables 2 and 3: the estimated coefficients on district size and the interaction term are both positive suggesting that the positive relationship between base salaries and district size is stronger in states that prohibit collective bargaining.[26]

Free Riding versus Capitalization

The results reported in Tables 2 through 5 provide relatively compelling evidence that the bargaining power of teacher unions increases with district size. We now turn to examining the mechanism behind this relationship. As we noted in Section II, there are two alternative explanations for why unions may be more powerful in larger districts. The theory developed by Rose and Sonstelie (2010) predicts that unions are more powerful in larger districts because of the free rider problem. Political contributions and the time and energy spent promoting a particular school board member are public goods and as district size increases, the incentive to free ride off of the contributions and efforts of other voters increases. In contrast, unions can overcome the free rider problem by taxing their members in the form of dues. In the theory developed by Hoyt (1999) unions are more powerful

[25] We further note that our instruments pass several standard tests of validity, namely the Kleibergen-Paap under-identification test and the Hansen J-statistic test for over-identifying restrictions.

[26] We also estimated OLS models that do not instrument for the teacher-pupil ratio. For the two experience premium specifications the OLS results were quite similar to our 2SLS results. In the base salary specification, the estimated coefficients on district size and the interaction term were both positive and statistically significant, with point estimates of 0.010 and 0.017 respectively, while the estimate on the teacher-pupil ratio was negative and statistically significant with a point estimate of -0.048.

in larger districts because tax increases and inefficiencies in the allocation of resources are not fully capitalized into housing values in large districts. Consequently, homeowners have less of an incentive to monitor the behavior of school boards and unions in larger districts.

We attempt to differentiate between these two competing theories in several ways. We begin by exploiting the fact that 9 of the 33 states that mandate collective bargaining are right-to-work states.[27] In right-to-work states, teachers are not required to join a union or pay union dues. Consequently, teacher unions in states with right-to-work laws face free rider problems similar to those faced by voters. That in turn suggests our district size effects should be muted by the ability of teachers to free ride. To examine that possibility, we limited our sample to districts located in states that mandate collective bargaining and then created a dichotomous variable that takes the value of unity if a district is located in a right-to-work state and a value of zero otherwise. We then interacted this dichotomous variable with district size and included the interaction term in our specifications. The results of that exercise are reported in Panel A of Table 6.[28]

The results reported in the top panel of Table 6 provide some evidence that right to work laws mute the relationship between district size and outcomes of the collective bargaining process.[29] In both of the experience premium specifications the estimated coefficient on the interaction term is negative and statistically significant. The magnitude of the estimated coefficients reported in columns 2 and 3 suggest that the relationship between district size and the premium paid to experienced teachers is about one half as large in states with right-to-work laws. These results are generally consistent with the free riding hypothesis of Rose and Sonstelie (2010).

As an alternative method of differentiating between the two theories we examine whether our results are sensitive to the housing supply elasticity of metropolitan areas. Recall that the theory developed by Hoyt (1999) predicts that unions will be more powerful in larger districts because inefficiencies in the allocation of resources are not fully capitalized into housing prices in larger districts,

[27]The nine states are: Florida, Iowa, Idaho, Kansas, North Dakota, Nebraska, Nevada, South Dakota, and Tennessee.

[28]As in previous tables we report only the estimated coefficients on the number of eligible voters and the interaction term between the number of eligible voters and location in a right-to-work state, but note that all the models we estimate include the full set of control variables.

[29]A potentially confounding factor is that union members might be more likely to live in larger jurisdictions. In support of that notion, unionization rates (number of union members/voting age population) are higher larger counties (based on authors calculations). Thus, even though workers are not required to join a union in right-to-work states, larger districts in right-to-work states may have higher union membership rates than smaller districts, leading to greater union bargaining power in these large right-to-work districts. Unfortunately, data on district-level unionization rates is unavailable and thus we are unable to test that hypothesis.

while they are fully capitalized in smaller districts. Thus, the theory developed by Hoyt (1999) is relevant for areas where a more inelastic supply of housing leads to capitalization. Consequently, if the capitalization hypothesis of Hoyt (1999) is the mechanism behind our results, we would expect district size to have little effect on the bargaining power of teacher unions in areas with an elastic supply of housing.[30]

To examine that possibility, we estimate models similar to those reported in Panel A of Table 6, except we replace the interaction term between district size and the indicator for right to work states with an interaction between district size and an estimate of the housing supply elasticity in each metro-area in our sample. The metro-area housing supply estimates we employ are those reported by Saiz (2010) and represent estimates of the housing supply elasticity in 269 major metropolitan areas. Note that since our initial sample contains both metro and non-metro areas (for which we have no housing supply estimates) and Saiz (2010) only reports housing supply estimates for larger metro areas, our sample size is reduced.[31] Thus, for comparison purposes Panel B presents estimates of the effect of district size on bargaining outcomes based on the sample for which we have housing supply elasticity estimates. In general the estimated coefficients on district size reported in Panel B are quite similar to those reported in Table 2. The one exception is the estimated coefficient in the teacher-pupil specification which is roughly half the magnitude of the estimate reported in Table 2.[32]

In Panel C we report estimates based on specifications that interact district size with the housing supply elasticity estimates. With the exception of the base salary specification, all of the the estimated coefficients on district size reported in Panel C are statistically significant and similar in magnitude to those reported in Panel B. Furthermore, none of the estimated coefficients on the interaction term are statistically significant. As a robustness check in Panel D we report estimates based on specifications that interact district size with an indicator variable that takes the value of unity if a district is located in a metro-area with a housing supply elasticity above the median

[30]Brasington (2002) finds that the capitalization of school quality into housing values is significantly weaker near the edge of urban areas where housing supply is more elastic. Similarly, using data from New Hampshire, Lutz (2009) finds no evidence that intergovernmental grants were capitalized into housing values for the state as a whole (which is primarily rural), but strong evidence that grants were capitalized into housing values within the Boston suburban ring. Finally, using data from Massachusetts, Hilber and Mayer (2009) find that the capitalization of school spending into housing values is sensitive to the amount of developed land. Similar to Lutz (2009) they find no evidence that school spending is capitalized into housing values in communities with significant amounts of developable land but strong evidence that spending is capitalized into housing values in communities with little developable land.

[31]In addition, because only a relatively small number of the metropolitan areas located in states that prohibit collective bargaining are included in the Saiz (2010) sample, we focus our analysis solely on districts located in states the mandate collective bargaining.

[32]The decline in the magnitude of the estimated coefficient on district size in the teacher-pupil ratio specification is most likely attributable to the fact that our sample now contains only larger metropolitan areas and thus fewer small and rural districts.

elasticity of 2.24. Once again with the exception of column 1, where we observe a positive and statistically significant effect of the housing supply elasticity on base salaries, none of the estimated coefficients on the interaction term are statistically significant while all of the estimated coefficients on district size are similar in magnitude to those reported in Panel B and statistically significant. Thus, the results reported in Panels C and D provide little evidence that our results are sensitive to metro-area housing supply elasticities. That finding cast doubt on whether the capitalization hypothesis of Hoyt (1999) could be an explanation for our results.

1.6 Conclusion

Studies that examine the effect of teacher unionization on educational outcomes and the allocation of school resources typically do so by comparing districts with and without unions. However, even within the context of a setting where all districts are unionized, the effect unions have on educational outcomes and the allocation of school resources may be quite heterogeneous. In this paper we provide new evidence consistent with that notion. Building on the work of Rose and Sonstelie (2010) we examine how district size affects the bargaining power of teachers unions and the subsequent allocation of school resources. We find strong support for the prediction that union bargaining power increases with district size. In states that mandate collective bargaining, base salaries and the premium paid to experienced teachers increase with district size while the ratio of teachers to pupils declines with district size. In contrast, in states that prohibit collective bargaining we find a negative relationship between district size and the premium paid to experienced teachers. Furthermore, the positive relationship between district size and base salaries is stronger while the negative relationship between district size and the teacher-pupil ratio is weaker in states that prohibit collective bargaining. Collectively, these results suggest that larger and more powerful unions bargain for higher salary premiums for experienced teachers at the expense of staffing ratios and base salaries. Finally, we attempt to differentiate between two competing explanations for the observed relationship between district size and union bargaining power. While not definitive, our results appear more consistent with the free riding hypothesis of Rose and Sonstelie (2010) than the capitalization hypothesis of Hoyt (1999).

Our findings concerning the relationship between district size and union bargaining power may have several important policy implications. Over the last four decades, there has been a substantial reduction in the number of school districts operating within the United States. A common argument for district consolidation is that it leads to cost savings due to economies of scale. Our results point to an alternative and unintended consequence of school district consolidation, namely an increase in union bargaining power that may lead to a reallocation of resources that favors salary increases for senior teachers over other productive school inputs. Consistent with that notion, Barrow and Rouse (2004) examine how district size affects school district efficiency. They find that there is a significant difference between small and large districts, suggesting that large districts are more likely to overspend than small districts. While the authors could not identify the source of the inefficiency difference, our results point to a possible mechanism: larger districts tend to have more powerful unions which results in differences in the allocation of school resources across large and small districts.

Table 1.1: Summary Statistics

	Districts in States with Mandatory Collective Bargaining		Districts in States where Collective Bargaining is Prohibited	
Dependent Variables	Mean	St. Dev.	Mean	St. Dev.
Base Salary ($1,000)	33.244	5.458	34.090	4.606
Experience Premium [BA 10 / BA 0]	1.284	0.158	1.220	0.094
Experience Premium [MA 10 / MA 0]	1.312	0.156	1.219	0.093
Teacherpupil ratio	0.071	0.023	0.070	0.013
District or County Independent Variables	0.000	0.000	0.000	0.000
Number of Voters (1,000)	30.249	91.156	57.309	90.425
County Democrat Presidential Vote 2004	0.440	0.124	0.388	0.121
Fraction Homeowner	0.738	0.116	0.712	0.110
Fraction Students Nonwhite	0.242	0.260	0.455	0.262
Fraction Free or Reduced Price Lunch	0.380	0.209	0.485	0.185
Median Household Income ($1,000)	52.595	18.893	45.860	14.957
Fraction College Educated	0.316	0.132	0.273	0.111
Index of Racial Heterogeneity	0.213	0.176	0.365	0.159
Income Gini Coefficient	0.413	0.044	0.435	0.043
Total Revenue per Pupil ($1,000)	12.141	4.286	11.044	3.004
Average Teacher Experience	14.013	3.095	12.849	2.229
Rural District	0.412	0.492	0.481	0.500
Elementary District	0.082	0.274	0.004	0.066
High School District	0.025	0.157	0.000	0.000
Number of States	33		5	
Number of Commuting Zones	408		111	
Number of Districts	2,880		460	

Notes: Sample sizes are rounded to the nearest 10 to comply with restricted-use data reporting requirements.

Table 1.2: Baseline Estimates

	(1) Base Salary	(2) Exp. Premium [BA 10 /BA 0]	(3) Exp. Premium [MA 10 /MA 0]	(4) Teacher Pupil Ratio
Log of Number of Voters	0.015***	0.008***	0.012***	-0.091***
	(0.003)	(0.002)	(0.002)	(0.007)
Log of Number of Voters*CB Prohibited	0.015***	-0.024***	-0.028***	0.046***
	(0.005)	(0.004)	(0.004)	(0.008)
County Democrat Presidential Vote 2004	0.059*	0.020	0.043*	-0.036
	(0.032)	(0.031)	(0.024)	(0.052)
Fraction Homeowner	-0.034	0.012	0.009	-0.113*
	(0.024)	(0.030)	(0.030)	(0.062)
Fraction Students Nonwhite	0.018	0.018	0.012	-0.013
	(0.018)	(0.015)	(0.015)	(0.034)
Fraction Free or Reduced Price Lunch	0.002	0.009	0.014	0.195***
	(0.018)	(0.020)	(0.017)	(0.048)
Log of Median Household Income	0.054***	-0.005	0.006	-0.119***
	(0.014)	(0.016)	(0.016)	(0.042)
Fraction College Educated	0.058**	0.034	0.034	0.440***
	(0.029)	(0.035)	(0.036)	(0.076)
Index of Racial Heterogeneity	0.052**	-0.021	-0.007	0.014
	(0.022)	(0.025)	(0.027)	(0.050)
Income Gini Coefficient	0.079*	-0.006	-0.021	0.111
	(0.047)	(0.049)	(0.046)	(0.162)
Elementary District	0.003	0.009	0.008	-0.010
	(0.008)	(0.007)	(0.006)	(0.022)
High School District	0.056**	0.036**	0.033**	0.024
	(0.023)	(0.015)	(0.016)	(0.024)
Observations	3,360	3,360	3,360	3,360
R-Square	0.854	0.660	0.715	0.734

Notes: Estimates are based on sample of school districts located in states with mandatory collective bargaining laws and states that prohibit collective bargaining. Each column presents estimates from a separate regression where the dependent variable is listed in the top row. All specifications include commuting zone and state fixed effects. Sample sizes are rounded to the nearest 10 to comply with restricted-use data reporting requirements. Robust, clustered standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

Table 1.3: Robustness Checks

	(1)	(2)	(3)	(4)
	Base Salary	Exp. Premium [BA 10 /BA 0]	Exp. Premium [MA 10 /MA 0]	Teacher Pupil Ratio
A. Baseline (N=3,360)				
Log of Number of Voters	0.015***	0.008***	0.012***	-0.091***
	(0.003)	(0.002)	(0.002)	(0.007)
Log of Number of Voters*CB Prohibited	0.015***	-0.024***	-0.028***	0.046***
	(0.005)	(0.004)	(0.004)	(0.008)
B. No Controls (N=3,360)				
Log of Number of Voters	0.023***	0.009***	0.013***	-0.070***
	(0.002)	(0.002)	(0.002)	(0.006)
Log of Number of Voters*CB Prohibited	0.016***	-0.024***	-0.028***	0.040***
	(0.005)	(0.004)	(0.004)	(0.008)
C. Dropping Large Districts High Minority Concentration (N=3,140)				
Log of Number of Voters	0.016***	0.009***	0.013***	-0.100***
	(0.003)	(0.002)	(0.002)	(0.007)
Log of Number of Voters*CB Prohibited	0.014**	-0.024***	-0.029***	0.056***
	(0.006)	(0.004)	(0.004)	(0.011)
D. Dropping Large Districts HighFree or Reduced Price Lunch Concentration (N=3,240)				
Log of Number of Voters	0.016***	0.008***	0.012***	-0.096***
	(0.003)	(0.002)	(0.002)	(0.007)
Log of Number of Voters*CB Prohibited	0.013**	-0.021***	-0.026***	0.054***
	(0.006)	(0.005)	(0.004)	(0.009)
E. Dropping Districts in CB Mandatory States that Don't Bargain (N=3,100)				
Log of Number of Voters	0.014***	0.009***	0.012***	-0.085***
	(0.002)	(0.002)	(0.002)	(0.006)
Log of Number of Voters*CB Prohibited	0.014***	-0.024***	-0.027***	0.041***
	(0.005)	(0.004)	(0.004)	(0.008)
F. Adding Potentially Endogenous Controls (N=3,100)				
Log of Number of Voters	0.015***	0.009***	0.013***	-0.049***
	(0.002)	(0.002)	(0.002)	(0.004)
Log of Number of Voters*CB Prohibited	0.016***	-0.024***	-0.029***	0.040***
	(0.005)	(0.004)	(0.004)	(0.006)

Notes: Table provides estimates based on the sample of school districts located in states with mandatory collective bargaining laws and states that prohibit collective bargaining. Each column presents estimates from a separate regression where the dependent variable is listed in the top row. All specifications include commuting zone and state fixed effects and, with the exception of Panel B, the full set of control variables listed in Table 1. Panel A reproduces the results from Table 3. Panel B provides estimates when no control variables other than locational fixed effects (commuting zone, state, and rural fixed effects) are included. Panel C drops large districts (over 70,000 students) with a high concentration on minorities (40% or over). Panel D drops large districts (over 70,000 students) with a high concentration of free or reduced price lunch students (50% or over). Panel E drops districts located in mandatory collective bargaining states that do not engage in collective bargaining. Panel F presents results based on specifications that also include controls for the log of total revenue per pupil and average teacher experience. Sample sizes are rounded to the nearest 10 to comply with restricted-use data reporting requirements. Robust, clustered standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

Table 1.4: Nonlinear District Size Estimates

	(1)	(2)	(3)	(4)
	Base Salary	Exp. Premium [BA 10 /BA 0]	Exp. Premium [MA 10 /MA 0]	Teacher Pupil Ratio
District Size Q2	0.024***	0.016***	0.016***	-0.175***
	(0.005)	(0.005)	(0.005)	(0.012)
District Size Q3	0.039***	0.026***	0.029***	-0.212***
	(0.007)	(0.006)	(0.006)	(0.015)
District Size Q4	0.042***	0.030***	0.039***	-0.227***
	(0.007)	(0.007)	(0.008)	(0.016)
District Size Q2*CB Prohibited	0.040*	-0.077***	-0.071***	0.089**
	(0.021)	(0.016)	(0.016)	(0.034)
District Size Q3*CB Prohibited	0.073***	-0.110***	-0.111***	0.095***
	(0.022)	(0.015)	(0.016)	(0.035)
District Size Q4*CB Prohibited	0.086***	-0.121***	-0.131***	0.085**
	(0.025)	(0.016)	(0.016)	(0.036)
Observations	3,360	3,360	3,360	3,360
R-Square	0.854	0.662	0.714	0.724

Notes: Estimates are based on sample of school districts located in states with mandatory collective bargaining laws and states that prohibit collective bargaining. Each column presents estimates from a separate regression where the dependent variable is listed in the top row. All specifications include commuting zone and state fixed effects and the full set of control variables listed in Table 1. Sample sizes are rounded to the nearest 10 to comply with restricted-use data reporting requirements. Robust, clustered standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

