ABSTRACT

Title of Document: ESSAYS ON ENERGY EFFICIENCY AND FOREST CONSERVATION

Joseph Andrew Maher, Doctor of Philosophy, 2015

Directed By: Distinguished University Professor, Richard E. Just, Department of Agricultural and Resource Economics

This dissertation is composed of three essays in environmental economics related to residential energy efficiency and forest conservation.

My first paper assesses the effectiveness of energy-efficient technologies in the setting of a utility rebate program. To date, the energy savings from energy-efficiency building retrofits are assessed using ex-ante engineering models. My analysis provides the first evaluation of engineering models that uses residential billing data, combined with data on observable characteristics of each residence, to assess the accuracy of engineering predictions across nine retrofit technologies used in Gainesville, Florida.

My second essay presents the first causal evidence that trees have a major impact on consumer demand—with large shade trees reducing household electricity use by more than 20 percent. This work contributes to the existing literature on the energy saving potential of urban forests by implementing a quasi-experimental design to identify a causal link between tree shade and energy use. Results suggest that the energy savings from tree shade are an order of magnitude greater than other energy-efficiency policy measures, providing new evidence that tree ordinances may serve as effective demand-side management policies.

My third essay assesses the effectiveness of forest conservation policies in reducing carbon emissions from deforestation. To date, the effectiveness of protected areas has been assessed using cross-sectional methods. In this essay, new quasi-experimental models using panel data on annual deforestation are used to reveal new insights into the importance of government oversight of protected areas with findings that counter economists' prior notions of the avoided deforestation of new parks. I extend the analysis to estimate avoided carbon emissions, a key policy metric that varies considerably from deforestation trends.

ESSAYS ON ENERGY EFFICIENCY AND FOREST CONSERVATION

By

Joseph Andrew Maher

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2015

Advisory Committee:
Professor Richard Just, Chair
Professor Maureen Cropper
Professor Roberton Williams III
Assistant Professor Sebastien Houde
Professor Kenneth McConnell

© Copyright by

Joseph Andrew Maher

2015

Dedication

I dedicate this work to my loving, caring, and always supportive wife, Rachel Maher.

Table of Contents

List of Tables

List of Figures

ESSAY 1 : **SUBSIDIES FOR ENERGY EFFICIENT BUILDING RETROFITS**

Buildings account for 42 percent of energy use and 38 percent of CO_2 emissions in the United States (US Green Building Council, 2011). In recent years, State and Federal governments have increased funding programs that subsidize energy-efficient retrofits to existing buildings. For example, the American Recovery and Reinvestment Act of 2009 (ARRA) included $17 billion for energy-efficiency programs, which helped to initiate $54 billion in energy-related home improvements in 2009 (von Schrader 2010). In 2013, President Obama announced a new goal, "Let's cut in half the energy wasted by our homes and businesses over the next 20 years. We'll work with the states to do it."[1]

Despite the widespread implementation of retrofit rebate programs and calls for increased investment in demand side management programs, surprisingly little is known about whether energy-efficiency retrofits are an effective way to reduce energy consumption. Engineering simulations provide most of the evidence, but simulated predictions, even if based on sound models, do not account for installation quality or behavioral responses (Allcott and Greenstone 2012). Hence, there is an important and timely need for empirical research that uses field data to evaluate more fully the effects of energy-efficiency retrofits on energy consumption.

---

[1] Statement delivered on February 12, 2013, in a State of the Union Speech. Obama continues to promise that "Those states with the best ideas to create jobs and lower energy bills by constructing more efficient buildings will receive federal support to help make that happen."

This essay answers the basic question: How much do energy efficiency investments reduce energy use? Next, it compares these estimated energy savings with engineering models, to answer the question: How do engineering model biases contribute to the energy efficiency gap? Then it compares aggregate benefits against private and social costs of retrofit installations, to answer the questions: How cost-effective are rebate programs in reducing energy demand? What is the private return-on-investment for retrofit installations?

EMPIRICAL SETTING AND DATA COLLECTION

A central barrier for this type of research is the difficulty of obtaining data necessary to better understand energy efficiency investment behavior and its implications for energy use. Furthermore, empirical assessments are context specific. Generalizing the findings from one energy use and energy user context to other users and uses may have limited validity.