Table 1.5: Compensating Differential Estimates

	(1)	(2)	(3)
	Base Salary	Exp. Premium [BA 10 /BA 0]	Exp. Premium [MA 10 /MA 0]
Log of Number of Voters	0.007	0.009**	0.013***
	(0.005)	(0.004)	(0.005)
Log of Number of Voters*CB Prohibited	0.019***	-0.024***	-0.029***
	(0.005)	(0.004)	(0.004)
Log of Teacher-Pupil Ratio	-0.088*	0.012	0.015
	(0.048)	(0.038)	(0.045)
Observations	3,270	3,270	3,270
R-Square	0.855	0.660	0.715

Notes: Estimates are based on sample of school districts located in states with mandatory collective bargaining laws and states that prohibit collective bargaining. Each column presents 2SLS estimates from a separate regression where the dependent variable is listed in the top row. All specifications include commuting zone and state fixed effects and the full set of control variables listed in Table 1. The teacher-pupil ratio in each specification is instrumented for using prior enrollment growthestimates. Sample sizes are rounded to the nearest 10 to comply with restricted-use data reporting requirements. Robust, clustered standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

Table 1.6: Mechanisms

	(1)	(2)	(3)	(4)
	Base Salary	Exp. Premium [BA 10 /BA 0]	Exp. Premium [MA 10 /MA 0]	Teacher Pupil Ratio
A. Districts in Collective Bargaining Mandatory States (N=2,900)				
Log of Number of Voters	0.016***	0.009***	0.014***	-0.089***
	(0.003)	(0.003)	(0.003)	(0.009)
Log of Number of Voters* Right to Work	-0.003 (0.005)	-0.005* (0.003)	-0.008** (0.003)	-0.012 (0.011)
B. Districts in CB Mandatory States with Housing Supply Elasticity Estimates (N=1,390)				
Log of Number of Voters	0.011***	0.008***	0.013***	-0.050***
	(0.003)	(0.002)	(0.002)	(0.007)
C. Districts in CB Mandatory States with Housing Supply Elasticity Estimates (N=1,390)				
Log of Number of Voters	0.005	0.009*	0.013***	-0.046***
	(0.005)	(0.005)	(0.005)	(0.015)
Log of Number of Voters* Housing Supply Elasticity	0.003 (0.002)	-0.001 (0.002)	-0.000 (0.002)	-0.002 (0.005)
D. Districts in CB Mandatory States with Housing Supply Elasticity Estimates (N=1,390)				
Log of Number of Voters	0.007**	0.008**	0.013***	-0.050***
	(0.003)	(0.003)	(0.003)	(0.011)
Log of Number of Voters* Above Median Housing Supply Elasticity	0.011** (0.005)	-0.001 (0.005)	-0.001 (0.005)	0.001 (0.011)

Notes: Panels A through D provides estimates based on the sample of school district located in states with mandatory collective bargaining laws. Panel A reports estimates when district size is interacted with an indicator variable that takes the value of unity for right to work states. Panel B reports estimates based on the sample for which we have estimates of the metropolitan housing supply elasticity from Saiz (2010). Panel C reports estimates when district size is interacted with the metropolitan housing supply elasticity estimates of Saiz (2010). Panel D reports estimates when district size is interacted with an indicator variable that takes the value of unity for districts located in metropolitan areas with an above median housing supply elasticity (elasticity=2.24). Each column presents estimates from a separate regression where the dependent variable is listed in the top row. All specifications include commuting zone and state fixed effects and the full set of control variables listed in Table 1. Sample sizes are rounded to the nearest 10 to comply with restricted-use data reporting requirements. Robust, clustered standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

Chapter 2

Does Malaria Affect Long Term Growth? Enlightening Results.

2.1 Introduction

Malaria is one of the worlds most devastating diseases, with a high burden of mortality and morbidity. The disease is also heavily concentrated in poor countries. Over the last 15 years, there has been an intense debate among economists regarding the role that malaria plays in inducing poverty.

Proponents of the view that malaria has large economic effects have cited several possible channels. One is the direct impact on adult productivity from malaria episodes. Second is damage to human capital of children (Lucas (2010), Cutler, Fung, Kremer, Singhal, and Vogl (2010), Bleakley (2010), and Ashraf, Lester, and Weil (2008)). Third has been the relocation of economic activity away from potentially productive sites. However, a number of critics have questioned the importance of malaria for economic outcomes. Malaria mortality and morbidity are highly concentrated among the young; most adults in endemic areas are largely immune to the disease. Malaria is concentrated in the tropics, which might have low income for other reasons, including low agricultural productivity and the presence of other tropical diseases. Finally, some critics (Acemoglu, Johnson, and Robinson (2001)) have argued that the correlation of malaria with low income results from the disease's impact on institutions historically, but that there is no current effect of the disease.

Some of the most important evidence regarding malaria's impact on economic activity comes from a series of cross country studies. (Gallup and Sachs (2001) and Sachs (2003)). Gallup and Sachs regress malaria prevalence on GDP per capita at the country level. Their measure of malaria is the percentage of a country's population living in areas with a risk of malaria transmission. They find a strong negative correlation between economic activity and this measure. Sachs (2003) goes a step beyond this by using an instrument to correct for the endogeneity of malaria. Sachs (2003) specifically estimates the effect of malaria on the log of GDP per capita, at the country level, instrumenting for malaria with the share of a country's population in the temperate zone and an index of malaria ecology. This malaria ecology index measures the stability of malaria transmission. It was created by determining the dominant malaria carrying mosquito in a region and using it's biting and breeding characteristics. This paper finds that a 20 percentage point decrease in malaria prevalence would lead to a 28% increase in GDP per capita.

Critics have raised three major concerns about the results presented by Sachs.[1] The first critique is that the unit of observation is the country level. We would expect country level regressions to suffer from omitted variable bias since countries have important unobservables such as health, education, culture, and institutions.

The second major criticism is that the malaria prevalence measure does a poor job of measuring the actual impact of the disease. The malaria prevalence measure is based on a World Health Organizations (WHO) report which estimates possibility of transmission at the regional level from which Sachs then constructs a weighted average at the national level.[2] The data underlieing the WHO estimates is in turn the number of reported cases which, according to the WHO, is less then ten percent of actual cases. Even worse the WHO only has reports from less then half of the most malarious countries in Africa. For the areas that they do not have data from it is unclear how the WHO determines if the area has high, limited or no risk.

The third critique is that the literature is unable to rule out the possibility the estimated impact of malaria is being biased upward by the correlation with other mosquito born diseases. Mosquito born diseases such as malaria, west nile, dengue fever and yellow fever are highly correlated in their location throughout the world. Therefore the coefficient on malaria may just be picking up the impact of other diseases. These other diseases lead to over 60 million infections a year and, in some

[1] Acemoglu, Johnson, and Robinson (2001); Rodrik, Subramanian, and Trebbi (2004); Easterly and Levine (2003)
[2] WHO (1997)

areas, are more prevalent than malaria.[3]

On the other side of the debate, researchers argue that malaria should have little impact on economic activity both because in endemic areas adults would have acquired significant immunity over their lifetimes and a significant fraction of the population carries genetic resistance to the disease. The only impact they expect malaria to have on households is through the mortality rates of children, which they do not expect to impact economic activity.

Critics of the Sachs view argue that the observed relationship between malaria and poverty is not through productivity but through institutions. AJR show that locations with a high prevalence of malaria had a larger fraction of the initial colonial settlers die which they call settler mortality.(Acemoglu, Johnson, and Robinson (2001) (AJR)) They argue that because the initial settlers had higher mortality they did not want to stay for long and therefore set up extractive institutions. In their studies they find that once you control for malaria's impact on institutions, malaria has no significant impact on economic activity. (Rodrik, Subramanian, and Trebbi (2004); Easterly and Levine (2003))

These papers suffer from similar issues to Sachs (2003) in that they also use countries as the unit of observation, the same malaria measure that he does, and do not control for other mosquito borne diseases. A further problem they have is that either they do not use an instrument for malaria, or their instrument does not meet the exclusion restriction. RST have no instrument for malaria, while AJR instruments for malaria with latitude, mean temperature and distance to coast, which have been shown to impact GDP through other channels, violating the exclusion restriction.

In this paper I use new data and techniques to overcome all of the issues in the literature and present robust findings on the impact of malaria on economic activity.

First, to solve the problem of omitted variable bias common in cross country regressions, I use data on a finer geographic scale. This finer scale is generated by breaking the world up into a grid where each quarter degree by quarter degree grid square represents one observation in the dataset. This is only possible because of the availability of a new finer geographic measure of economic activ-

[3]World wide there are approximately 219 million cases with 660,000 deaths from malaria.(http://www.who.int/features/factfiles/malaria/en/)

ity, nighttime lights, which has been shown to be a good proxy for GDP by Henderson, Storeygard, and Weil (2012), Hodler and Raschky (2014), and Michalopoulos and Papaioannou (2011). Using this data, along with gridded population, allows me to use fixed effects to remove the bias from omitted variables at the country level.

Second, previous research used a malaria prevalence measure that is based on likely transmission to a population rather then actual prevalence. In order to overcome this drawback this paper uses the percentage of the population between two and ten who have the malaria parasite in their blood. This measure was created by the Malaria Atlas Project by combining multiple surveys of random or community level blood sampling and then extrapolating these points to the entire earth.[4]

Third, the problem of other mosquito borne diseases being correlated with malaria has never been addressed in the literature. I take advantage of the fact that the genus of mosquito that carries malaria is different from the genus that carries all other mosquito born diseases. Further only the genus that carries malaria can survive above 1,000 meters. By looking at a sub sample from above this altitude I am able to estimate the impact of malaria on economic activity without any bias from other diseases.

Fourth, to solve the issue of poor instruments, I use malaria ecology, one of the instruments in Sachs (2003). While Sachs used this measure aggregated up tot the level of countries I take advantage of the continuous measure originally created. The malaria ecology data comes from Kiszewski, Mellinger, Malaney, Spielman, Ehrlich, and Sachs (2004) who provide a continuous map of the stability of malaria transmission based on geographic, temperature and climate data. This measure was designed to fit the exclusion restriction and I show evidence that this is the case.

After correcting for all of these issues, I find a significant negative impact of malaria on economic activity. The results suggest that a 20 percentage point decrease in malaria would cause an 80% increase in nighttime lights. Using the estimated conversion from nighttime lights to GDP from HSW and Holder and Raschky this would mean that approximately a 20 percentage point decrease in malaria would lead to a 20% increase in GDP. These results are robust to limiting my sample to only places where malaria carrying mosquitoes live and a range of other robustness checks.

[4]Hay, Guerra, Gething, Patil, Tatem, Noor, Kabaria, Manh, Elyazar, Brooker, Smith, Moyeed, and Snow (2009)

The rest of the paper is organized as follows: Section 2 describes the data. In Section 3, I create an empirical specification to test my model correcting for problems in the previous literature. In Section 4, I begin by repeating the cross country analysis that has been done in the literature, finding similar results. Next I use new sub national data to estimate my model with country fixed effects. In Section 5, I check the robustness of these results in order to rule out a host of alternative explanations. In Section 6, I bring to bear data on population at the sub national level in order to separate the estimated effect of malaria on lights into a piece representing the impact of malaria on lights per capita and a piece representing the impact of malaria on population density. Section 7 concludes.

2.2 Data

To estimate this model, I use data from 7 different sources to provide information on nighttime lights, malaria prevalence, malaria ecology, land suitability, distance to the coast, distance to fresh water, absolute latitude, and biomes. Summary statistics of this data are presented in Table 1.

Lights

To measure economic activity, I follow Henderson, Storeygard, and Weil (2012) (HSW) using lights seen from space, as shown in Figure 2.1, provided by the National Geophysics Data Center (NGDC) at NOAA. The data was collected from US Air force weather satellites which circle the globe 14 times a day collecting data on light intensity everywhere on earth. The data is collected from 8:30pm to 10:00pm, between 65 degrees North and 65 degrees South, and averaged yearly across cloud free nights. The scale of the data is from 0 to 63 with higher values representing higher light intensity. The data is censored at 63, but this constraint affects an extremely small number of data points, almost none of which are in areas where malaria is prevalent. [5] The data were collected from 1992 to 2009 with one pixel representing approximately one square kilometer at the equator.

[5] All results were checked for robustness to excluding the top coded areas.

Figure 2.1: Lights Seen From Space (2009)

HSW demonstrated a strong correlation between nighttime lights data and GDP, suggesting nighttime lights can be used as a proxy for GDP. HSW shows that the value of lights as a proxy is particularly useful in developing regions where national accounts data is relatively poor. Malaria is primarily confined to the developing world. Using nighttime lights provides a significant benefit over other GDP measures in that it is continuous over almost the entire surface of the earth allowing for a detailed examination of determinants of GDP village by village instead of aggregated at a country level.

Malaria

The endemicity of malaria, shown in Figure 2.2, is from the Malaria Atlas Project (MAP) (2009). MAP represents the probability of an individual between the age of 2-10 having the malaria parasite in his or her blood. This data was created by using 7,953 surveys from around the world with each survey representing either a random sample or a community level survey. The data was collected from 1985-2008 and was then predicted forward using a standardization model with Bayesian updating. The data was then turned into a continuous map by applying a Monte Carlo simulation to create a distribution for each grid cell based on the location of each survey and the number of people surveyed. In this paper, I focus on the mean of this distribution, but also include the standard deviation for a robustness check. The variables used to create this continuous mapping were the information provided by the surveys, and a measure of urban, peri-urban and rural. The measure of urbanization entered into the equation as two dummy variables and no other environmental, social, or economic variables were used for the creation.[6]

[6] For a detailed explanation see Hay, Guerra, Gething, Patil, Tatem, Noor, Kabaria, Manh, Elyazar, Brooker, Smith, Moyeed, and Snow (2009)

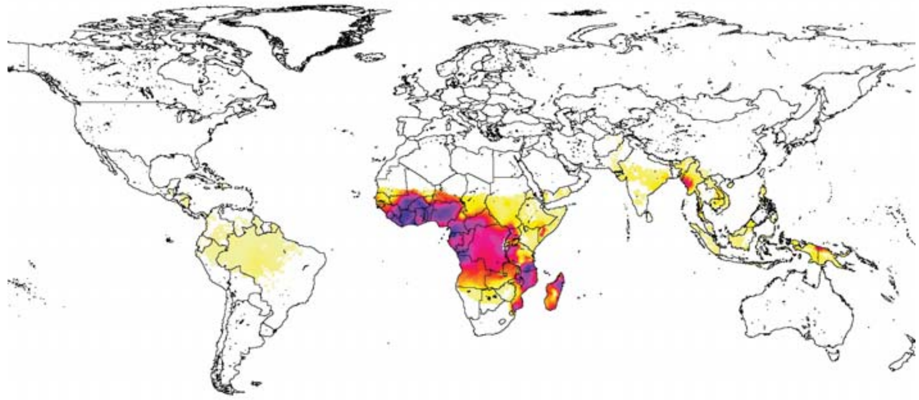

Figure 2.2: Endemicity of Malaria from the MAP (2009)

The malaria measure presents a few possible concerns. First, it is not based on a survey from every grid cell but rather is a smoothed function over areas. Second, the surveys used are not all from the same year and therefore they were time standardized using a standardization model with Bayesian updating. Regarding the first concern, I do multiple corrections for spatial autocorrelation. This helps to reduce the weight of individual surveys while not removing all variation. In terms of the second concern, given that the time shifting is a nonlinear procedure, I am unable to fully address what bias it would introduce into my estimates. I argue that, given that I am looking at GDP and not the growth or changes in GDP, this would not bias my results conclusively in one direction or the other, but I can provide no evidence. Since the malaria surveys were not uniformly taken across the world, some grid squares are more noisy than others. However, since MAP provides the standard deviation, I can correct for this increase in uncertainty by weighting my estimates. Overall, I do not expect these problems to systematically bias my results. Any resulting attenuation bias should be removed by instrumenting.

The malaria ecology data is from Kiszewski, Mellinger, Malaney, Spielman, Ehrlich, and Sachs (2004). It provides a continuous map of the stability of malaria transmission in an area. Malaria ecology gives a measure of how well malaria would survive and spread in an area if humans were placed there. It is based on individual mosquito characteristics, temperature, altitude, the number of days with frost, ecosystem classifications, and rainfall data. The data is constructed by first determining the major malaria carrying mosquito for each part of the world and then collecting individual characteristics on that mosquito such as percentage of the mosquito's diet derived from humans, survivability and daily mortality. Interacting these variables with geographic characteristics allows for the creation of a continuous map of the stability of malaria transmission.

Other Measures

Other measures used in this paper are national boundaries, latitude, land suitability, Ecoregions, distance to the coast, distance to fresh water, and elevation. National boundaries, distance to the coast, and latitude were acquired from the Gridded Population of the World and Global Rural-Urban Mapping Project. The Land Suitability measure is the probability that a grid cell will be cultivated as a function of only soil quality and climate variables.[7] Ecoregion data was collected from the World Wildlife Foundations "Terrestrial Ecoregions of the World: A New Map of Life on Earth". Of particular use was the biomes data which provided information on desert locations around the world.(Olson, Dinerstein, Wikramanayake, Burgess, Powell, Underwood, D'amico, Itoua, Strand, Morrison, Loucks, Allnutt, Ricketts, Kura, Lamoreux, Wettengel, Hedao, and Kassem, 2001) Distance to water was created by combining bodies of water over 50km^2 from Lehner and Doll (2004), and World Major Rivers provided by ESRI. Elevations were acquired from the Iscience Elevation and Depth map.[8]

All the data was aggregated up to the level 0.25 degree latitude by 0.25 degree longitude[9], with each data point representing the mean of the underlying data, with the exception of the national boundaries which were aggregated using the median. For datasets that only reported data in non-zero regions, missing data was replaced with zeros for aggregation but the results were tested for robustness to this imputation.

[7]Ramankutty, Foley, Norman, and McSweeney (2002)

[8]http://geosolutions.isciences.com/news/isciences.html

[9]This divides the world up into 2,073,600 grid squares of which approximately 259,200 are land of which 230,000 are included in my sample.

2.3 Empirical Specification

As mentioned above, previous research has concentrated on the country level impact of malaria on economic activity, while controlling for institutions. The problem with this past empirical specification is that the results could suffer from instruments that violate the exclusion restriction, country level omitted variables, and an inability to separate out malaria from other mosquito borne diseases. As a starting point, I use the empirical specifications of Sachs (2003) and AJR, regressing malaria on economic activity, instrumenting for malaria and instrumenting for average expropriation risk with settler mortality. I use these specifications to compare the malaria measure and the instrument I use to what has been done before to see if they are consistent with other measures.

Beyond using the empirical specification from previous research (Sachs 2003 and AJR 2001) I am also able to use country fixed effects to correct for concerns of omitted variables. It may be the case that much of the previous research was picking up omitted health, political, education or cultural characteristics at the country level. By including fixed effects I can look at the within country impact while controlling for national unobservables. Therefore, the main specification of the model is given by:

$$Y_j = \beta_0 + \beta_M * \hat{M}_j + \beta_x * X_j + \beta_c * C_j + \varepsilon_j \tag{2.1}$$

where Y is the log of one plus the average lights from 2004-2009 divided by one plus population, my proxy for economic activity, M is malaria prevalence for 2009, E is malaria ecology, X is a set of spatial control variables, and C is country fixed effects for each quarter degree grid square j. The first stage is given by:

$$M_j = \delta_0 + \delta_E * E_j + \delta_x * X_j + \delta_c * C_j + \nu_j \tag{2.2}$$

where M is instrumented by malaria ecology, E. Given that the lights data is left censored with a significant mass at zero, an IV Tobit model is used for the estimation.

Malaria ecology, as mentioned above, is constructed from purely geographic characteristics and has no economic or population level variables entering into it ameliorating concerns that it might be correlated with local population characteristics. Still concern might remain about malaria ecology being correlated with other geographic characteristics which impact economic activity such as food production or trade. Given how malaria ecology was constructed, it captures just the environmental characteristics favorable for the survival of a particular mosquito type, and those characteristics are extremely non linear and unlikely to be correlated with geographic conditions that may affect GDP per capita. To address any lingering concerns I control for distance to fresh water, distance to the coast, absolute latitude and land suitability and show that the results are not generated through these other possible channels.

Another concern that has not been addressed in the literature is that there are many mosquito borne diseases and therefore the impact of these diseases are correlated. It might be the case that malaria or malaria ecology is simply picking up the impact of another mosquito borne disease. This is unlikely the case because malaria ecology is based on solely those mosquitoes that are malaria carriers and each mosquito's specific breeding and biting pattern. Mosquitoes that carry diseases can be broken down into two genera, *Anophles* and *Aedes*, with *Anophles* only carrying malaria and *Aedes* carrying all other major diseases. Specifically, it is based on only the sub species of the *Anophles* genus that can carry malaria, 40 out of the 460 species in *Anophles*.

To address any remaining concern that malaria is highly correlated with other mosquito borne diseases, I exploit biological differences between the species. There are many biological difference between the major disease carrying genera of mosquitoes. Specifically, I exploit the characteristic that *Aedes* is restricted to elevations below one thousand meters.[10]

2.4 Results

2.4.1 Cross Country

Before running my full specification I repeat the specifications previously used in the literature to see how using my measure of malaria and instrument impact the results. Following AJR I use settler mortality as an instrument for expropriation risk and the log of GDP per capita for dependent variable.

Table 2a presents AJR's specification using data that they provide online and which allows me to exactly reproduce their results in Column 1. Column 2 presents the results using my instrument. I find that malaria becomes extremely significant and slightly larger in magnitude. In Column 3

[10]This is stated in many papers as being between 1000 meters and 1200 meters. A few examples include: WHO (2009) and Delatte, Dehecq, Thiria, Domerg, Paupy, and Fontenille (2008).

I use my measure of malaria and their instruments. I find a slightly more negative coefficient on malaria but, as with their specification, it is not significant. Finally I use both my instrument and malaria measure giving me a more negative coefficient on malaria which is significant at the 1 % level. I find that the main result from their paper, that expropriation risk plays a significant role in economic activity, holds but I also find that malaria plays a significant role which they did not. The difference is mostly due to the poor instruments they use for malaria.

Sachs (2003) does not provide a preferred specification overall but does provide a specification he uses to compare to the work done in AJR, which is what I use. Following his specification I instrument for expropriation risk with settler mortality and instrument for malaria with malaria ecology and the share of a country's population in the temperate zone.

Table 2b presents Sachs' specification using data that he provides online. In Column 1 I was unable to exactly reproduce his results as the data is not the same as what was used in his paper.[11] The results are similar but I get a less significant coefficient on malaria. Again, as was done for AJR, I replace his instrument with mine (Column 2) and find that the coefficient becomes more significant and larger, similar to AJR. Next I use my measure of malaria (Column 3) and find, again similar to AJR, that the coefficient becomes more negative but the significance is not increased over the base specification. Finally when using both my malaria measure and instrument I get a more negative and significant coefficient on malaria (Column 4).

Given the consistency of the results from reproducing previous studies, it seems that the malaria measure provides a stable representation of the malaria levels across multiple samples and that my instrument for malaria ecology seems to be better than instruments previously used. To further test malaria ecology as an instrument, I move away from the cross country regressions, to a within country framework.

2.4.2 Within Country

Instead of using cross country data as in the previous literature, I exploit the finer geographic level of my data and run all future regressions at the grid square level. The dependent variable in my

[11] His exact results had a coefficient on average expropriation risk of 0.45 and a coefficient on malaria of -1.07. The biggest difference being he finds a large coefficient on malaria of -1.07 versus my value of -0.711 and significant at the 5 percent level versus my 10 percent level.

regressions is the log of one plus the average nighttime lights from 2004 to 2009 divided by one plus population. In Table 2.3 Column 1 I run my first specification with malaria ecology instrumenting for malaria, and no controls. In the first stage of this regression I get a large positive coefficient on malaria ecology with an F statistic well above ten.[12] In the second stage the coefficient on malaria is negative and significant at well beyond the 1% level. The coefficient on malaria implies that an area with 10% less malaria would have 29% more lights. Going from the worst malaria area (73% prevalence) to an area with no malaria would increase lights by 211%.

One of the major concerns is that malaria might be be picking up other factors, and therefore violate the exclusion restriction. It might be the case that malaria is correlated with fertile soil, better living conditions or better conditions for trade. To test these concerns, I include Land Suitability, a measure of how good the land is for farming; latitude and distance to fresh water, which are measures of how hospitable a place is to live; and distance to coast, a measure of how easy it is to get to the coast to trade. Columns 2 to 5 show that the inclusion of these controls has little influence on the size or significance of the coefficient on malaria or my instrument.

Overall this does not address one of the major concerns with the previous literature, that unobserved country characteristics might be playing a major role. It could simply be that some unobserved health care or education characteristic are driving the results. To address this issue, I add in country fixed effects in Table 2.3 Column 6 and find that the coefficient on malaria remains stable and significant.

Finally, due to the spatial nature of my data, spatial correlation also may be a concern. In Table 3 Column 7 I cluster at a local level. Specifically I broke the world up into three by three grids and clustered by these grids. The results remain significant at the 1% level.[13] Thus, after solving many of the major issues that have plagued the literature we can see that the impact of malaria on economic activity is not just at the country level but, even after controlling for country level characteristics, malaria still has a large impact on economic activity. I treat Column 7 as the baseline case in all future tables.

[12] MAP provides the standard deviation for the malaria data, therefore weights for each grid square based on the noise of the estimate can be constructed. The results presented are robust to a range of weights using the standard deviation.

[13] Given the large extent of the data, a simple spatial weighting matrix is infeasible to construct, the matrix would have over 60 billion points. To solve this I separate the world into grids of 0.75 degree by 0.75 degree (three observations by three observations) and cluster my results by them.

The results found in Column 7 state that areas with 20 percentage points decrease in malaria would have 80% more lights. To really understand the meaning of these results it is best to convert them in terms of the effect on GDP. To do so, I use the HSW and Hodler and Raschky result that a unit change in the log of lights is associated with approximately a 0.25 change in the log of GDP. Therefore areas with 20 percentage point decrease in malaria would have 20% higher GDP, which constitutes a huge impact for many of the most malarious regions.

2.5 Robustness

2.5.1 Spatial Autocorrelation

As always with spatial data, Tobler's law applies; "Everything is related to everything else, but near things are more related than distant things." Above, I tried to correct for the possibility of spatial correlation by clustering. However, clustering does not account for many different types of spatial relations. Three types of spatial models that might not be captured by clustering are a spatial lag of the dependent variable, a spatial lag in the error or a combination with a lag in the error and independent variable.[14] Theory provides no clear indicator of which model should be applied and therefore statistical tests are used to determine the best fit. Fortunately, to determine which model to use, there are three Lagrange multiplier (LM) tests that allow me to check for a spatial lag, a spatial error or both. Unfortunately the tests and corrections for spatial correlations are not available with an IV tobit but only IV two stage least squares. Therefore, for the rest of this section I will use two stage least squares for all of the regressions and tests.

Using the three tests mentioned above I can determine which model best fits the data and therefore which model to apply. After running a robust spatial lag and robust spatial error test I get that both are significant with a P-value of zero. Therefore, I run a LM test for a combination (SARMA) model and find that this too has a P-Value of 0. Therefore, I run a SARMA model but I also include the results from a spatial lag and spatial error model.

To provide comparison against which to judge the spatial model I use as a baseline IV two stage

[14]The weights matrix used is a rook weights matrix; defined as a one for the North, South, East, and West neighbors and a zero everywhere else. The weights matrix is row standardized and robust to using a queen instead.

least squares with clustering in Table 2.4 Column 1. In Column 2 I run a spatial error model and in Column 3 I run a spatial lag model. I find that the malaria coefficient still remains significant at the 1% level and the same magnitude in both specifications. Finally in Column 4 I run the SARMA model. I find that the coefficient on malaria is the same magnitude and significant at the 1% level. Even though the tests suggests spatial autocorrelation exists in the data, it does not impact the malaria coefficient significantly and therefore the results do not seem to be driven by spatial auto-correlation.

2.5.2 Other Diseases

Another major concern might be that the exclusion restriction does not hold because of other mosquito borne diseases. The instrument is specifically constructed to pick up the ecology of malaria carrying mosquitoes, but mosquitoes can carry a range of other diseases such as Dengue Fever, Yellow Fever, and Arboviral Encephalitides. Mosquitoes of *Aedes* genus carry Dengue Fever, Yellow Fever and Arboviral Encephalitides while the *Anophles* genus carries only malaria. As a general rule, the environment that supports one mosquito supports the other, which might call into question the exclusion restriction. An important exception to this rule is that the genus *Aedes* cannot survive above 1000 meters while *Anophles* can.

This important environmental restriction allows me to test the exclusion restriction by restricting the sample to elevations greater than 1000 meters. Table 2.5 Column 2 shows that, after limiting our sample to over 1000 meters, the results hold and are significant. The coefficient of interest increases which is unexpected if I am removing a positively correlated variable, but it may be due to the exclusion of deserts, as will be discussed next.

2.5.3 Sample Restrictions

So far I have shown that malaria ecology and lights have a strong negative correlation. An exception to this would be in desert regions where the two are positively correlated. In desert regions I would expect that the only place people live are places that have access to fresh water. Similarly the only place I would expect malaria ecology to be positive in deserts is in areas where there is fresh water. Thus the positive effect of the water dominates the negative effect of the malaria confounding the expected correlation.

To determine how much of an impact deserts have, in Table 2.5 Column 3, I drop all the areas with biomes defined as desert. The results are more significant and more negative, as to be expected if the positive correlation between malaria ecology and lights dominates in deserts. The coefficient almost doubles, meaning that the impact malaria plays is even more significant than mentioned above.

Another concern might be that the results are simply picking up the difference between having any malaria or having no malaria. It could simply be the case that having any malaria significantly hurts the economy but once there is any malaria, increases in the prevalence do not matter much. To check if this is driving my results I can limit my sample to only places with malaria. As shown in Column 4, after removing areas with no malaria, the results remain negative and significant but drop in magnitude. [15]

2.5.4 Intensive versus Extensive Margins

The dependent variable used so far in this paper has been the log of one plus lights divided by population plus one, but variation in this can come from two possible sources. The level of light can be estimated at both an intensive and extensive margin. In other words, one can estimate whether an area has any level of light at all as a zero or one variable, or one can estimate the amount of light per capita in a lit area.

One hypothesis is that people avoid living in areas with malaria, which would result in malaria having an impact on the extensive margin. To test this hypothesis, I can estimate the impact of an area having any lights at all using a dummy variable for being lit versus unlit. Column 5 shows that malaria has a large impact on the location of economic activity.

On the other hand, one might hypothesize that malaria is impacting productivity rather than simply impacting the location choices of individuals. Therefore, I can test this hypothesis by restricting the sample to only places with non-zero level of lights. I find that, conditional on an area being lit, malaria still has a negative and significant impact on the level of lights but a smaller impact. (Table 2.5 Column 6) These two results suggest that malaria impacts not only the location

[15] A similar estimation is done keeping only places with positive malaria ecology and the results are similar.

of economic activity, but also the amount of economic activity in a location.

2.6 Population

The goal of this paper is to look at malaria's impact on economic activity and so far the results that have been presented have only considered the impact on lights as an indicator of GDP. Another possible measure of GDP is population density, which is also available on the grid square basis. Using population density allows me to break down economic activity into it's two parts; where people live and how productive they are.

Before using population I think it is important to mention some of the problems associated with the population density measure. First, in areas with malaria the population data is based on large administrative districts with an average size of a 20,000 km^2. Unfortunately these countries present a large fractions of places with malaria. Second to go from administrative districts to grid squares the population in the administrative area is uniformly spread across the grid squares. These first two facts mean that if there is a large district with a major city that city's population would be spread over all the grid squares in that area rather then in just the area with the city. Overall I expect this to bias the standard errors downwards, but it is difficult to exactly determine the impact.Therefore I hesitate to put much weight on the population results.

As mentioned above economic activity can be broken down into two parts, how productive people are and where people live, I start by looking at the latter. Using the log of one plus the population density as the outcome variable I get a negative significant coefficient on malaria which means that malaria is playing a major role in individuals's decision where to live. (Table 2.6 Column 2)

Next to look at the impact on the standard of living I use the log of one plus lights as the dependent variable but add the log of one plus the population as an independent variable. This allows me to control for populations's impact on economic activity and look at malaria's impact beyond this. The results of this are shown in Column 3. The coefficient on malaria decreases significantly in magnitude but remains significant and negative. This points to malaria playing a role in productivity on top of it's impact on population.

One concern is that the coefficient on population in Column 3 is not equal to one, which is what would be expected if population density is correlated with income per capita. Again the coefficient being less then one is probably due to the quality of the population data rather then a true economic relationship. To test how this impacts the results I can use the log of one plus lights per capita as the dependent variable thereby forcing the coefficient on population equal to one. When I do this regression I find that coefficient on malaria is significant and negative but smaller then any other results. The population results seem to point to malaria making a larger impact in where people locate then in the productivity of individuals, but both results are significant and non negligible.

2.7 Conclusion

Economists have long debated the issue of malaria's impact on economic activity. All of the previous literature that has addressed this question has struggled from three major problems. First all previous literature has relied on cross country regressions, and therefore could not properly disentangle the effect of institutions and other country specific characteristics from the effect of malaria. Second, the literature has struggled with finding a good measure of malaria prevalence, with the previous measures being based off of reported prevalence extrapolated to the country level. Third they have been unable to come up with a plausible identification strategy. My identification strategy and the data employed address all of these issues.

First, by exploiting the novel nighttime lights data combined with the data for malaria prevalence and malaria ecology available at a grid square level, I can include country fixed effects in the regression specification. This controls for country specific characteristics and allows for heterogeneity within country.

Second, I use a measure of malaria that is based on actual prevalence in an area rather then reported prevalence. My measure is based off of random or community level blood tests and therefore provide a much better measure of the actual diseases environment.

Third, by using an instrumental variable approach where malaria is instrumented by malaria ecology, I address the issue of endogeneity of malaria. I provide evidence that this instrument satisfies the exclusion restriction. I specifically show that it is robust to other mosquito born diseases by limiting my sample to where only malaria mosquitoes live. I also show that it is robust to including land suitability, a measure of how nice an area is and a host of other controls.

The main conclusion is that malaria plays a role in economic activity, not just at the aggregate country level but also at the grid square level. By the arguments given above we can be confident that these results are not driven by country characteristics, other diseases, or land suitability. I also show that malaria has an impact on the extensive and intensive margin. I find that that malaria is playing a significant role in where people chose to locate, the extensive margin. I also find that malaria is playing a role in how much economic activity is going on in areas where people are living, the intensive margin. On top of this I use lights per capita to show that malaria is impacting standard of living. However as mentioned above there are many problems with the population data and therefore more research is needed to determine the validity of these results.