Despite these challenges and limitations, empirical ex-post assessments of energy savings from efficiency investments present a tremendous opportunity to improve demand-side management policies. Empirical research quantifying the size and nature of energy savings in different energy uses and user contexts is critical to understanding the social benefits of efficiency investments. Understanding heterogeneity of such benefits helps reveal which policy interventions are effective in achieving program goals. Putting this information to action by targeting policies towards specific energy uses and energy users will improve policy effectiveness. This essay develops empirical estimates of energy savings from a wide range of energy efficiency investments and housing contexts using data on household energy consumption and efficiency investments. The study exploits three key datasets for Gainesville, Florida. The econometric approach uses difference in difference methods.

The purpose of this essay is to evaluate retrofit-specific residential rebate programs based on observed household-level consumption data. I can identify the energy savings from nine retrofit rebate programs in Gainesville, Florida, using a panel dataset of electricity and natural gas consumption and building characteristics for 30,000 residences. The difference-in-difference method compares changes in energy use within a residence before and after an energy-saving retrofit intervention (treatment group) with changes in energy use within a similar residence that did not receive improvements (control group). A unique feature of the data, which is central to my identification strategy, is that I have monthly billing data combined with time-variant and time-constant characteristics of each residence.

The essay tests two hypotheses to identify key variations in energy savings. First, I examine seasonal variation in energy savings, testing whether energy savings for each technology vary across seasons. Second, I focus on heterogeneity across user contexts, testing whether the energy savings of a technology varies across building characteristics. Results suggest that ex-post assessments provide rich and valuable information about the heterogeneity of size and timing of energy savings across different technologies and energy users.

This work makes general contributions to the literature assessing energy efficiency programs. First, this is the first assessment of a retrofit rebate program to apply difference-in-difference methods linking billing data and housing characteristics for every customer within a utility service area. Second, by assessing nine retrofit programs, this study explores heterogeneity across a diverse range of retrofit options using actual billing data.

The Energy-Efficiency Rebate Program

The Gainesville Regional Utilities (GRU) is a municipal utility and exclusive supplier of electricity and natural gas for over 30,000 households in Gainesville, Florida. As the 5th largest electric utility in Florida, GRU is responsible for almost all of the electricity generation for the city, as well as transmission and distribution. GRU is also an award-winning national leader in energy conservation with the ability to demonstrate the potential of well-designed retrofit rebate programs.[2]

The Gainesville Residential Energy-Efficiency Incentives (GREEN) was a utility-run subsidy program that provided rebates and low-cost loans for the purchase of energy-efficient products. Established in 2008, GREEN was developed from federal funding under the American Reinvestment and Recovery Act of 2008 (ARRA). In 2013, all GREEN subsidy programs were terminated following the expiration of ARRA funding, a short life-cycle shared by dozens of utility-run subsidy programs funded by ARRA. Nonetheless, the GREEN program spurred investment across more than a dozen retrofit products, which provides a unique opportunity for quasi-experimental analysis of retrofit programs.

GREEN had a large number of participants for a utility-level program. By December 2012, GREEN provided subsidies for approximately 25,000 residential energy efficiency replacements. In total, more than 14,000 households participated in the

---

[2] Awards include the 2005 Green Power Beacon Award presented by the U.S. Environmental Protection Agency, the U.S. Department of Energy, and the Center for Resource Solutions. (http://www.epa.gov/greenpower/documents/2005awards.pdf; https://www.gru.com/AboutGRU/NewsReleases/Archives/Articles/news-2005-10-28.jsp)

program, with many participants qualifying for subsidies for multiple retrofit products. Program take-up was unusually high, with approximately half of eligible GRU customers participating in GREEN over the five-year program lifetime. Most rebate programs were available to all GRU customers.

GREEN reimbursement required submission of original billing receipts from a professional contracting company. GRU recorded the exact date of installation, total project costs, company name, and details of equipment and materials detailed on rebate application forms. Participants were required to hire GRU certified professional contractors to be eligible for rebates, ensuring a basic level of quality assurance and accurate reporting by reputable businesses. A precise and accurate treatment date permits a clean timing for my difference-in-difference model. Total project costs, separately reported for each retrofit, provide additional information about the homeowners' willingness to pay for retrofits.

Each rebate program issued rebates based on a formula that reflects treatment intensity. The program provides both fixed cash rebates and sliding capped rebates. All programs have a maximum rebate level. Exact formulas, rebate rates, and maximum caps vary from year to year with each program. Program inception dates and program termination dates vary across retrofit types. Households could accept rebates as a cash subsidy, or as a credit to future utility bills. Program participants were reimbursed for the subsidy about one month after the file was completed, which required verified receipts for the total amount of the project cost. Funding for specific programs varied from year to year, and rebates were issued on a first-come-first-served basis until annual funding was exhausted.

Energy and Building Data

The major innovations of this essay stem from unusually detailed information about program participants, linked to natural gas and electricity billing data for all GRU customers. Furthermore, data are linked to building characteristics and home improvement permits.

Structural data are used to identify credible control households for each treatment household. Matching methods use observable characteristics of households, in addition to energy use patterns in pre-treatment periods, to match treatment homes that install retrofits with identical control homes that do not. To ensure the effects estimated in this essay are caused by energy-efficiency retrofits, my sample excludes any households with home improvements unrelated to GREEN.

Rebate Program Participation Data

Program participant data include the date, rebate amount, company name, and technical details of retrofit intervention, in-house engineering estimates, and total cost of installation.

The GREEN participation data has several distinguishing features, including, (i) diversity of retrofit types, (ii) detail of technical specifications about retrofit equipment and materials, (iii) project-level estimates of expected energy savings, and (iv) project-level information of total retrofit costs and rebate amounts.

Table 1.1 (column 2) provides descriptive statistics of GREEN participation data. My analysis examines participants from fifteen retrofit programs (number of participants reported in parenthesis), including high-efficiency room air-conditioning replacement (812), high-efficiency central air-conditioning replacement (3,884), super

SEER central air-conditioning replacement (792), air-conditioner maintenance (3,158), duct leakage repair (2,702), natural gas water heater replacement/conversion (1,657), attic insulation installment (2,033), refrigerator buyback (3,702), pool pump replacement (797).

Each retrofit includes energy-saving predictions using GRU engineering simulations. These predictions, which are used for program evaluation purposes, allow comparisons between my estimates of observed energy-savings with engineering predictions. Table 1.1 (column 3) illustrates the average predicted energy savings by retrofit type.

Detailed technical specifications characterize each retrofit installation. For most retrofit programs, variables include continuous measures of the energy-efficiency ratings, which approximate treatment intensity. For example, attic insulation includes the R-value of added insulation, square footage of coverage, and type of insulation material. For equipment replacements, such as air-conditioners, data include the brand, model number, capacity, certification number, and the seasonal energy efficiency ratio (SEER) rating of installed equipment.[3]

Empirical Methods

I employ a difference-in-difference model to estimate the effect of retrofits on household energy consumption. Covariate matching is used to select control homes which do not receive energy-efficiency improvements but have a similar distribution of structural characteristics as the treatment group. The large size of treatment and control

[3] Data do not generally include information about the old equipment that is replaced. The refrigerator buyback program is an exception that includes data about removed equipment.

household groups allow me to estimate effects precisely while using highly nonparametric specifications.

The simple case of a single retrofit per household would entail estimating the following two-way fixed effects model:

$$y_{it} = \lambda_t + c_i + W_{it}\tau + \varepsilon_{it}\,, \qquad t = 1, \dots, T \tag{1}$$

where y_{it} is electricity consumption for house i in month t; λ_t is a month-specific effect for time t that is introduced with time-period indicator variables to capture city-wide trends that affect electricity consumption over time, such as weather fluctuations; c_i is a household-specific effect for house i that are introduced with house indicator variables to capture all time-constant factors of a house that affect electricity consumption; $W_{it} = [w_{it1}, \dots, w_{itJ}]$ is a row vector of retrofit-specific treatment indicators variables associated with observation y_{it}, where each element w_{itj} is equal to 1 for all months t after a retrofit of type j is installed in house i among houses making energy efficient upgrades, and equal to 0 otherwise; τ is a J-vector of retrofit-specific treatment effect coefficients that represent the average monthly energy savings from each retrofit type j, that is assumed to persist over time; T is a constant equal to the total number of billing months; J is a constant equal to the total number of rebate programs; and ε_{it} is an error term clustered by house and represents unmeasured time-variant factors affecting electricity consumption.

A two-way fixed effects model imposes very weak identifying assumptions. Specifically, my model maintains the standard fixed effect assumption of strict exogeneity of treatment, expressed as $Cov(\omega_{it}, \varepsilon_{it}) = 0 \;\forall\; i, t$. Endogeneity problems violating this assumption arise from four common sources: period effects, measurement error, unobserved heterogeneity, and simultaneity.

The model directly addresses period effects. Month fixed effects eliminate time-variant period effects common to all households in Gainesville within a billing period, as non-participant households serve as counterfactuals that isolate period effects. Treatment effects are identified from remaining *within variation beyond period effects*.