Table 2.1: Summary Statistics

	Mean	Standard Deviation	Minimum	Maximum
Log(1+(Lights/(Population+1)))	.0973362	.2865848	0	4.133134
Log(Population)	1.321822	1.654588	0	10.76423
Malaria Prevalence	0.0338582	0.1086078	0	0.74496
Malaria Ecology	1.817958	5.163224	0	38.08099
Land Suitability	0.2651857	0.3178787	0	1
Distance to Fresh Water	2.921479	3.355775	0	69.17053
Distance to the Coast	5.147349	5.117196	0	24.75505
Absolute Latitude	39.64882	21.97211	0	83.75
Elevation	596.413	783.1346	-1313.281	6156.419
N	256,921			

Table 2.2: Previous Research

(a) AJR

	Base	My instrument	My malaria measure	My instrument and malaria measure
Average Expropriation Risk	0.689**	0.501***	0.757***	0.560***
	(0.258)	(0.137)	(0.180)	(0.130)
Malaria (malfal94)	-0.623	-0.920***		
	(0.685)	(0.328)		
Malaria (MAP)			-0.886	-1.699***
			(1.010)	(0.637)
N	60	65	58	66
Instruments	Latitude, Mean Temperature, Distance to Coast	Malaria Ecology	Latitude, Mean Temperature, Distance to Coast	Malaria Ecology

(b) Sachs

	Base	My instrument	My malaria measure	My instrument and malaria measure
Average Expropriation Risk	0.586***	0.468***	0.656***	0.557***
	(0.159)	(0.175)	(0.146)	(0.174)
Malaria (malfal94)	-0.711*	-0.961***		
	(0.372)	(0.324)		
Malaria (MAP)			-1.354*	-1.802**
			(0.790)	(0.680)
N	59	60	59	61
Instruments	Malaria Ecology(Sachs)	Malaria Ecology	Malaria Ecology(Sachs)	Malaria Ecology

The dependent variable is the log of GDP adjusted for purchasing power parity in 1995. Malaria is measured in two ways. Malfal94 is the amount of the population living in areas with a risk of malaria. MAP is a measure of the prevalence of malaria among 2-10 year olds. Average expropriation risk is the risk of expropriation of foreign investment by governments. All Columns use 2SLS with malaria and average expropriation risk instrumented. All Columns use settler mortality to instrument for average expropriation risk. In AJR Columns 1 and 3 malaria is instrumented with latitude mean temperature and distance from the coast. In Sachs Columns 1 and 3 malaria is instrumented with malaria ecology and the share of a country's population in temperate ecozones. In Columns 2 and 4 malaria is instrumented with malaria ecology. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.3: Controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Malaria	-1.041***	-1.225***	-1.273***	-1.278***	-1.041***	-1.341***	-1.341***
	(0.0159)	(0.0176)	(0.0178)	(0.0178)	(0.0197)	(0.0582)	(0.109)
Land Suitability		0.584***	0.533***	0.523***	0.548***	0.460***	0.460***
		(0.00427)	(0.00423)	(0.00426)	(0.00438)	(0.00464)	(0.00920)
Distance to Fresh Water			-0.0124***	-0.0155***	-0.0123***	-0.0234***	-0.0234***
			(0.000608)	(0.000623)	(0.000615)	(0.000870)	(0.00164)
Distance to the Coast				-0.0101***	-0.0107***	-0.00323***	-0.00323***
				(0.000247)	(0.000247)	(0.000340)	(0.000726)
Latitude					0.00192***	-0.0111***	-0.0111***
					(0.0000766)	(0.000216)	(0.000433)
Country FE	No	No	No	No	No	Yes	Yes
Clustering	No	No	No	No	No	No	Yes
First Stage							
Malaria Ecology	0.0158***	0.0157***	0.0157***	0.0157***	0.0146***	0.00666***	0.00666***
	(0.0000773)	(0.0000774)	(0.0000775)	(0.0000773)	(0.0000844)	(0.0000707)	(0.000178)
N	230,844	229,537	217,912	217,912	217,912	217,912	217,912

The observations are 0.25 degree squares for the entire world. The dependent variable is log of one plus the average lights for 2004-2009 divided by one plus population. Malaria is a measure of the prevalence of malaria among 2-10 year olds. Land Suitability is a 0 to 1 measure of how suitable land is for cultivating. Latitude is the absolute latitude taken at the top left corner of each grid square. Malaria Ecology is the ability of malaria carrying mosquitoes to survive, reproduce and bite humans. The specifications are estimated using a 2SLS tobit censored at 0. Standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.4: Robustness Tests-Spatial Autocorrelation

	(1) Baseline	(2) Spatial Error	(3) Spatial Lag	(4) Spatial Combo
Malaria	-2.084***	-1.984***	-1.47***	-1.53***
	(0.70)	(0.0526)	(0.0526)	(0.0404)
Land Suitability	0.770***	0.716***	0.556***	0.556***
	(0.091)	(0.007)	(0.006))	(0.007)
Distance to Water	-0.018**	-0.016***	-0.013***	-0.013***
	(0.008)	(0.0008)	(0.0006)	(0.0007)
Distance to the Coast	-0.018***	-0.017***	-0.015***	-0.015***
	(0.0063)	(0.0004)	(0.0004)	(0.0003)
Latitude	-0.010	-0.00897***	-0.0059***	-0.0054***
	(0.006)	(0.0003)	(0.0002)	(0.0002)
Log(avglights+1)[a]			0.44***	0.45***
			(0.007)	(0.0081)
Country FE	Yes	Yes	Yes	Yes
Clustering	Yes	No	No	No
N	226,071	226,071	226,071	226,071

[a] Represents variables that are the average of the Neighbors values. Neighbors being defined as North South East and West.

The observations are 0.25 degree squares for the entire world. The dependent variable is log of one plus the average lights for 2004-2009. Malaria is a measure of the prevalence of malaria among 2-10 year olds. Land Suitability is a 0 to 1 measure of how suitable land is for cultivation. Latitude is the absolute latitude taken at the top left corner of each grid square. Malaria Ecology is the ability of malaria carrying mosquitoes to survive, reproduce and bite humans. The specifications are estimated using a 2SLS. Column 1 is the baseline specification from Table 3 Column 7 done with 2sls instead of an IV tobit. Column 2 adds in a spatial error. Column 3 uses a spatial lag of the dependent variable. Column 4 uses both a spatial lag of the dependent and a spatial error. Standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.5: Robustness Tests-Sample Restrictions

	(1) Baseline	(2) Elevation>1000m	(3) No-Deserts	(4) Lights>0	(5) Lights Dummy	(6) Malaria>0
Malaria	-1.341***	-1.952***	-3.077***	-0.233**	-4.857***	-0.696***
	(0.109)	(0.558)	(0.212)	(0.113)	(0.333)	(0.123)
Land Suitability	0.460***	0.241***	0.483***	-0.00168	2.151***	0.0445***
	(0.00920)	(0.0160)	(0.0101)	(0.00746)	(0.0286)	(0.00670)
Distance to Fresh Water	-0.0234***	-0.00498**	-0.0174***	-0.00344***	-0.0718***	-0.00372***
	(0.00164)	(0.00205)	(0.00195)	(0.00126)	(0.00470)	(0.00104)
Distance to the Coast	-0.00323***	0.00497***	-0.00295***	-0.00298***	-0.0299***	-0.00836***
	(0.000726)	(0.00109)	(0.000819)	(0.000669)	(0.00173)	(0.000890)
Latitude	-0.0111***	-0.0126***	-0.0135***	0.00255***	-0.0313***	0.000374
	(0.000433)	(0.000924)	(0.000468)	(0.000488)	(0.00106)	(0.000870)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Clustering	Yes	Yes	Yes	Yes	Yes	Yes
First Stage						
Malaria Ecology	0.00666***	0.00432***	0.00472***		0.00666***	0.00222***
	(0.000178)	(0.000642)	(0.000204)		(0.000176)	(0.000233)
N	217,912	42,191	176,399	99,417	221,022	38,738

The observations are 0.25 degree squares for the entire world. The dependent variable is log of one plus the average lights for 2004-2009 for Columns 1-4 and 6 and a dummy variable for lit for Column 5. Malaria is a measure of the prevalence of malaria among 2-10 year olds. Land Suitability is a 0 to 1 measure of how suitable land is for cultivating. Latitude is the absolute latitude taken at the top left corner of each grid square. Malaria Ecology is the ability of malaria carrying mosquitoes to survive, reproduce and bite humans. The specifications are estimated using a 2SLS tobit censored at 0, besides in Column 4 which is estimated using 2SLS. Column 1 is the baseline specification from Table 3 Column 7. Column 2 has all grid squares with an elevation less then 1000m removed from the sample. Column 3 has all grid squares that fall under the category of desert removed from the sample. Column 4 has all observations with lights equal to zero removed from the sample. Column 5 uses a dummy for whether a grid square has lights. Column 6 has all observations with malaria equal to zero removed from the sample. Standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.6: Population

	(1) Baseline	(2) Lights	(3) Population	(4) Lights
Malaria	-1.341***	-4.097***	-1.049***	-2.803***
	(0.109)	(0.273)	(0.307)	(0.203)
Log of Population				0.464***
				(0.00357)
Land Suitability	0.460***	1.476***	1.976***	0.543***
	(0.00920)	(0.0208)	(0.0254)	(0.0155)
Distance to Fresh Water	-0.0234***	-0.0696***	-0.0658***	-0.0311***
	(0.00164)	(0.00392)	(0.00356)	(0.00263)
Distance to the Coast	-0.00323***	-0.0322***	-0.0674***	0.00219*
	(0.000726)	(0.00154)	(0.00157)	(0.00121)
Latitude	-0.0111***	-0.0282***	-0.0369***	-0.0126***
	(0.000433)	(0.000951)	(0.00110)	(0.000736)
Country FE	Yes	Yes	Yes	Yes
Clustering	Yes	Yes	Yes	Yes
First Stage				
Malaria Ecology	0.00666***	0.00666***	0.00666***	0.00666***
	(0.000178)	(0.000176)	(0.000178)	(0.000177)
N	217,912	226,076	217,912	217,912

The observations are 0.25 degree squares for the entire world. The dependent variable is log of one plus the average lights for 2004-2009 in Column 1, and 3. The dependent variable in Column 2 is the log of one plus the population for that grid square. The dependent variable in Column 3 is the log of one plus average lights for 2004-2009 divided by one plus the population density. Malaria is a measure of the prevalence of malaria among 2-10 year olds. Land Suitability is a 0 to 1 measure of how suiTable land is for cultivation. Latitude is the absolute latitude taken at the top left corner of each grid square. Malaria Ecology is the ability of malaria carrying mosquitoes to survive, reproduce and bite humans. The specifications are estimated using a 2SLS tobit censored at 0. Column 1 is the baseline specification from Table 3 Column 7. Column 2 uses a dependent variable of the log of population in that grid square to show that malaria impacts other measures of economic activity. Column 3 controls for population to show that not all the impact of malaria on economic activity is through population. Standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Chapter 3

The Impact of Access to Electricity on Education: Evidence from Honduras

3.1 Introduction

Worldwide there are over 1.3 billion people who still lack access to electricity which is seen a major hurdle to economic and human development. This issue was brought to the forefront of policy discussions in 2014 when President Obama pledged $7 billion dollars to build electrical infrastructure in sub-Saharan Africa. The President said electricity is "the lifeline for families to meet their most basic needs and it's the connection needed to plug Africa into the grid of the global economy. You've got to have power."

One of the major benefits many proponents of electricity extoll is increased educational attainment. Christine E. Kimes, Acting Head, World Bank Bangladesh said, "Access to energy promotes economic growth and prosperity and has a positive impact on income, expenditure and education." The European Commission Development Commissioner Piebalgs said, "The benefits of rural electrification are manifold - by connecting people to clean energy, we'll improve healthcare, education, and opportunities to make a living in the area." Even though many suggest that greater access to electricity will improve educational attainment, it is not clear theoretically or empirically that this will be true.

The impact of access to electricity on educational attainment is theoretically unclear as there could be multiple mechanisms at work. One possible mechanism is that access to electricity increases demand for low skilled labor. This would increase the opportunity cost for students to stay in school and would lead to a drop in educational attainment. Another could be that access to electricity brings in manufacturing jobs. This would require more high skilled labor increasing the returns to human capital, which would cause students to be more likely to stay in school. There are a myriad of other possible mechanisms, which makes the impact of electricity on educational attainment uncertain.

There is also not a consensus in the empirical literature on the impact of access to electricity on educational attainment. While some papers do find a positive effect, many find no effect. Barron and Torero (2014) and Khandker, Samad, Ali, and Barnes (2012) find an increase in hours spent studying, but Bensch, Kluve, and Peters (2011) finds no effect. As for the impact on enrollment, Barron and Torero (2014) finds no effect while Khandker, Barnes, and Samad (2013), Khandker, Samad, Ali, and Barnes (2012) and Khandker, Barnes, and Samad (2009) find increases in both enrollment and attainment.[1]

In this paper, I study the electrical expansion in Honduras from 1992 to 2005 in order to estimate the impact of access to electricity on educational attainment. This expansion was instigated by a major energy crisis in 1993, which lead to widespread reform of the electrical industry in 1994. This reform forced the electrical company to expand coverage of electricity to as many people as quickly as possible. I exploit this increase in access to electricity, which occurred at different points in time for each municipality, to estimate the impact of electricity on education.

To measure electricity's impact on education I focus on three measures: attendance, attainment and the hazard of dropping out. Looking first at attendance I find that access to electricity reduced attendance by 4 percentage points. Next, looking at educational attainment I find that the number

[1]The research on electricity's impact on health and income also has conflicting results (Lipscomb, Mobarak, and Barham, 2013; Barron and Torero, 2014). Research on household income is also varied (Barron and Torero, 2014; Chakravorty, Pelli, and Marchand, 2014; Lipscomb, Mobarak, and Barham, 2013; Khandker, Barnes, and Samad, 2013; Bensch, Kluve, and Peters, 2011). The labor market outcomes that researchers have focused on are employment, hours worked, and wages (Lipscomb, Mobarak, and Barham, 2013; Dinkelman, 2011; Khandker, Samad, Ali, and Barnes, 2012; Grogan, 2013; Khandker, Samad, Ali, and Barnes, 2012; Khandker, Barnes, and Samad, 2013).

of years exposed to electricity as a child reduced completed years of schooling by approximately 0.1 years. Finally, by using a hazard approach I can allow for heterogeneity in the dropout rate by year of schooling. I find that access to electricity increased the hazard of dropping out in the first few years of school while actually decreasing the hazard of dropping out in the last few years. Allowing for this heterogeneity by year of education, the number of students in school decreased by 6.5 percentage points with access to electricity while the number of students who finish schooling dropped by 4.5 percentage points. Using each of these measures I find a similar impact of electricity but from the hazard analysis I am able to find heterogeneous results suggesting there may be multiple mechanisms at work. I focus on three possible mechanisms.

The first mechanism I examine is childhood participation in the labor market to see if access to electricity is reducing education through increased employment opportunities for children. I find that access to electricity is associated with a 2.4 percentage point increase in employment for children, which is a similar magnitude to my estimates of the effect of electricity on the number of students dropping out.

The second mechanism I look at is electricity's impact on adult labor markets. I show that with access to electricity there is an increase in employment for women. If more women are being drawn into the labor market, children may need to stay home to take care of younger siblings. I find that having a sibling under the age of 5 reduces the probability that children continue in school. Suggesting that the increased demand for children to stay home may be driving some of the results.

The third mechanism I analyze is access to electricity impact on the demand for human capital. If the demand for completed education is increasing there would be larger returns for individuals near the end of the schooling to stay in school. This might explain why I find a decrease in the hazard of dropping out at the end of individuals schooling. While I do see a return to human capital before access to electricity, I do not find that this return changed with access to electricity.

Finally, to reduce concerns of endogeneity in the year of access to electricity, I introduce a new instrument, distance along the distribution grid from a municipality to the nearest substation. This instrument is an improvement upon previous instruments because I use information on the location of the distribution grid which others have not had. First, I show that distance to the grid is highly correlated with the year municipalities gain access to electricity. Furthermore, distance along the distribution grid is also uncorrelated with population growth, levels or other labor market variables. Finally, using this new instrument I find that my results hold and remain significant.

The rest of the paper is organized as follows: Section 2 provides background information on the education system, electrical grid and electrical company in Honduras. Section 3 describes the data used for the analysis. Section 4 provides information on municipalities gaining access to electricity and the impact it has on education. Section 5 explores the mechanisms at work. Section 6 provides robustness analysis with an instrument for the year of electricity and information on migration patterns. Section 7 concludes.

3.2 Electrification and Education Background

The electrical network in Honduras is monopolized by the government owned company Empresa Nacional de Energia Electrica (ENEE). ENEE heavily subsidizes electricity to households, such that the main obstacle to having electricity is access to the grid. In 1993, a drought and lack of generation capacity brought about a huge energy crisis in Honduras, causing massive blackouts, and economic and political turmoil. In response to this catastrophe, the government passed a law (Ley marco del subsector eléctrico) in 1994 to create more competition, increase generating capacity, and increase the electrification rate throughout the country. The law succeeded in the last two goals, but the ENEE remains the sole distributor of electricity. The law created the Fondo Social de Desarrollo Eléctrico (FOSODE) which was responsible for increasing electrical access in rural and poor communities. FOSODE along with the ENEE drastically increased access to electricity across a broad swath of the country. The fraction of the population with electricity went from 40% in 1992 to 67% in 2005 (the end of the sample for this paper) and 83% in 2011.

The education system in Honduras is provided by the government and organized as primary, lower secondary and upper secondary. The average years of schooling is 3.57 years in my sample which is lower but similar to the 4.5 years the government found in their 2001 census. Primary school is from ages six to twelve or grades one to six. In my data secondary schooling is almost never observed, less than one percent, therefore I limit the analysis to just primary schooling. The reason that secondary schooling attendance is so low in Honduras is that in many places there is simply no access to secondary schools. Beyond this many secondary schools also charge therefore providing even larger barriers to attendance. While few students go beyond 6 years of education, initial attendance is extremely high, over 95%. On top of this, average attendance in my sample is approximately 75% with around 33% getting 6 years of education.

3.3 Data

3.3.1 Household Survey

The individual level data used in this paper is from the Encuesta Permanente de Hogares de Propósitos Múltiples (EPHPM) which is produced by the Instituto Nacional de Estadística (INE) in Honduras. The main objective of the survey is to collect data on labor, housing, education and household composition.[2] On average there are 40,000 individuals per survey with a total of 905,000 individuals over the entire study period from 1992 to 2005. Of the approximately 300 municipalities in Honduras, the survey covered an average of 180 municipalities each year. The municipalities fall into three general categories: municipalities surveyed throughout the entire study period, municipalities surveyed in 1992, 1993, and 2001 onward, or municipalities surveyed from only 2001 onward. The survey has been conducted since 1990 with one or two surveys conducted each year.[3] Beginning in 1992, the survey asked how a household gets its electricity, which will allow me to determine when each household received access to the electrical grid.

The survey data was not originally designed for matching municipalities across years but rather for creating summary statistics for the national government. Because of this, the surveys have neither consistent variable names (i.e. p4 from one year is p8 in another) nor consistent coding across months and years. To create the dataset used in this paper, I matched all variables across time based on the question asked and recoded all variables for consistency.[4] This unique dataset allows me to follow municipalities across time.

3.3.2 Electrical Grid and Geographic Characteristics

Electrical grid data for Honduras was provided by the ENEE. The ENEE provided an ArcGIS file of the electrical grid for 2012 containing all distribution lines, transmission lines, and substations

[2] INE website: http://www.ine.gob.hn/index.php/censos-y-encuestas/encuestas-todos-las-encuestas-de-honduras/encuesta-permananente-de-hogares

[3] Survey data is unavailable for 1994 and 2000.

[4] The categories were combined to the smallest number asked and top coding was removed.

(Figure 3.1). They also provided a PDF of the map of transmission lines and substations for 1997 (Figure 3.2).

The land gradient variable was constructed from the 90-meter Shuttle Radar Topography Mission Global Digital Elevation Model.[5] The average land gradient was calculated for each municipality using GIS software to calculate the gradient (defined in degrees from 0 to 90).

3.4 Analysis

3.4.1 Access to the Electrical Grid

The main variable of interest in this paper is whether or not a municipality has access to the electrical grid in a given year. Unfortunately, there is no official record of when a municipality was added to the grid. Therefore, to determine the timing of access, I create a binary variable for each municipality which is based on individual reports of having access to the electrical grid. This variable is defined as one after a municipality gains access and zero before. Once a municipality gains access to electricity they have access for the rest of the sample, regardless of reporting. Figure 3.3 shows the percentage of the population reporting access to the electrical grid in municipalities before and after I define them as having access to electricity. This graph shows that when a municipality gets hooked up to the electrical grid a significant fraction of the population gets connected.

To generate municipality level access to the electric grid, I must account for individual reporting errors. For example, in one municipality one household reports having electricity while no other household that year or for the next 3 years report having access to electricity. To correct for these errors I can choose a minimum percentage of the population and a minimum number of individuals that are required for a municipality to be defined as having access to electricity.[6] To determine the appropriate cutoffs, I compare the cutoff's effect on the increase in the percentage of the population reporting access to electricity from before I define an area as having electricity to after. Figure 3.4 shows different percent cutoffs versus the increase in the percentage of the population reporting access to electricity associated with each cutoff. This graph shows that a cutoff of 5 percent maximizes

[5]http://landcover.org/data/

[6]Both percentage and number of individuals are used for a cutoff because the size of municipalities is extremely heterogeneous. With a percent cutoff large districts might under report electricity while with an amount of people small districts might under report.

the difference between the before and after periods. I calculated the difference between before and after for the full set of percent and number of people cutoffs and the values that maximizes the difference are 5 percent and 16 people, as seen in Figure 3.5.

3.4.2 Education

As mentioned before, the impact of access to electricity on education is not clear. On one hand, many studies argue that electricity can increase educational attainment by reducing the amount of manual labor needed in the home or by extending the number of daylight hours allowing for more time to study. On the other hand, electricity can reduce educational attainment by increasing the opportunities of students in the labor market, raising the opportunity cost of staying in school. To determine what effect is dominating in Honduras during my time period I look at three outcomes.

First, I estimate the impact of access to electricity on school attendance. I restrict the sample to children between the ages of 7 and 16, because most students have started schooling by age 7 and more than two thirds of the students are no longer in school after age 16. My specification is as follows:

$$attend_{i,m,y} = \delta_0 + \delta_e ect_{m,y} + \delta x X_{i,m,y} + \gamma_m + \sigma_y + \nu_{i,m,y} \tag{3.1}$$

where *attend* is a binary variable for if students are in school, *ect* is an indicator variable equal to one if that municipality had electricity during that year and zero otherwise, X is a vector of individual characteristics containing sex, age, gender of the household head, and total number of individuals in the household, γ_m is municipality fixed effects, and σ_y is year fixed effects. Table 3.1 presents the impact electricity had on whether or not children were attending school. I find that electricity had a negative and significant impact on school attendance overall and for males, and no significant effect for females. The decrease in male attendance due to electricity is 4.3 percentage points which is not a small drop in an area that has 75% attendance on average.

Next, I look at the relationship between total years of schooling for individuals and the total number of years with access to electricity during their childhood. In order to get an accurate picture of total years of schooling, I restrict my sample to only those people who were under 24 years of age. The reason for this is that no one over the age of 24 had electricity during their schooling so they are not a good comparison group. Finally, since I am interested in the cumulative impact of electricity, I estimate the impact of the total number of years of electricity during childhood on total years of schooling. The final specification is the following:

$$eduyr_{i,m,y} = \delta_0 + \delta_e ect_{m,y} + \delta_c ectedu_{i,m,y} + \delta x X_{i,m,y} + \gamma_m + \sigma_y + \nu_{i,m,y} \tag{3.2}$$

where *eduyr* is the number of years of schooling, *ectedu* is the number of years that the person had electricity between the ages of 6 and 12 and all other variables are defined as above. The reason I restrict the sample to ages 6 to 12 is because most of the population starts school at 6 and almost the entire sample gets at most 6 years of schooling. The results of this regression are presented in Table 3.2. The results show that if an individual had electricity during their entire childhood (6 years) they would have gotten on average a half less year of schooling.

So far, I have assumed that the impact of access to electricity is constant across years of schooling. However, electricity might impact students at the beginning of their education more than students at the end of their education, because students at the beginning of their education have a higher cost to finishing their schooling. Alternatively, it may be the case that access to electricity has a larger effect at the end of an individual's education, when the returns to labor might be higher. To determine whether there were heterogeneous impacts of electricity by year of schooling, I estimate a Cox Proportional Hazard model:

$$h(y, x, ect) = h_0(y) exp(\delta_p ect_{m,y} * YearofSchooling_i + \delta_x X_i + \gamma_m + \sigma_y) \tag{3.3}$$

where $h(y, x, ect)$ is an indicator variable equal to one in the year that the student drops out of school and zero before, $h_0(y)$ is the hazard rate for dropping out of school for each year of schooling, X is a vector of individual characteristics containing sex, gender of the household head, and total number of individuals in the household, *ect* is an indicator variable equal to one if that municipality had electricity during that year and zero otherwise, and $YearofSchooling$ is an indicator equal to one if the student is in a specific year of school (1-5). The results are presented in Table 3.3. The coefficients in the earlier years is significant and positive while the coefficients in the later years become negative and significant. This suggests that electricity initially increases the hazard of students to drop out of school but this effect is mitigated in the later years. Directly interpreting the coefficient is extremely difficult due to the non-linearity in the specification but graphical representations make the impact clear.

In Table 3.4 I present two graphical representations of the coefficients from the analysis above. The first graph is the change in the hazard of dropping out if a location in the sample were to go from having access to electricity to not having access to electricity. These hazards are the predicted value from the Cox regression above, with the values being for the average municipality. By using this hazard, I can calculate the probability that a student would make it to each year. The second graph shows that as students progresses in school the probability they make it to the next year of schooling is lower with electricity. Specifically, with electricity there is a 4.5 percentage point drop in the number of students who finish their education. This would also translate into a 6.5 percentage point decrease in the number of students in school after a location gains electricity.

3.5 Mechanisms

3.5.1 Childhood Labor

I have so far shown that gaining access to electricity reduced educational attainment, but I have not yet presented evidence for the mechanism through which this might occur. It seems likely that a change in employment opportunities may have increased students' hazard of dropping out. Atkin (2012), Duryea and Arends-Kuenning (2003), Kruger (2007), and Shah and Steinberg (2013) present evidence that increases in labor market opportunities reduce childhood educational attainment in a wide variety of settings. Honduras in the 1990s struggled with high rates of child labor(Cruz, 2002). These rates are similar to the rates today in many developing countries. Specifically, in Honduras, melon farming is a major industry and a major employer of children. On melon farms "as each fruit nears its maturation workers turn it over a total of three times to guarantee equal exposure in the sun and the appearance that is pleasing to consumers."(Harwood and Mull, 2002, p. 20) This is an easy task for young children who comprise 80 % of the workforce on melon farms. Another common practice in Honduras is to pay workers either for a set quota or by piece. In this situation, it is common for workers to bring their entire family to the worksite to increase the amount collected or to reach the quota in a shorter time. Exacerbating this tradeoff between school and work is that many rural families "reject the importance of their children acquiring a basic education."(Harwood and Mull, 2002, p. 17)

To test if child labor is a driving mechanism behind the increase in drop out rates I compare childhood employment rates before and after electricity.[7] Table 3.5 shows a significant increase in overall employment and a significant increase in female employment. For male employment, I find a positive effect similar in magnitude but insignificant. These results suggests that there was an increase in child labor market opportunities following electrification and, consequently, an increase in the hazard for dropping out of school.

Finally, the hazard model from above gives a prediction of the percent of students who would drop out with and without access to electricity. Comparing the overall increase in employment, 2.4 percentage points, I find that this is similar in magnitude but smaller than the drop out rate from above suggesting there may be other mechanisms at work.

3.5.2 Adult Employment

For adult labor markets, one of the accepted outcomes in the electricity literature is that with access to electricity there is an increase in female labor force participation. [8] I find that access to electricity increased female employment, hours worked and earnings, as shown in Table 3.6. Increases in female labor force participation would lead to changes in the household dynamics, because household tasks may be redistributed. One of the ways this might impact educational attainment is if a mother goes off to work and requires a child to stay home to take care of a younger sibling. Therefore, I estimate the impact of electricity interacted with the number of toddlers, children under 5, in the household. In Table 3.7, I find that with access to electricity the number of toddlers in a household increases the probability of dropping out of school.

3.5.3 Human Capital

The baseline hazard specification demonstrates that access to electricity raised the hazard of dropping out at the beginning of an individual's education but decreases it at the end. So far I have shown evidence for the increase in the hazard of dropping out but not for the decrease. One of the possible explanations for the decrease is that there is an increase in the returns to finishing school. To determine if this is the case, I can look at the returns to education. In Table 3.8 I show

[7] I restrict the sample to children between the age of 10-15 because children under the age of 10 are not asked labor market questions in many of the surveys and children over 15 are mostly not in school.
[8] Dinkelman (2011); Grogan (2013); Lipscomb, Mobarak, and Barham (2013)

that as schooling increases there is an increase in the wages paid to workers. Next, I can examine whether the returns increased after a municipality received access to electricity. In Table 3.9, I find that there was not a statistically significant increase in the returns to education for individuals who gained access to electricity. This is unsurprising given that many of the jobs that enter into communities when they gained access to electricity are low skilled manufacturing or agricultural jobs.

3.6 Robustness

3.6.1 Instrumental Variable Analysis

If electricity is endogenously allocated or if the timing of this allocation is endogenous, then the previous estimates will be biased. In this paper, I compare within-municipality changes over time, so endogenous allocation of electricity to certain municipalities will not bias my results. Unfortunately, the timing of access to electricity may still be endogenous. The timing of electrical access in municipalities may be correlated with other public works projects, a powerful politician coming into office, or another major unobservable event. Therefore, I instrument for the year each municipality gained access to electricity with the distance along the electrical network.

Electrical networks provide electricity from power plants to consumers through a four step process. First, power is produced at the power plant. Second, it is transmitted along transmission lines to substations. Transmission lines are extremely high voltage lines and there is no benefit to being near them. Third, at the substation, the voltage is lowered to a level that can be distributed to consumers. Finally, it is distributed along distribution lines from substations to the consumer.[9]

The instrument I propose is distance along the distribution network from the centroid of a municipality to the nearest substation, which I refer to as the "grid distance." The grid distance is calculated for each municipality by first finding the nearest point on the distribution line to the centroid of the municipality. Next, I calculate the distance from the point on the distribution line to the nearest substation, along the grid. An example of this is presented in Figure 3.6. The location of the substations and distribution lines are based on the 2012 electricity distribution network. One

[9]The process is actually much more complicated than this with multiple transmission line voltages, substations at intersections of transmission lines and pole transformers to connect to houses. The multiple transmission lines and substations on these transmission lines are ignored for the analysis in this paper. Even though pole transformers are needed to attach to a distribution grid these are assumed to be costless with the majority of the cost of hooking up to electricity being access to a distribution line.

concern with using the 2012 network is that substations might be endogenously placed during the sample. Therefore, I use a map of the network in 1997 to remove any substations or transmission lines built after 1997.[10]

Since the major push to increase access to electricity came with the passage of the 1994 law, mentioned above, the best way to create the instrument would be to use the distance from the end of the grid in 1994 to the centroid of the municipality. Unfortunately, because of the 1994 law, the ENEE stopped keeping records of where their distribution grid was and there are no maps that exist on the location of the distribution network prior to 2012. Therefore by creating the instrument in this way I am making two assumptions. First I am assuming that places that are farther along a distribution line in 2012 were farther from the network in 1994. Second I am assuming that places were hooked up in the order of their distance from the network in 1994.

Even though I cannot prove either of these there are three pieces of evidence that seem to support these assumptions. First, the law that was passed in 1994 created FOSODE which along with the ENEE were the major drivers of increasing access to electricity. FOSODE was created to give access to poor and rural areas while ENEE was left to increase access to everyone else. Given the information available it seems their incentives would be to hook up places as quickly as possible therefore targeting the closest places first. Second, since distance along the grid is uncorrelated with other outcome variables and population, there is no evidence that access was being targeted at populous or rich regions. Finally, since my instrument strongly predicts the year of electrification it is most likely the case that places that are farther along the line in 2012 were farther from the network in 1994. While none of these proves the assumptions above it does provide strong evidence that they are valid.

Previous research has used a range of instruments for access to electricity but they broadly fall into two categories: geographic based and network based.[11] The most common geographic measure is based on slope or gradient while the most common network characteristic is based on distance to a substation. I create these two other instruments to compare to my instrument. The average slope measure was created following methodology in Dinkelman (2011). The distance to the nearest substation was calculated using straight-line distance for all municipalities, even those without

[10] By doing so, only two substation and one transmission line were removed. Although data from 1992 is not available, the small number of substations and lines built after 1997 suggests that few if any substations were built between 1992 and 1997.

[11] Dinkelman (2011); van de Walle, Ravallion, Mendiratta, and Koolwal (2013); Chakravorty, Pelli, and Marchand (2014); Grogan (2012) andKhandker, Barnes, and Samad (2009)

distribution lines, since that is most common in the literature.

To see whether there are any obvious violations of the instrumental variables approach, I first test whether or not the instrument is a good predictor of the endogenous variable, the year of electrification. In Table 3.10, I regress all three instruments, grid distance, straight distance, and slope, against the year of electrification. Both the grid distance and the substation distance are significant, while the slope variable is insignificant.[12] This is unsurprising since slope was used to estimate which places would gain access to electricity rather than when an area gains access to electricity. Because of this, I concentrate on the other two measures.

The second test compares each instrument to the outcome variables before electricity. If the instrument satisfies the exclusion restriction we would expect it to impact the year a municipality gets electricity, but not the outcome variables before electrification. Table 3.11 shows that the grid distance does not significantly impact the outcome variables before electrification, but substation distance does significantly impact many of these variables.

The final test checks whether or not the instrument is simply picking up population density or population growth rates. I regress the instruments against various measures of population and population growth, as shown in Table 3.12. For all four measures, the impact on grid distance is insignificant, while the impact on substation distance is marginally significant for the population growth rate.

Overall these tests suggest that grid distance has a strong first-stage impact on electrification and there is no evidence it fails the exclusion restriction. Also, the evidence suggests that the other two instruments either are not relevant or may fail the exclusion restriction. Therefore, I will use grid distance to instrument for the year of electrification.

My estimation strategy is as follows, I instrument for the year of electrification, $eyear$, with the grid distance, $GridDistance$, while controlling for municipality characteristics, X, average age, percent female, average size of households and percent of households with female heads. The first

[12] This is robust to estimation using the average slope and the average slope along the path of the electrical line.

stage of my analysis is:

$$eyear_m = \beta_0 + \beta_c GridDistance_m + \delta x X_m + \varepsilon_m \tag{3.4}$$

Next, I predict the year of electrification and generate a variable, $Elect\hat{r}icity$, that is equal to one if a municipality is past the year of electrification and zero if it is before. I use this predicted value in the second stage, as follows:

$$h(y, x, pect) = h_0(y, \alpha) exp(\delta_p Elect\hat{r}icity_{m,y} + \delta_x X_i + \gamma_m + \sigma_y) \tag{3.5}$$

where the variables are defined as in Section 4.2. The entire process is bootstrapped and the results are robust to bootstrapping each stage.[13]

As done above, I interact electricity with year of schooling. The results, presented in Table 3.13, show that access to electricity had a positive effect for the first few years of education and decreases to a negative value for later years. Similar to the original estimates but larger. For communities gaining electricity this would reduce the probability of completing schooling by 5 percentage points. Estimating the total effect on attendance, as done above, I find that access to electricity is causing a 15 percentage point drop in the number of students in school. These results are larger but similar in magnitude and direction to the uninstrumented case.

3.6.2 Migration

One concern with the results presented so far is that there may be differential migration before and after electricity. If families with kids that work are more likely not to move after gaining access to electricity there might be a decrease in education and an increase in childhood labor simply due to migration. To see if there is differential migration I can look at cohorts before and after gaining access to electricity. I generate a dependent variable which is the fraction of individuals in a five age range divided by the number of individuals that were in that cohort 2 years ago. I then compare places that got access to electricity in that time period to places that did not. From the graph of the coefficient on electricity for each regression, Figure 3.7, there is no significant pattern of migration

[13]Since this is not the standard two stage set up I am basing the theory on previous literature which has shown that Hazard models are consistent when using a two stage analysis. Beyond this since I am using a generated regressor I simulate the estimation method to show that in the case of an endogenous variable at the municipality level my estimation strategy produces consistent results. Theses simulations are being done now.

during the childhood years.

3.7 Conclusion

With all of the recent investment in electrical grid infrastructure in the developing world, it is important to understand how access to electricity impacts households. I provide substantial evidence suggesting that access to electricity decreases educational attainment. Furthermore, I show evidence that the impact of electricity on attendance is heterogeneous across years of schooling. Specifically, the probability of dropping out of school increases for the first few years of schooling, but decreases for the last few years of schooling. The results of which is a 6.5 percentage point decrease in school attendance and a 4.5 percentage point decrease in students completing their schooling. As a possible mechanism for the decrease in attendance, I show that access to electricity increases the employment of children. Moreover, the magnitude of the increase in childhood employment is similar to the magnitude of the decrease in educational attendance. I provide evidence that another mechanism driving the drop in attendance is an increased draw for students to stay home. This is created by increased job opportunities for adults. Next I show that in the context of Honduras there does not seem to be increases in the returns to education with electricity. Finally, to remove concerns about the endogeneity in the year of electrification, I instrument for the year and find the results hold and are significant.