To address other forms of endogenity, this basic model assumes: (i) non-participants in the comparison group never receive unobserved retrofit interventions (a potential measurement error); (ii) treatments are uniform for each retrofit type (another potential measurement error) (iii) retrofit interventions are uncorrelated with time-variant changes affecting energy consumption, such as changes in occupancy or home-remodels (potential omitted variable bias); and (iv) pre-retrofit shocks to energy bills never trigger retrofit interventions (potential simultaneity bias). Extensions of model (1) tests for each of these additional sources of endogenity.

In addition, the sample of households is restricted in several ways. First, the analysis is restricted to single family households that have a single customer account during the four-year time period. Second, all treatment households are required to have at least 6 months of billing data pre-retrofit and post-retrofit to ensure consistent estimation of treatment effects. Third, the treatment group only includes households that receive a single retrofit intervention; households that receive multiple retrofits are excluded from the analysis.

EMPIRICAL ANALYSIS OF ENERGY SAVINGS

Table 1.2 presents results of the energy savings from retrofit installation. Electricity usage is reported in kilowatt hours per month. Coefficients reflect average monthly treatment effects across retrofit types, represented in equation (1). If retrofit installations increase efficiency, then energy use should decrease post retrofit for participating households. This is, in fact, the case for most retrofits.

Results confirm expectations, suggesting most retrofits reduce electricity use. At the upper extreme, a retrofit can save up to 190 kWh per month, or 17 percent of the median household energy consumption of 1,100 kWh per month.[4] Five programs have results significant at the 1 percent level (energy savings reported in parentheses), including refrigerator buyback (79 kWh per month), pool pump replacement (189 kWh per month), super SEER air-conditioner replacement (139 kWh per month), duct leakage repair (97 kWh per month), and attic insulation installment (64 kWh per month). Energy savings for air-conditioner maintenance (17 kWh per month) and natural gas water heater conversion (99 kWh per month) are significant at the 5 percent level. Room air-conditioner replacements may yield small energy savings, but coefficients are not different from zero at any reasonable level of statistical significance.

Seasonal Variation in Electricity Savings

Retrofit technologies offer varying energy services, and some services vary with weather, while others do not. Retrofits related to climate control should yield seasonally variable energy savings, whereas retrofits related to all-season appliances should yield

[4] The median energy consumption for the median house averaged 1,100 kWh across all months. Energy consumption varies seasonally.

constant levels of energy savings throughout the year. Results in Table 1.3 test this hypothesis by including an interaction between post-treatment and a seasonal indicator variable equal to one for electricity bills in the months of April through September. Seasonal interactions indicate the time of year that different retrofits provide energy savings, which is information useful for managing peak demand loads.

Energy savings are constant for products with year-round energy services, providing evidence that confirms energy changes are caused by retrofit installation. The seasonal interaction coefficient is not significant for both refrigerator and hot water heater retrofits—products with year-round energy services—suggesting constant energy savings during warm and cool seasons.

On the other hand, energy savings vary seasonally for products that provide climate control services. Attic insulation provides double the energy savings during warmer seasons, consistent with a humid tropical climate associated with extremely hot summer months and mild winter months. Air-conditioners provide almost all energy savings during warmer months, a result that holds for both room units and high-efficiency central units. Similarly, duct leakage repair and super SEER air-conditioners provide more energy savings during warmer months, but provide substantial energy savings throughout the year, perhaps due to combined use of central air systems such as heat pumps and air dehumidifiers.

Some retrofits create seasonally varying savings that are counter-intuitive because energy savings are produced by eliminating wasteful off-season appliance use. For example, pool pumps filter water to prevent algae growth need to be operated more intensively during summer months than winter months. So single-speed pumps have over-sized motors that waste energy during winter months. In contrast, high-efficiency pool pump motors have variable speeds and programmable timers, which are features that save energy during winter months by reducing hours of operation and lowering motor speeds. Similarly, air-conditioning maintenance often involves reprogramming thermostats to reduce wasteful off-season use and increase desirable cooling services during warmer months.

Building Size Effects on Electricity Savings

Although engineering estimates are reasonably accurate on average, engineering models may not capture variations in energy savings across different types of residential buildings. Due to structural differences between buildings, different homes may respond differently to the same types of retrofits. Understanding which households benefit most from retrofits has direct policy relevance as policy makers can target homeowners with the largest energy savings potential. Furthermore, identifying building characteristics that reduce the accuracy of engineering estimates can help engineers to improve the accuracy of the next generation of ex-ante models.