Even though I find that access to electricity decreases educational attainment this is in no way a reason to decrease or stop programs to increase access to electricity. The key take away is that not all effects of electricity are positive. Therefore policy makers should be aware of the perverse incentives electricity may cause and target programs to try to reverse these incentives.

Figure 3.1: Honduran Electrical Network in 2012

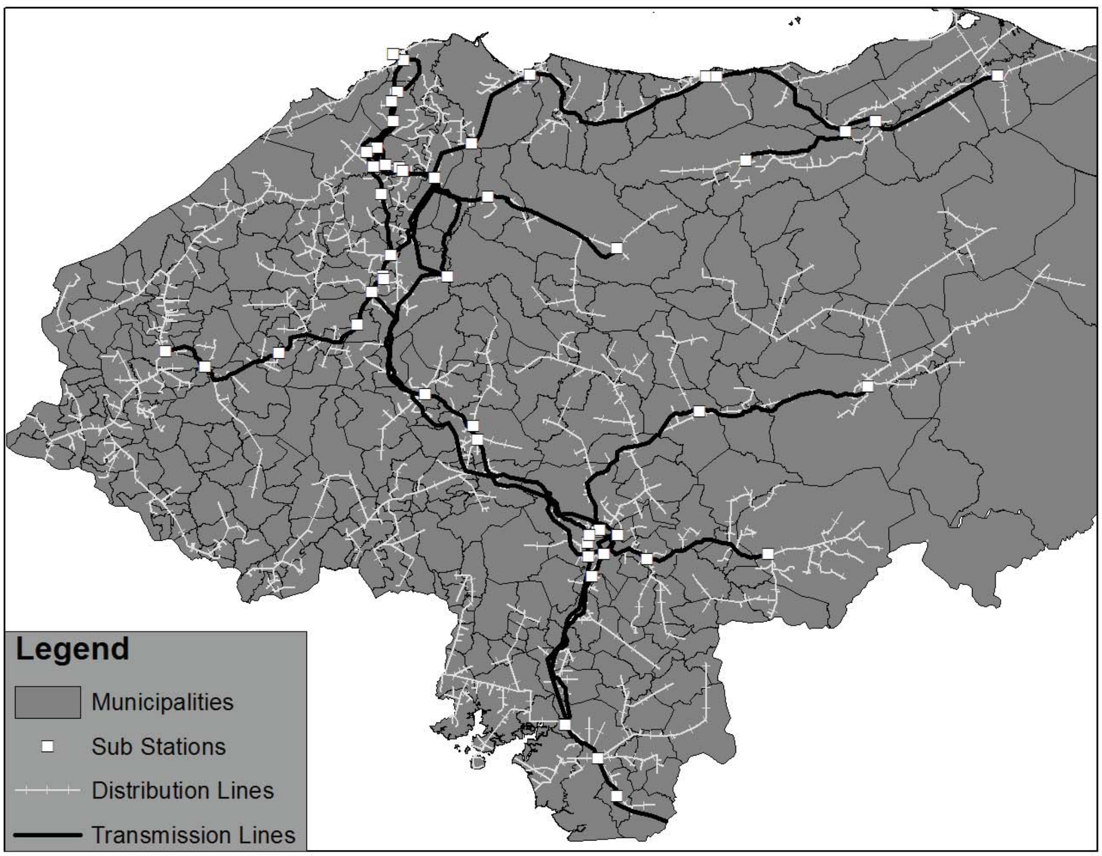

This map represents the electrical grid as it was in 2012.

Figure 3.2: Honduran Electrical Network in 1997

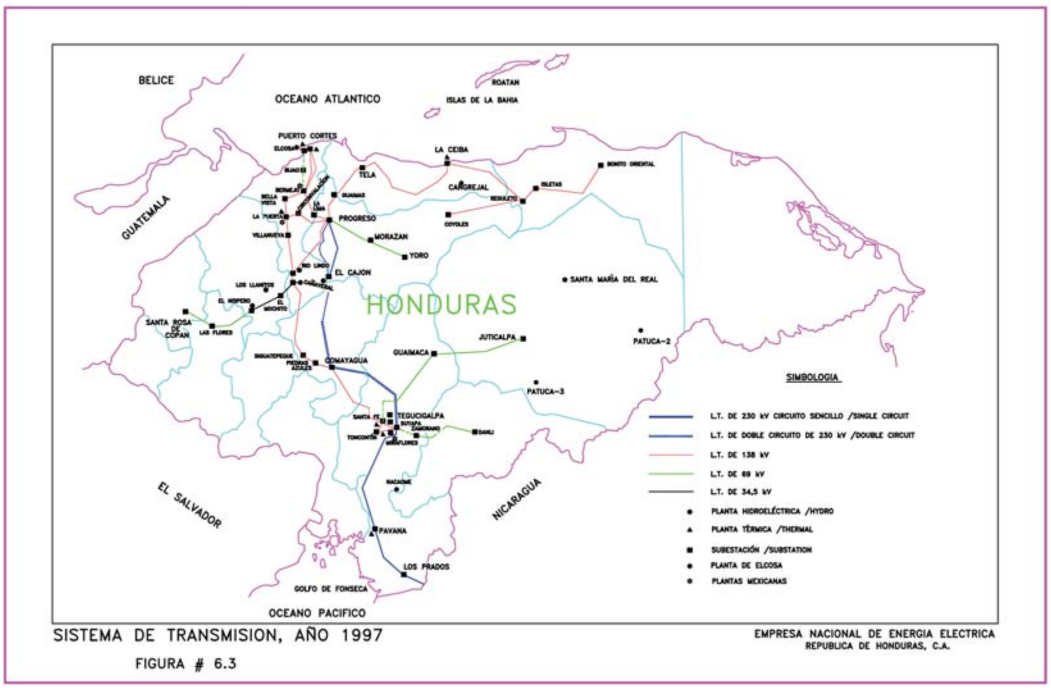

This map represents the electrical grid as it was in 1997. The purple line outlines the country while the light blue line outlines the districts. The thick blue, green and orange lines represent transmission lines, substations are represented by squares and unfortunately distribution lines are not shown. This map was provided by the ENEE in pdf format and is the only other network map available.

Figure 3.3: Percentage of the Population Reporting Electricity

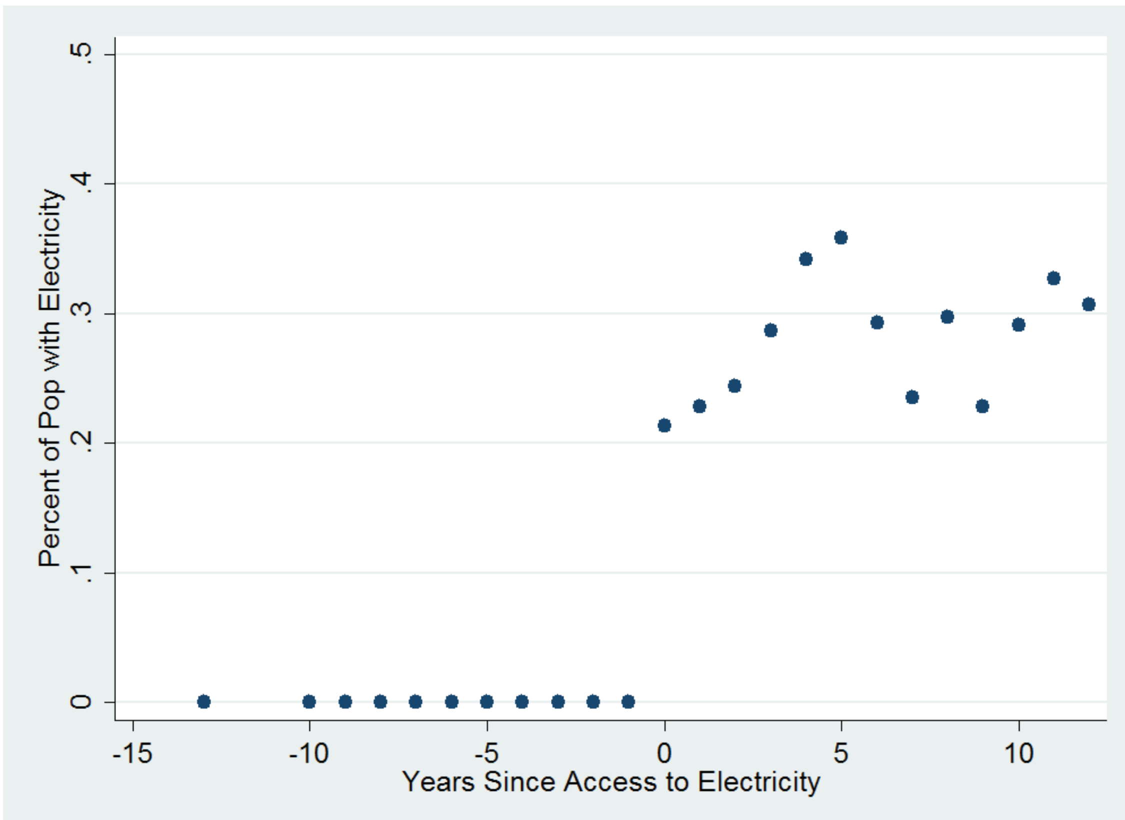

Each point represent the percentage of the population who has access to a electricity. The year of electricity for each individual is defined as the year that the municipality the individual is in had at least one person reporting access to electricity.

Figure 3.4: Percentage Change in Access

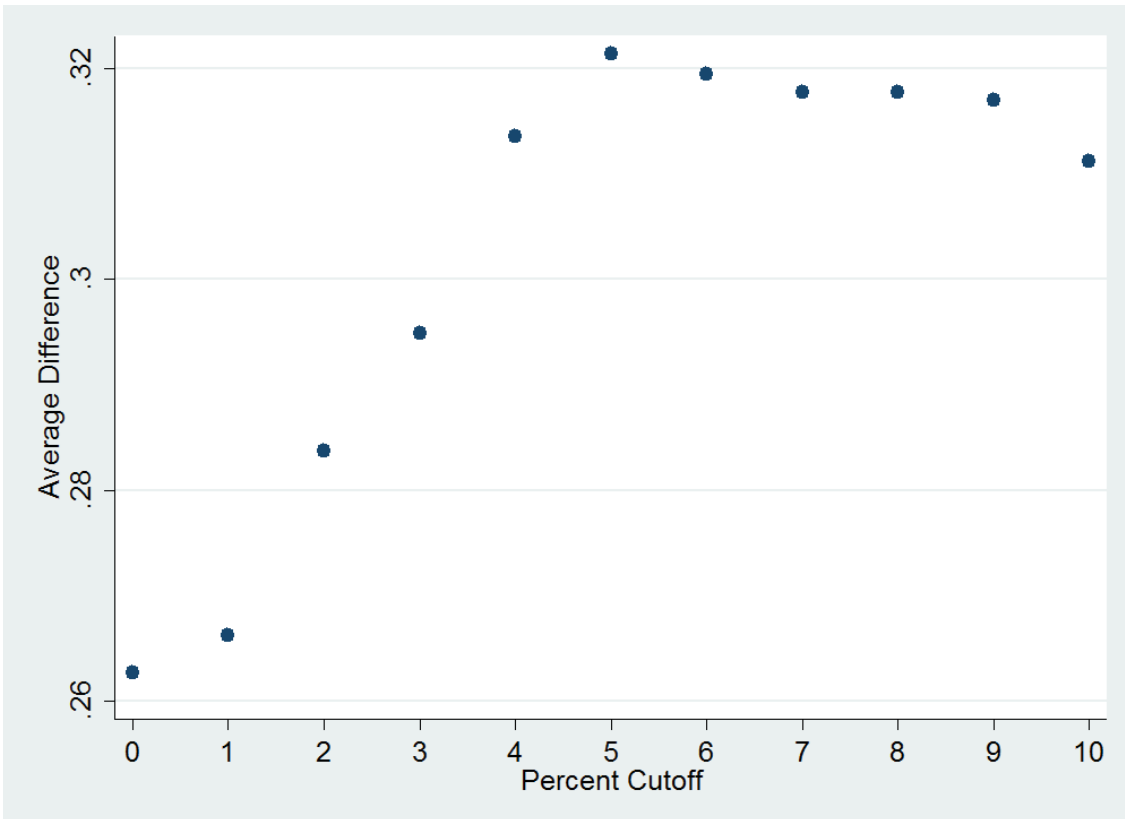

Each point represent the jump in percentage of the population reporting access to electricity given a specific cutoff. The x axis represents the cutoff used to define when a municipality gained access to electricity.

Figure 3.5: Percentage of the Population Reporting Electricity

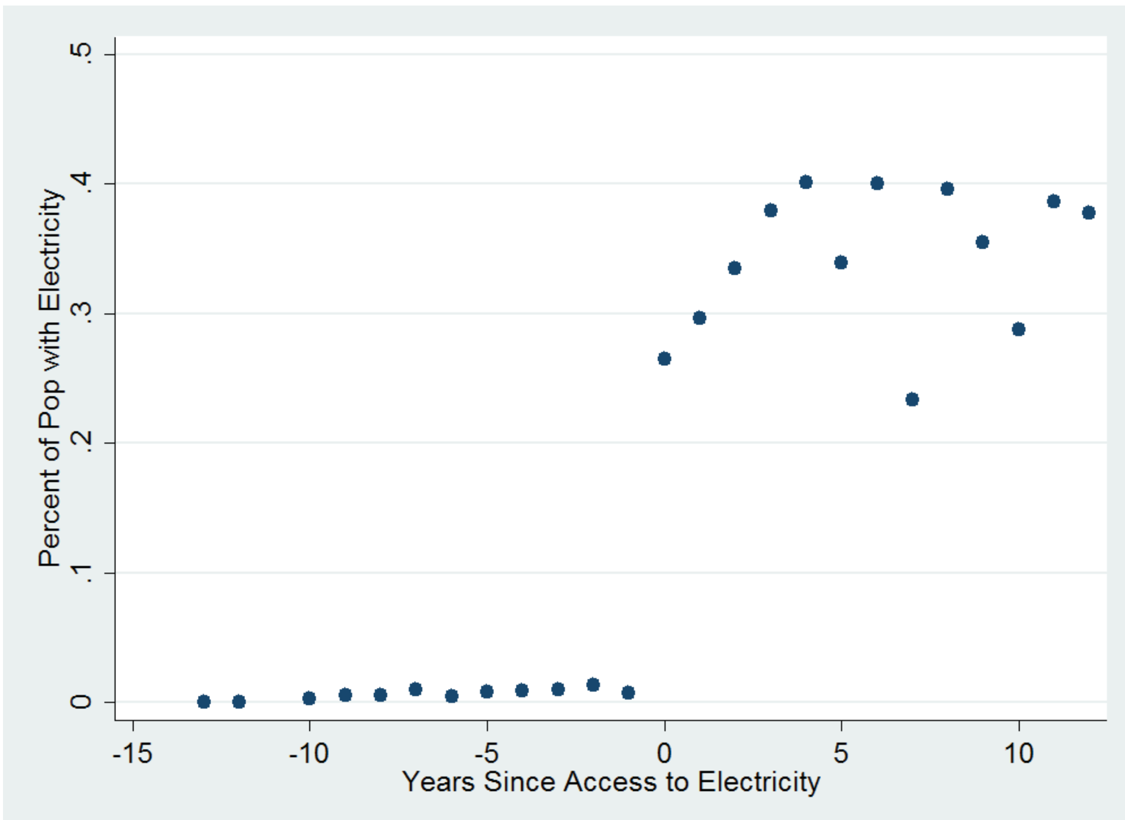

Each point represent the percentage of the population who has access to a electricity. The year of electricity for each individual is defined as the year that the municipality the individual is in had at least five percent or 16 people reporting having access to electricity.