Regressions in Table 1.4 test whether building size affects energy savings by including an interaction between post-treatment and a "big house" indicator variable equal to one for homes with more than 2,000 square feet of conditioned living space. Results indicate that larger homes experience substantially larger energy savings from retrofits.

VALIDATION OF ENGINEERING ESTIMATES

A central debate within the energy efficiency literature concerns how energy savings should be measured. Most policies rely on engineering simulations. Home owners rely on ex-ante energy audits to guide investment decisions. Policy makers, in turn, apply the same engineering estimates to conduct ex-ante program evaluations. As a result, policies continue to be justified based on limited-information engineering estimates of the savings that the technologies could deliver.

Engineering predictions may be prone to bias from several sources. Faulty assumptions create a gap between realized and predicted energy savings. For example, engineering simulations tend to assume perfect installation and maintenance of energy efficient upgrades, thereby overstating the projected energy savings. Second, even if based on sound models, simulations fail to account for behavioral responses. In some cases, this can arise from failure to account for interactions between energy uses. For example, efficient light bulbs radiate less heat than incandescent bulbs, and as a result efficient lighting upgrades may increase usage of heating systems.

Typical engineering models assume constant utilization of energy services before and after efficiency investments. Higher efficiency reduces the marginal price of the energy services provided by a product and, consequently, leads to increased consumption of these energy services, a response called the *rebound effect*. For example, a new (more efficient) air-conditioner will lower the marginal price of cool air, leading homeowners to set thermostats at a lower temperature. If true, model bias may have implications for an energy-efficiency paradox and related cost-effectiveness.

More than 30 years of literature finds that engineering simulations rely on faulty assumptions that create a gap between realized and predicted energy savings. For example, Metcalf and Hassett (1999) find that engineering simulations over-predict the energy savings from attic insulation by over 500% when compared with actual savings in household energy bills. More recently, Davis et al. (2014) find that aggregate engineering figures overestimate effects of refrigerator replacements by 250% or more.

The large magnitude of engineering bias for residential retrofits is surprising, especially considering that validations in similar settings, such as residential building codes and commercial lighting retrofits, find engineering estimates are quite accurate (Jacobsen and Kotchen 2013, Lang and Siler 2013).

One possibility may be that residential retrofit validation studies represent "apples-to-oranges" comparisons whereby engineering bias may be exaggerated due to poorly defined retrofit parameters that differ from retrofits actually installed. Metcalf and Hassett (1999) and Davis et al. (2014) are both validation studies of residential retrofits compare empirical estimates to *a-priori* engineering estimates based on a hypothetical scenario of technology adoption; neither study compares *ex-ante* engineering models calibrated to estimate energy savings from the actual sample of households and technologies used for empirical estimates. *A-priori* engineering estimates have limited policy relevance; in fact, most ex-ante program evaluations use predictions from engineering models parameterized according the precise technology improvements and building characteristics as the true sample of adopters. Hence, a timely and policy-relevant need exists for empirical validations of household-level (or project-level) ex-ante engineering estimates of energy savings from retrofits compared against ex-post empirical estimates using billing data.

Empirical Comparison with Engineering Estimates

Table 1.5 compares my results with household-level engineering estimates collected by GRU. Ex-ante bias calculations report the percentage that engineering estimates deviate from difference-in-difference estimates derived from observational data.

Results confirm an upward bias in engineering estimates for three technologies previously studied in empirical literature: attic insulation, refrigerator replacement, and room air conditioners. However, on average, the magnitude of engineering bias is smaller than prior studies suggest. Engineers overestimate energy savings from attic insulation by 100%—a substantial bias—however, even this large bias appears modest compared to a bias of over 500% reported in past literature (Metcalf and Hassett, 1999). Similarly, engineers overestimate energy savings from refrigerator buybacks by 60%, a minor error compared to a bias of 250% reported for a similar program in Mexico (Davis et al., 2014).[5] While engineers estimate energy savings from room air conditioners, estimates in this study find no statistically significant or economically important energy savings, but do not go as far as to suggest that room air conditioner replacements increase energy consumption as reported in Mexico (Davis et al. 2014). It

[5] Some of the difference in engineering overestimates for refrigerator programs may be due to differences in specific program requirements in Mexico and Gainesville. In particular, the Mexican refrigerator replacement program issues rebates at the time of a new refrigerator purchase, while the GRU refrigerator buyback program only requires removal of an old refrigerator. Due to differences in refrigerator replacement requirements, and possibly in how engineering estimates are calculated, valid comparisons between these programs may be limited.

is worth noting, however, that comparison of the magnitude of engineering bias across studies is difficult. Such comparisons are complicated by idiosyncratic differences in climate, program regulations, and other confounding factors that complicate generalizations about the accuracy of engineering modes. Hence, when compared to previous literature, these estimates should be viewed as additional evidence rather than an improved assessment of engineering bias.

Engineering estimates are impressively accurate for retrofits novel to this study, although in some cases, engineering estimates are too conservative. Engineering estimates related to central air-conditioning systems are very accurate, with duct leak repair work achieving 87% of estimated energy savings, and super SEER air-conditioner replacements achieving 80% of estimated savings. Energy savings estimates for high-efficiency air-conditioner replacements are very conservative—*under* predicting energy savings by 60%. Similarly, engineers also *under* predict energy savings from pool-pumps by 22%. Ex-ante engineering models must forecast future weather conditions, which are inherently unpredictable; thus it is surprising to find accurate modeling of climate-dependent energy systems, such as air-cooling and pool circulation.

In sum, the accuracy of engineering models varies across retrofit programs, with a tendency to modestly over predict energy savings. However, inasmuch as this study measures actual performance of retrofits as implemented by residential energy users, the results of this study likely have more relevance for assessing the benefits of policies designed to foster residential energy retrofits than engineering studies.

Variation of Engineering Bias with Building Size

Despite strong empirical evidence that energy savings from retrofits increase with building size, engineering estimates fail to account for differences in energy savings between small and large buildings. Table 1.6 shows that houses larger than 2,000 square feet have nearly double the energy savings as houses smaller than 2,000 square feet. In most cases, engineering estimates vary only by 1 to 5 kWh per month between small and large houses, whereas empirical estimates vary by 60 to 120 kWh per month between small and large houses.

COST-EFFECTIVENESS OF REBATE PROGRAMS

Measuring the welfare effects from energy efficiency subsidies presents many challenges. In particular, such analyses require a framework to evaluate subsidy programs for energy efficient investments that properly accounts for program-induced benefits and costs.

The cost effectiveness results in this section do not constitute a rigorous evaluation of retrofit programs but, instead, provide a simple measure of cost effectiveness. This limited analysis only considers cost-effectiveness in terms of utility demand-side management objectives, as opposed to overall welfare effects. Furthermore, I impose several limiting assumptions that bias cost-effectiveness estimates. I assume 100% of energy savings observed by participating households is attributed to subsidy programs, a claim that implicitly assumes a counterfactual world without any energy efficiency investments. I also assume that subsidies are the only costs of operating programs, which ignores promotional campaigns and administrative costs that account for about half of total program expenditures. Together these assumptions will lead to an over estimate of energy savings and an under estimate of costs, effects that lead to a bias that underestimates the cost per kilowatt hour of electricity saved. Finally, to account for future energy savings, I also impose rigid assumptions about the lifetime of installed retrofits. Despite the limitations imposed by these assumptions, simple cost-based calculations reveal important insights not yet considered in the empirical literature evaluating the cost effectiveness of energy efficiency incentive programs.

Current program evaluation literature relies heavily on the econometric analysis of utility-level datasets. A review of econometric literature by Gillingham et al. (2009) highlights the wide range of Demand-Side Management (DSM) cost-effectiveness estimates reported in recent studies, with costs as low as 1 cent per kWh and as high as 26 cents per kWh. Research by Loughran and Kulick (2004) and Auffhammer et al. (2008) demonstrate the sensitivity of models using aggregate utility-level DSM spending and electricity sales data, as both studies analyze the same data but find substantially different estimates of program cost-effectiveness. A major pitfall of these studies is a reliance on utility-level data rather than household data, which aggregates effects over all types of retrofit rebates.

Given the wide menu of retrofit options—including air conditioner replacements, attic insulation upgrades, and refrigerator buybacks—information about the relative cost-effectiveness across different types of retrofits can guide policy makers to improve the efficacy of DSM rebate programs. This following simple analysis provides the first evaluation of product-specific cost effectiveness evaluation of energy efficiency retrofits.