Table 3.1: Educational Attendance

Outcome: School Attendance						
	Full		Women		Men	
Electricity	-0.0386*	-0.0400*	-0.0372	-0.0379	-0.0415*	-0.0428*
	(0.0204)	(0.0203)	(0.0235)	(0.0232)	(0.0222)	(0.0222)
Household Size		-0.00436**		-0.00502**		-0.00370
		(0.00188)		(0.00243)		(0.00236)
Age		0.0127***		0.0110***		0.0148***
		(0.00254)		(0.00357)		(0.00279)
Female Head of Household		0.00720		0.00739		0.00642
		(0.00832)		(0.00987)		(0.0113)
Sex		0.0112*				
		(0.00654)				
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes	Yes	Yes
N	18359	18312	9018	9001	9341	9311

The dependent variable is one if an individual says they are in school and zero otherwise. Electricity is a binary variable with one being if the municipality has electricity and zero otherwise. The sample is limited to children between the age of 7 and 16. The first 2 columns include the entire sample, columns 3-4 limit the sample to only females and columns 5-6 limit the sample to only males. All errors are clustered at the municipality level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.2: Educational Attainment: Years of Education

Outcome: Years of Education			
	Full	Women	Men
Electricity	-0.262*	-0.347**	-0.168
	(0.144)	(0.169)	(0.163)
Years of Childhood Electricity	-0.115***	-0.116***	-0.106***
	(0.0299)	(0.0384)	(0.0334)
Household Size	-0.00860	0.0211	-0.0373**
	(0.0107)	(0.0135)	(0.0145)
Age	-0.0348**	-0.0524***	-0.0170
	(0.0134)	(0.0150)	(0.0164)
Sex	0.292***		
	(0.0644)		
Female Head of Household	-0.0611	-0.0156	-0.111
	(0.0704)	(0.0965)	(0.0886)
Year FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes
N	19257	8715	10542

The dependent variable is the number of years of completed schooling an individual has. Electricity is a binary variable with one being if the municipality has electricity and zero otherwise. Years of childhood electricity is the number of years an individual had electricity between 6 and 12. The sample is limited to individuals who have finished schooling, are under the age of 24 and have not moved. Columns 2, 4 and 6 limit the sample to those individuals who had at least one year of electricity as a child. The first 2 columns include the entire sample, columns 3-4 limit the sample to only females and columns 5-6 limit the sample to only males. All errors are clustered at the municipality level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.3: Educational Attainment: Hazard of Dropping Out

Outcome: Hazard of Dropping Out			
	Full	Women	Men
Electricity1	1.139***	1.390***	0.882***
	(0.117)	(0.135)	(0.167)
Electricity2	0.344***	0.297**	0.378***
	(0.0833)	(0.141)	(0.0983)
Electricity3	-0.0606	-0.103	-0.0251
	(0.0640)	(0.0649)	(0.105)
Electricity4	-0.0667	-0.0606	-0.0712
	(0.118)	(0.139)	(0.130)
Electricity5	-0.198*	-0.303**	-0.105
	(0.110)	(0.122)	(0.143)
N	19909	9476	10433
Year FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

The dependent variable is one if an individual drops out of school that year and zero otherwise. Electricity1, Electricity2, Electricity3, Electricity4, and Electricity5 is equal to one if the student is in a specific year of school (1-5) and the year is after the year of electrification and zero otherwise. The model is estimated using a Cox proportional hazard framework. The sample is limited to individuals who have finished their schooling, have non-zero years of schooling and have not moved. The first column includes the entire sample, columns 2 limits the sample to only females and columns 3 limits the sample to only males. All errors are clustered at the individual level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.4: Graph of Heterogeneous Cox Results

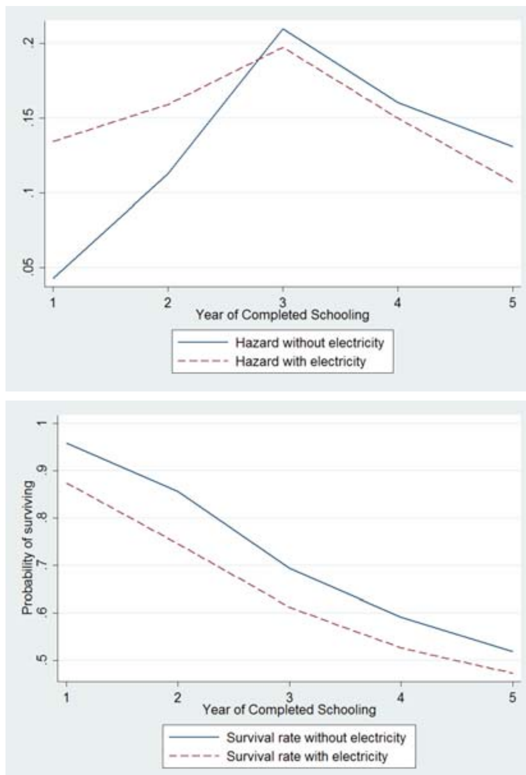

These graphs represent different predicted values from the first column of table 3.3. The hazard with(without) electricity represents the hazard the average individual would face if they did(not) have electricity. Similarly the survival rate with(without) electricity is the probability of completing each year of schooling with(without) access to electricity. To create the hazard without electricity predicted Xb was generated from the previous regression. I then subtracted off electricity times the coefficient on electricity. Finally the hazard displayed is the exponential of this multiplied by the negative of the baseline hazard. For the hazard with electricity I did the same thing except instead of adding electricity times the coefficient I simply added one minus electricity times the coefficient. For the two survival rates I calculated the exponential of the negative baseline hazard for each period and raised that to the hazard with or without electricity. Finally I multiplied each period by the period before to calculate the survivorship.

Table 3.5: Childhood Employment

Outcome: Employment			
	Full	Women	Men
Electricity	0.0238*	0.0289***	0.0143
	(0.0138)	(0.00958)	(0.0245)
Household Size	0.00420***	-0.00165	0.00902***
	(0.00159)	(0.00160)	(0.00213)
Age	0.0819***	0.0326***	0.127***
	(0.00191)	(0.00267)	(0.00301)
Female	-0.288***		
	(0.0125)		
Female Head of Household	-0.00761	0.0332***	-0.0474***
	(0.00981)	(0.0109)	(0.0130)
N	22923	11090	11833
Year FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes

The dependent variables is whether an individual is employed. Electricity is a binary variable with one being if the municipality has electricity and zero otherwise. The controls are household size, age, age squared, sex, and the gender of the head of the household. The first column includes the entire sample, column 2 limits the sample to only females and column 3 limits the sample to only males. All errors are clustered at the municipality level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.6: Female Employment

	Employed	Hours Worked	Log(Earnings)	Log(Wage)
Electricity	0.0564***	4.088***	0.150**	0.0311
	(0.0161)	(1.038)	(0.0736)	(0.0643)
Household Size	-0.0105***	-0.212*	-0.0250***	-0.0175***
	(0.00141)	(0.115)	(0.00750)	(0.00637)
Age	0.0200***	-0.108	0.0410***	0.0487***
	(0.00104)	(0.0978)	(0.00849)	(0.00862)
Age Squared	-0.000242***	0.0000167	-0.000549***	-0.000561***
	(0.0000105)	(0.00111)	(0.0000935)	(0.0000940)
Female Head of Household	0.189***	0.161	-0.0288	-0.0389
	(0.00991)	(0.658)	(0.0429)	(0.0409)
N	31168	9615	8663	8369
Year FE	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes

The dependent variables are whether an individual is employed, the hours an individual worked, the log of one plus labor earnings, and the log of one plus wage. Electricity is a binary variable with one being if the municipality has electricity and zero otherwise. The sample is limited to only females above the age of 18. All errors are clustered at the municipality level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.7: Impact of Childcare

Outcome: Hazard of Dropping Out			
	Full	Women	Men
Electricity1	1.011***	1.149***	0.903***
	(0.131)	(0.175)	(0.180)
Electricity2	0.256**	0.275*	0.239*
	(0.106)	(0.166)	(0.138)
Electricity3	-0.128*	-0.137	-0.124
	(0.0714)	(0.0919)	(0.129)
Electricity4	-0.117	-0.145	-0.0979
	(0.134)	(0.191)	(0.154)
Electricity5	-0.215	-0.330*	-0.122
	(0.137)	(0.187)	(0.152)
Electricity1*Toddler	0.101	0.173*	-0.0323
	(0.0748)	(0.0939)	(0.116)
Electricity2*Toddler	0.0733	0.00866	0.124*
	(0.0607)	(0.0945)	(0.0724)
Electricity3*Toddler	0.0602	0.0206	0.102
	(0.0514)	(0.0555)	(0.0828)
Electricity4*Toddler	0.0491	0.0700	0.0314
	(0.0484)	(0.0904)	(0.0713)
Electricity5*Toddler	0.0169	0.0225	0.0211
	(0.0886)	(0.112)	(0.0948)
Toddler	0.0712***	0.0858***	0.0569***
	(0.00916)	(0.00980)	(0.0125)
N	89246	42887	46359
Year FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

The dependent variable is one if an individual drops out of school that year and zero otherwise. Electricity1, Electricity2, Electricity3, Electricity4, and Electricity5 are equal to one if the student is in a specific year of school (1-5) and the year is after the year of electrification and zero otherwise. Toddler is equal to the number of children under five in a household. The model is estimated using a Cox proportional hazard framework. The sample is limited to individuals who have finished their schooling, have non-zero years of schooling and have not moved. The first column includes the entire sample, columns 2 limits the sample to only females and columns 3 limits the sample to only males. All errors are clustered at the individual level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.8: Returns to Education

Outcome: Log of Wage			
	Full	Women	Men
Years of Schooling 1-2	0.220***	0.284**	0.215**
	(0.0766)	(0.117)	(0.0879)
Years of Schooling 3-4	0.262***	0.368**	0.232***
	(0.0506)	(0.154)	(0.0472)
Years of Schooling 5-6	0.503***	0.753***	0.433***
	(0.0704)	(0.131)	(0.0646)
Household Size	-0.00347	-0.00675	-0.00378
	(0.00812)	(0.0186)	(0.00929)
Age	0.0531***	0.0403**	0.0570***
	(0.00650)	(0.0162)	(0.00642)
Age Squared	-0.000538***	-0.000373*	-0.000589***
	(0.0000751)	(0.000198)	(0.0000709)
Female	-0.205***		
	(0.0654)		
Female Head of Household	-0.0788	-0.121	-0.0369
	(0.0517)	(0.0880)	(0.0797)
N	4490	978	3512
Year FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes

The dependent variables is the log of one plus an individuals wage. Years of Schooling 0, 1-2, 3-4, and 5-6 are binary variables which are one if an individuals maximum years of schooling is equal to the number and zero otherwise. The sample is limited to places before they gain access to electricity. The first column includes the entire sample, column 2 limits the sample to only females and column 3 limits the sample to only males. All errors are clustered at the municipality level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.9: Returns to Education with Electricity

Outcome: Log of Wage			
	Full	Women	Men
Electricity*Years of Schooling 0	-0.0189	-0.0205	0.00373
	(0.0544)	(0.0770)	(0.0621)
Electricity*Years of Schooling 1-2	0.0679	0.155	0.0312
	(0.0593)	(0.103)	(0.0657)
Electricity*Years of Schooling 3-4	0.0699	0.128	0.0435
	(0.0690)	(0.125)	(0.0686)
Electricity*Years of Schooling 5-6	0.0256	-0.0674	0.0430
	(0.0617)	(0.0856)	(0.0619)
Years of Schooling 1-2	0.208***	0.242**	0.217**
	(0.0752)	(0.114)	(0.0845)
Years of Schooling 3-4	0.284***	0.379***	0.267***
	(0.0490)	(0.140)	(0.0476)
Years of Schooling 5-6	0.508***	0.733***	0.466***
	(0.0619)	(0.112)	(0.0597)
Household Size	-0.0137***	-0.0111*	-0.0142***
	(0.00358)	(0.00627)	(0.00408)
Age	0.0519***	0.0514***	0.0505***
	(0.00344)	(0.00794)	(0.00311)
Age Squared	-0.000515***	-0.000521***	-0.000498***
	(0.0000353)	(0.0000833)	(0.0000328)
Female	-0.0586*		
	(0.0312)		
Female Head of Household	-0.0824***	-0.0378	-0.122***
	(0.0228)	(0.0408)	(0.0251)
N	30145	8161	21984
Year FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes

The dependent variables is the log of one plus an individuals wage. Ect is a binary variable with one being if the municipality has electricity and zero otherwise. Years of Schooling 0, 1-2, 3-4, and 5-6 are binary variables which are one if an individuals maximum years of schooling is equal to the number and zero otherwise. The first column includes the entire sample, column 2 limits the sample to only females and column 3 limits the sample to only males. All errors are clustered at the municipality level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3.6: Instrument Creation

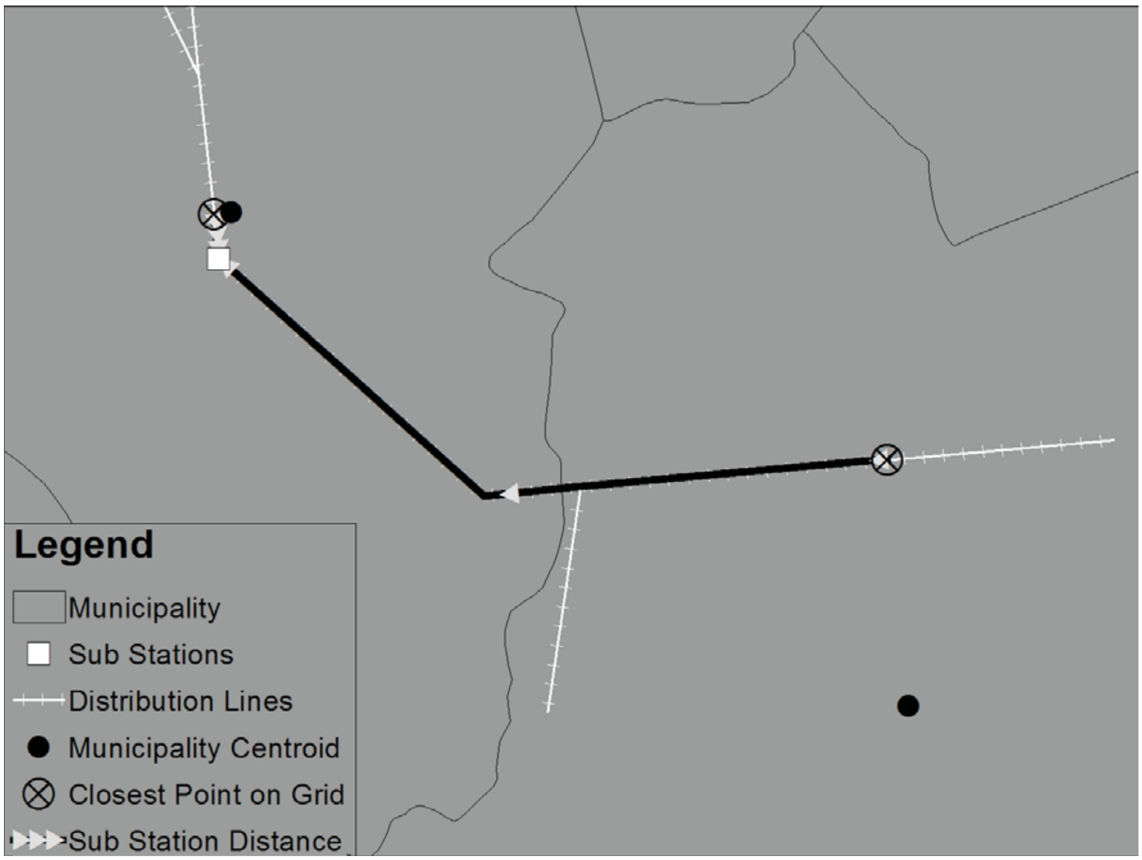

The instrument was created in two steps. First the nearest point, on the distribution line, to the centroid of the municipality was determined. Second the distance to the nearest substation from this point was calculated. All municipalities without a distribution line were removed and the final value is in kilometers.

Table 3.10: Test 1: Relevance

Outcome: Year of Electrification			
Grid Distance	0.0340***		
	(0.0108)		
Straight Distance		0.0439***	
		(0.0129)	
Average Slope			-0.0527
			(0.141)
Age	-0.0849	-0.103	-0.117
	(0.129)	(0.128)	(0.129)
Sex	-19.12**	-14.87*	-14.71*
	(7.921)	(8.190)	(8.535)
Household Size	-0.383	-0.652	-0.439
	(0.451)	(0.473)	(0.452)
Female Head of Household	-0.8897	-3.535	-2.776
	(3.692)	(3.756)	(3.928)
N	85	87	87

The dependent variable is the year each municipality got electricity. Grid distance is the distance along the distribution line from a substation to the closest point to the centroid of a municipality. Straight distance is the distance from the centroid of a municipality to the nearest substation. Average slope is the average slope in the municipality. The control variables, age, sex, household size, and gender of household head are the average for the municipality in the first year that data is available. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.11: Test 2: Exclusion Restriction

	Employment	Hours Worked	Log Earnings	Log Wage	School Attendance	Years of Schooling	Log of Household Income
Grid	0.0000145	-0.0208	0.000887	0.00132	0.000762	0.000718	-0.000514
Distance	(0.000202)	(0.0176)	(0.00204)	(0.00163)	(0.000513)	(0.00135)	(0.00146)
N	10505	6159	5336	5272	6185	6185	3838
Straight	0.000521***	0.0108	0.00657***	0.00594***	0.00113***	0.000656	0.00478***
Distance	(0.000108)	(0.00718)	(0.000858)	(0.000695)	(0.000229)	(0.000607)	(0.000797)
N	10617	6226	5397	5333	6262	6262	3880

The dependent variables are whether an individual is employed, the number of non-zero hours worked, the log of non-zero earnings, and the log of non-zero wages respectively averaged over individuals in a municipality. Grid distance is the distance along the distribution line from a substation to the closest point to the centroid of a municipality. Straight distance is the distance from the centroid of a municipality to the nearest substation. The sample is limited to observations before the municipality had electricity. All errors are clustered at the municipality level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.12: Test 3: Population Density and Growth

	Population 1988	Population 2001	Population Growth 1974-1988	Population Growth 1974-2001	Population Growth 2001-2010
Grid	-0.0627	-0.125	-0.00238	-0.00416	0.0000701
Distance	(0.0935)	(0.117)	(0.00409)	(0.00286)	(0.00380)
Straight	-0.0897	-0.155	-0.00303	-0.00621*	0.00987**
Distance	(0.114)	(0.139)	(0.00657)	(0.00316)	(0.00428)
Sample	All	All	All	All	Post 2001

The dependent variables are population and population growth rates by municipality. Rid distance and straight distance were run in seperate regressions. Grid distance is the distance along the distribution line from a substation to the closest point to the centroid of a municipality. Straight distance is the distance from the centroid of a municipality to the nearest substation. In column four the sample is limited to municipalities that get electricity post 2001. All errors are clustered at the municipality level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.13: Instrumented Educational Attainment

Outcome: Hazard of Dropping Out			
	Full	Women	Men
$\hat{Electricity}1$	1.628***	1.8271***	1.3879***
	(0.0987)	(0.1299)	(0.1549)
$\hat{Electricity}2$	0.5922***	0.7523***	0.4340***
	(0.0881)	(0.1197)	(0.1303)
$\hat{Electricity}3$	0.0511	0.0271	0.0598
	(0.0802)	(0.1265)	(0.1052)
$\hat{Electricity}4$	0.0154	0.02847	-0.0041
	(0.0926)	(0.1333)	(0.1300)
$\hat{Electricity}5$	-0.17998*	-0.2258	-0.1560
	(0.1015)	(0.1511)	(0.1390)
N	37302	18240	19062
		First Stage	
Grid	0.0340***	0.0340***	0.0340***
Distance	(0.0105)	(0.0105)	(0.0105)
Year FE	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Bootstrapped	Yes	Yes	Yes