Cost-Effectiveness of the Retrofit Rebate Program

The energy savings figures used to calculate cost effectiveness are based on estimates from essay 1. Table 1.7 presents basic calculations of costs per kilowatt hour saved based on utility rebates. Results suggest that program cost-effectiveness varies considerably across retrofit types. Following standard assumptions in the literature (see, e.g., Arimura, Li, Newell, and Palmer 2011), cost-effectiveness estimates assume a 5% discount rate and 10-year product lifetime.[6] Average cost savings per kilowatt hour range from 1 cent to 33 cents among programs based on rebate levels and energy savings estimated from electricity bills. Program-specific results include: refrigerator buyback (1 cent per kWh), pool pump replacement (2 cents per kWh), super SEER air-conditioner replacement (4 cents per kWh), high-efficient air-conditioner replacement (3 cents per kWh), duct leakage repair (4 cents per kWh), and attic insulation installment (7 cents per kWh), air-conditioner maintenance (3 cents per kWh), natural gas water heater conversion (3 cents per kWh), and room air-conditioner replacement (33 cents per kWh). Cost-effectiveness estimates change predictably under alternative assumptions. Assuming only a 5-year treatment effect, cost-effectiveness estimates increase by 78 percent.

Unlike previous studies, my results indicate that well-designed retrofit programs can achieve similar cost-effectiveness as other leading energy-efficiency programs. In Gainesville, a portfolio of rebate retrofits targeting refrigerator buybacks, pool pump replacements, super SEER central AC replacements, and duct leakage repair can achieve

[6] Davis et al. (2014) assume a 5-year treatment effect for an appliance retirement program, arguing that households naturally replace appliances after 5 years; rebates only prompt households to replace them earlier. Since GREEN rebates require participants to meet stringent energy-efficiency standards to qualify for rebates—standards that homeowners may not meet otherwise—the cost-effectiveness calculations used in this section assume a longer 10-year treatment effect suggested by others.

an energy savings at a cost between 1 and 5 cents per kilowatt hour avoided. By comparison, Allcott (2011) reports cost-effectiveness measures for peer-comparison programs from OPOWER to range from 2 to 5 cents per kilowatt hour saved. I am not aware of any other studies using household billing data that identify rebate programs with cost-effectiveness on par with OPOWER.

Private Benefits from Energy-Efficiency Investments

Homeowners also benefit from lower electricity bills. Based on GRU's block-rate pricing in 2009 and the distribution of monthly consumption levels observed in the data, the average marginal price that GRU charges consumers is 14.6¢. This consists of an 8.6¢ average energy charge plus a 6¢ fuel adjustment charge. These numbers imply that annual electricity savings per household vary considerably across retrofits: $86.65 for refrigerator buyback, $225.33 for pool pump replacement, $203.68 for super SEER air-conditioner replacement, $111.70 for duct leakage repair, $58.30 for attic insulation installment, and $187.50 for low-income weatherization. Programs without significant energy savings are cost-ineffective regardless of rebate levels.

Cost-effectiveness calculations suggest considerable heterogeneity in retrofit cost-effectiveness, an important finding missing from current literature. Energy savings per kilowatt hour range from 1 cent to 33 cents based on rebate levels and energy savings estimated from electricity bills, assuming a 10-year product lifetime and a 5 percent discount rate. These results highlight the contributions of even a basic analysis of product-specific cost-effectiveness, as they provide valuable information about optimizing utility conservation portfolios by eliminating inefficient programs and expanding efficient programs. The importance of heterogeneity across product-offerings within a utility suggests that estimates from analysis of aggregate utility-level data that dominate the current literature reflect only an average-cost effectiveness, which may not reflect the potential cost-effectiveness of subsidy programs that target best-practice products.

CONCLUSIONS

Despite the widespread implementation of retrofit rebate programs and calls for increased investment in demand side management programs, surprisingly little is known about whether energy-efficiency retrofits are an effective way to reduce energy consumption. Engineering simulations provide most of the evidence, but simulated predictions, even if based on sound models, do not account for installation quality or behavioral responses. Hence, there is an important and timely need for empirical research that uses field data to evaluate more fully the effects of energy-efficiency retrofits on energy consumption.

A major purpose of this essay is to evaluate retrofit-specific residential rebate programs based on observed household-level consumption data. My results identify the energy savings from nine retrofit rebate programs in Gainesville, Florida, using a panel dataset of electricity and natural gas consumption and building characteristics for 30,000 residences. The difference-in-difference method compares changes in energy use within a residence before and after an energy-saving retrofit intervention (the treatment group) with changes in energy use within a similar residence that did not receive improvements (the control group). A unique feature of the data, which is central to my identification strategy, is that I have monthly billing data combined with time-variant and time-constant characteristics of each residence.