The dependent variable is one if an individual drops out of school that year and zero otherwise. $\hat{Electricity}1$, $\hat{Electricity}2$, $\hat{Electricity}3$, $\hat{Electricity}4$, and $\hat{Electricity}5$ is equal to one if the student is in a specific year of school (1-5) and the year is after the predicted year of electrification and zero otherwise. The predicted year of electrification is created by instrumenting for the year a municipality got electricity with Grid distance. The model is estimated using a Cox proportional hazard framework. The sample is limited to individuals who have finished their schooling, have non-zero years of schooling and have not moved. The first column includes the entire sample, columns 2 limits the sample to only females and columns 3 limits the sample to only males. All errors are clustered at the individual level and the entire two stage process is bootstrapped. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3.7: Migration

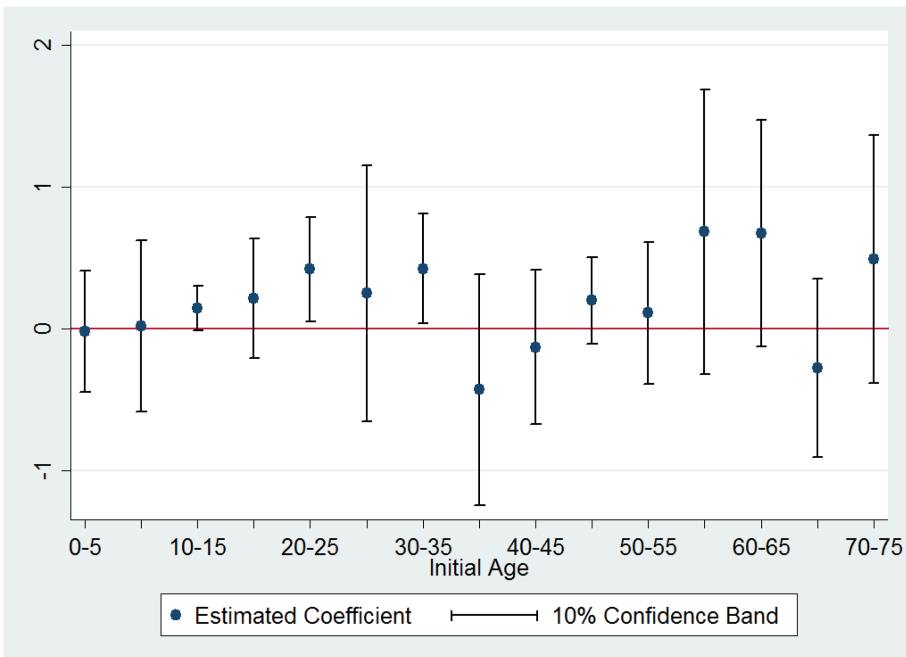

This graph represents the coefficient on electricity in a regression of electricity on the ratio of ages in a municipality. The ratio is ages 2 years divided by ages today. For example the amount of individuals who are 2-7 in two years divided by the number who are 0-5 today. Electricity is defined as 1 if a municipality gained access to electricity in those two years and zero if they did not have electricity in those two years.

Bibliography

Acemoglu, D., S. Johnson, and J. A. Robinson (2001): "The Colonial Origins of Comparative Development: An Empirical Investigation," American Economic Review, 91(5), 1369–1401.

Alesina, A., and E. L. Ferrara (2000): "Participation in Heterogeneous Communities," The Quarterly Journal of Economics, 115(3), pp. 847–904.

Andrews, M., W. Duncombe, and J. Yinger (2002): "Revisiting economies of size in American education: are we any closer to a consensus?," Economics of Education Review, 21(3), 245–262.

Antos, J. R., and S. Rosen (1975): "Discrimination in the market for public school teachers," Journal of Econometrics, 3(2), 123–150.

Ashraf, Q. H., A. Lester, and D. N. Weil (2008): "When Does Improving Health Raise GDP?," Working Paper 14449, National Bureau of Economic Research.

Atkin, D. (2012): "Endogenous Skill Acquisition and Export Manufacturing in Mexico," NBER Working Papers 18266, National Bureau of Economic Research, Inc.

Babcock, L., and J. Engberg (1999): "Bargaining unit composition and the returns to education and tenure," Industrial & labor relations review, 52(2), 163–178.

Ballou, D., and M. Podgursky (2002): "Returns to seniority among public school teachers," Journal of Human Resources, pp. 892–912.

Barron, M., and M. Torero (2014): "Short Term Effects of Household Electrification: Experimental Evidence from Northern El Salvador," Job Market Paper.

Barrow, L., and C. E. Rouse (2004): "Using market valuation to assess public school spending," Journal of Public Economics, 88(9), 1747–1769.

Baugh, W. H., and J. A. Stone (1982): "Mobility and wage equilibration in the educator labor market," Economics of Education Review, 2(3), 253–274.

Bensch, G., J. Kluve, and J. Peters (2011): "Impacts of Rural Electrification in Rwanda," IZA Discussion Papers 6195, Institute for the Study of Labor (IZA).

Bensch, G., J. Peters, and M. Sievert (2012): "Fear of the Dark? How Access to Electric Lighting Affects Security Attitudes and Nighttime Activities in Rural Senegal," Ruhr Economic Papers 0369, Rheinisch-Westflisches Institut fr Wirtschaftsforschung, Ruhr-Universitt Bochum, Universitt Dortmund, Universitt Duisburg-Essen.

Bernard, T. (2012): "Impact Analysis of Rural Electrification Projects in Sub-Saharan Africa," World Bank Research Observer, 27(1), 33–51.

Bifulco, R. (2010): "The Influence of Performance-Based Accountability on the Distribution of Teacher Salary Increases," Education Finance and Policy, 5(2), 177–199.

Bleakley, H. (2010): "Malaria Eradication in the Americas: A Retrospective Analysis of Childhood Exposure," American Economic Journal: Applied Economics, 2(2), 1–45.

Brasington, D. M. (2001): "Capitalization and community size," Journal of Urban Economics, 50(3), 385–395.

——— (2002): "Edge versus center: finding common ground in the capitalization debate," Journal of Urban Economics, 52(3), 524–541.

Brown, C. C., and J. L. Medoff (1988): "Employer size, pay, and the ability to pay in the public sector," in When public sector workers unionize, pp. 195–216. University of Chicago Press.

Calderon, C. A., and L. Servn (2008): "Infrastructure and economic development in Sub-Saharan Africa," Policy Research Working Paper Series 4712, The World Bank.

Chakraborty, K., B. Biswas, and W. C. Lewis (2000): "Economies of scale in public education: an econometric analysis," Contemporary Economic Policy, 18(2), 238–247.

Chakravorty, U., M. Pelli, and B. U. Marchand (2014): "Does the quality of electricity matter? Evidence from rural India," Journal of Economic Behavior & Organization, (0), –.

Clotfelter, C. T., H. F. Ladd, and J. L. Vigdor (2006): "Teacher-student matching and the assessment of teacher effectiveness," Journal of Human Resources, 41(4), 778–820.

Cruz, R. H. (2002): "National report on the results of the child labour survey in Honduras," Discussion paper, International Programme for the Elimination of Child Labour.

Cutler, D. M., W. Fung, M. Kremer, M. Singhal, and T. Vogl (2010): "Early Life Malaria Exposure and Adult Outcomes: Evidence from Malaria Eradication in India," American Economic Journal: Applied Economics, 2, 196–202.

Davis, M. L., and K. Hayes (1993): "The demand for good government," The review of economics and statistics, pp. 148–152.

Dee, T. S. (2004): "Are there civic returns to education?," Journal of Public Economics, 88(9), 1697–1720.

Delatte, H., J. Dehecq, J. Thiria, C. Domerg, C. Paupy, and D. Fontenille (2008): "Geographic Distribution and Developmental Sites of Aedes albopictus. During a Chikungunya Epidemic Event," Vector-Borne and Zoonotic Diseases, 8(1), 25–34.

Dinkelman, T. (2011): "The Effects of Rural Electrification on Employment: New Evidence from South Africa," American Economic Review, 101(7), 3078–3108.

Duflo, E., and R. Pande (2007): "Dams," The Quarterly Journal of Economics, 122(2), 601–646.

Duryea, S., and M. Arends-Kuenning (2003): "School Attendance, Child Labor and Local Labor Market Fluctuations in Urban Brazil," World Development, 31(7), 1165 – 1178, Economic Crises, Natural Disasters, and Poverty.

Easterly, W., and R. Levine (2003): "Tropics, germs, and crops: how endowments influence economic development," Journal of Monetary Economics, 50(1), 3–39.

Eberts, R. W. (2007): "Teachers unions and student performance: help or hindrance?," The future of children, 17(1), 175–200.

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

Fischel, W. A. (2001): The homevoter hypothesis: How home values influence local government taxation, school finance, and land-use policies. Harvard University Press Cambridge, MA.

Gallup, J., and J. Sachs (2001): "The economic burden of malaria," The American Journal of Tropical Medicine and Hygiene, 64(1 suppl), 85–96.

GOLDHABER, D., M. DEARMOND, AND S. DEBURGOMASTER (2010): "Teacher Attitudes about Compensation Reform: Implications for Reform Implementation. Working Paper 50.," National Center for Analysis of Longitudinal Data in Education Research.

GORDON, R. J., T. J. KANE, AND D. STAIGER (????): Identifying effective teachers using performance on the job. Brookings Institution Washington, DC.

GROGAN, L. (2012): "Household Electrification, fertility and employment: Evidence from the Colombian censuses," .

——— (2013): "Rural Electrification and Employment in Poor Countries: Evidence from Nicaragua," Working Papers 1303, University of Guelph, Department of Economics and Finance.

GROSSKOPF, S., K. J. HAYES, L. L. TAYLOR, AND W. L. WEBER (2001): "On the determinants of school district efficiency: Competition and monitoring," Journal of Urban Economics, 49(3), 453–478.

HANUSHEK, E. A., J. F. KAIN, D. M. O'BRIEN, AND S. G. RIVKIN (2005): "The market for teacher quality," Discussion paper, National Bureau of Economic Research.

HANUSHEK, E. A., AND J. LUQUE (2000): "Smaller classes, lower salaries? The effects of class size on teacher labor markets," Using What We Know: A Review of the Research on Implementing Class0Size Reduction Initiatives for State and Local Policymakers, Oak Brook, IL: North Central Regional Educational Laboratory, 35051.

HARWOOD, B., AND L. D. MULL (2002): "Education to Combat Abusive Child Labor: Planning Educational Response Strategies for Working Children in Honduras," Discussion paper, Creative Associates International, Inc., Washington, DC: Creative Associates International, for US Agency for International Development.

HAY, S. I., C. A. GUERRA, P. W. GETHING, A. P. PATIL, A. J. TATEM, A. M. NOOR, C. W. KABARIA, B. H. MANH, I. R. F. ELYAZAR, S. BROOKER, D. L. SMITH, R. A. MOYEED, AND R. W. SNOW (2009): "A World Malaria Map: Plasmodium falciparum Endemicity in 2007," PLoS Med, 6(3), e1000048.

HENDERSON, J. V., A. STOREYGARD, AND D. N. WEIL (2012): "Measuring Economic Growth from Outer Space," American Economic Review, 102(2), 994–1028.

Hess, F. M., and D. L. Leal (2005): "School house politics: Expenditures, interests, and competition in school board elections," Besieged: School boards and the future of education politics, 249.

Hilber, C. A., and C. Mayer (2009): "Why do households without children support local public schools? Linking house price capitalization to school spending," Journal of Urban Economics, 65(1), 74–90.

Hodler, R., and P. A. Raschky (2014): "Regional Favoritism," The Quarterly Journal of Economics, 129(2), 995–1033.

Hoxby, C. M. (1996): "How teachers' unions affect education production," The Quarterly Journal of Economics, pp. 671–718.

Hoyt, W. H. (1999): "Leviathan, local government expenditures, and capitalization," Regional Science and Urban Economics, 29(2), 155–171.

Kenny, L. W. (1983): "Economies of scale in schooling," Economics of Education Review, 2(1), 1–24.

Khandker, S. R., D. F. Barnes, and H. A. Samad (2009): "Welfare impacts of rural electrification : a case study from Bangladesh," Policy Research Working Paper Series 4859, The World Bank.

Khandker, S. R., D. F. Barnes, and H. A. Samad (2013): "Welfare Impacts of Rural Electrification: A Panel Data Analysis from Vietnam," Economic Development and Cultural Change, 61(3), 659 – 692.

Khandker, S. R., H. A. Samad, R. Ali, and D. F. Barnes (2012): "Who benefits most from rural electrification? evidence in India," Working Paper.

Kiszewski, A., A. Mellinger, P. Malaney, A. Spielman, S. Ehrlich, and J. D. Sachs (2004): "A Global Index Representing The Stability Of Malaria Transmission," The American Journal of Tropical Medicine and Hygiene, 70(5), 486–498.

Kleiner, M. M., and D. L. Petree (1988): "Unionism and licensing of public school teachers: Impact on wages and educational output," in When public sector workers unionize, pp. 305–322. University of Chicago Press.

KRUGER, D. I. (2007): "Coffee production effects on child labor and schooling in rural Brazil," Journal of Development Economics, 82(2), 448 – 463.

LANKFORD, H., AND J. WYCKOFF (1997): "The changing structure of teacher compensation, 1970–1994," Economics of Education Review, 16(4), 371–384.

LEHNER, B., AND P. DOLL (2004): "Development and validation of a global database of lakes, reservoirs and wetlands," Journal of Hydrology, 296(14), 1 – 22.

LEVINSON, A. M. (1988): "Reexamining teacher preferences and compensating wages," Economics of Education Review, 7(3), 357–364.

LIPSCOMB, M., A. M. MOBARAK, AND T. BARHAM (2013): "Development Effects of Electrification: Evidence from the Topographic Placement of Hydropower Plants in Brazil," American Economic Journal: Applied Economics, 5(2), 200–231.

LOVENHEIM, M. 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.

LUCAS, A. M. (2010): "Malaria Eradication and Educational Attainment: Evidence from Paraguay and Sri Lanka," American Economic Journal: Applied Economics, 2(2), 46–71.

LUTZ, B. F. (2009): Fiscal amenities, school finance reform and the supply side of the Tiebout market. Divisions of Research & Statistics and Monetary Affairs, Federal Reserve Board.

MICHALOPOULOS, S., AND E. PAPAIOANNOU (2011): "Divide and Rule or the Rule of the Divided? Evidence from Africa," Working Paper 17184, National Bureau of Economic Research.

MOE, T. M. (2005): "Teacher unions and school board elections," Besieged: School boards and the future of education politics, pp. 254–287.

——— (2006): "Political Control and the Power of the Agent," Journal of Law, Economics, and Organization, 22(1), 1–29.

MURNANE, R., J. D. SINGER, J. B. WILLETT, J. KEMPLE, AND R. OLSEN (2009): Who will teach?: Policies that matter. Harvard University Press.

NOOR, A. M., P. W. GETHING, V. A. ALEGANA, A. P. PATIL, S. I. HAY, E. MUCHIRI, E. JUMA, AND R. W. SNOW (2009): "The risks of malariainfection in Kenya in 2009," BMC Infectious Diseases, 9(1), 180.

Olson, D. M., E. Dinerstein, E. D. Wikramanayake, N. D. Burgess, G. V. N. Powell, E. C. Underwood, J. A. D'amico, I. Itoua, H. E. Strand, J. C. Morrison, C. J. Loucks, T. F. Allnutt, T. H. Ricketts, Y. Kura, J. F. Lamoreux, W. W. Wettengel, P. Hedao, and K. R. Kassem (2001): "Terrestrial Ecoregions of the World: A New Map of Life on Earth: A new global map of terrestrial ecoregions provides an innovative tool for conserving biodiversity," BioScience, 51(11), 933–938.

Pachauri, S., and D. Spreng (2011): "Measuring and monitoring energy poverty," Energy Policy, 39(12), 7497 – 7504, Clean Cooking Fuels and Technologies in Developing Economies.

Ramankutty, N., J. A. Foley, J. Norman, and K. McSweeney (2002): "The global distribution of cultivable lands: current patterns and sensitivity to possible climate change," Global Ecol Biogeography, 11(5), 377392.

Rockoff, J. E. (2004): "The impact of individual teachers on student achievement: Evidence from panel data," American Economic Review, pp. 247–252.

Rodrik, D., A. Subramanian, and F. Trebbi (2004): "Institutions Rule: The Primacy of Institutions Over Geography and Integration in Economic Development," Journal of Economic Growth, 9(2), 131–165.

Rose, H., and J. Sonstelie (2010): "School board politics, school district size, and the bargaining power of teachers unions," Journal of urban economics, 67(3), 438–450.

Sachs, J., and P. Malaney (2002): "The economic and social burden of malaria," Nature, 415(6872), 680–685.

Sachs, J. D. (2003): "Institutions Don't Rule: Direct Effects of Geography on Per Capita Income," Working Paper 9490, National Bureau of Economic Research.

Saiz, A. (2010): "The geographic determinants of housing supply," The Quarterly Journal of Economics, 125(3), 1253–1296.

Shah, M., and B. M. Steinberg (2013): "Drought of Opportunities: Contemporaneous and Long Term Impacts of Rainfall Shocks on Human Capital," Working Paper 19140, National Bureau of Economic Research.

Stone, J. A. (2000): "Collective bargaining and public schools," Conflicting missions? Teachers unions and educational reform, pp. 47–68.

Strunk, K. O., and J. A. Grissom (2010): "Do strong unions shape district policies? Collective bargaining, teacher contract restrictiveness, and the political power of teachers unions," Educational Evaluation and Policy Analysis, 32(3), 389–406.

van de Walle, D., M. Ravallion, V. Mendiratta, and G. Koolwal (2013): "Long-term impacts of household electrification in rural India," Policy Research Working Paper Series 6527, The World Bank.

West, K. L., and E. Mykerezi (2011): "Teachers unions and compensation: The impact of collective bargaining on salary schedules and performance pay schemes," Economics of Education Review, 30(1), 99–108.

WHO (1997): "World Malaria Situation in 1994. Part I," Discussion paper, Weekly Epidemiological Record.

——— (2009): "Dengue: guidelines for diagnosis, treatment, prevention and control," Discussion paper, World Health Organization.