This study makes three contributions to the literature assessing energy efficiency programs. First, this is the first assessment of a retrofit rebate program to apply difference-in-difference methods linking billing data and housing characteristics for every customer within a utility service area. Second, by assessing nine retrofit programs, this study explores heterogeneity across a diverse range of retrofit options using billing data. Third, this study is the first to use project-level data on engineering predictions and rebate-levels to identify retrofit-specific estimates of simulation bias and cost-effectiveness by use of a carefully matched control group.

The primary contribution of this first essay is the evaluation of retrofit-specific residential rebate programs based on actual billing data. I find that engineering simulations are surprisingly accurate when compared to empirical estimates, a stark contrast to previous studies. In four of nine retrofit programs, engineering models predicted energy savings within 20 percent of actual energy reductions. For the remaining five programs, engineering biases are modest relative to previous studies. Beyond providing new confidence in household engineering models, results also shed light on variation in engineering bias across retrofit types. This is information that could be used to improve the accuracy of future engineering models.

Additional results indicate that cost-effectiveness varies widely between 1 and 33 cents per kilowatt hour saved across retrofit programs. Seven out of nine programs achieve energy savings below 5 cents per kilowatt hour, roughly one-third of its cost of 18 cents per kilowatt hour for new electricity generation in Gainesville. Retrofits also benefit homeowners with reductions in annual electricity expenditures between $58 and $225.

These contributions provide new policy insights about the effectiveness and cost-effectiveness of specific retrofits. First, results inform policymakers about the relative efficacy of different retrofit rebates, allowing inefficient programs to be terminated and efficient programs to be expanded. Second, results provide new empirical evidence about the bias of engineering models, suggesting a need for future research to explain the cause of variation in model accuracy across retrofit types. Third, these evaluations empower homeowners to make informed decisions about energy-efficiency investments using credible information on the expected cost-savings from various retrofit options.

REFERENCES FOR ESSAY 1

Allcott, H. (2011), Social Norms and Energy Conservation. *Journal of Public Economics*, 95(9-10), 1982-1095.

Allcott, H. and M. Greenstone (2012). Is There an Energy Efficiency Gap? *Journal of Economic Perspectives*, 26(1), 3-28.

Allcott, H., and R. Sweeney (2014). Information Disclosure through Agents: Evidence from a Field Experiment," Working Paper 20048, National Bureau of Economic Research.

Angrist, J. and S. Pischke (2009), Mostly Harmless Econometrics: An Empiricist's Companion, Princeton University Press.

Auffhammer, M., C. Blumstein, and M. Fowlie (2008). Demand Side Management and Energy Efficiency Revisited. *Energy Journal,* 29(3), 91-104.

Boomhower, J. and L. Davis (2013), Free Riders and the High Cost of Energy-Efficient Subsidies, draft manuscript, University of California, Berkeley, March.

Crawford, M. and B. Fischer. (2011), Gainesville Regional Utilities with OPOWER, presented at the Home Energy Reporting Program DOE Technical Assistance Webinar, July 28.

Davis, L. W., A. Fuchs, and P. Gertler. (2014). Cash for Coolers: Evaluating a Large-scale Appliance Replacement Program in Mexico. *American Economic Journal: Economic Policy, 6*(4), 207-238.

Dubin, J., M., A., and D. McFadden. (1986). Price Effects of Energy-efficient Technologies: A Study of Residential Demand for Heating and Cooling. *Rand Journal of Economics*, 17, 310-25.

Gillingham, K. and K. Palmer (2014). Bridging the Energy Efficiency Gap: Policy Insights from Economic Theory and Empirical Evidence. *Review of Environmental Economics and Policy*, forthcoming.

Gillingham, K., G.N. Richard, and K. Palmer (2009). Energy Efficiency Economics and Policy. *Annual Review of Resource Economics,* 1, 14.1-14.23.

Jacobsen, G. D., and M.J. Kotchen. (2013), Are Building Codes Effective at Saving Energy? Evidence from Residential Billing Data in Florida. *Review of Economics and Statistics*, 95(1), 34-49.

Joskow, P.L. and D.B. Marron (1992), "What Does a MegaWatt Really Cost? Evidence from Utility Conservation Programs," *The Energy Journal*, 13, 41-73.

Lang, C. and M. Siler (2013). Engineering Estimates Versus Impact Evaluation of Energy Efficiency Projects: Regression Discontinuity Evidence from a Case Study, *Energy Policy*, 61, 360-370.

Loughran, D. S., and J. Kulick (2004). Demand-Side Management and Energy Efficiency in the United States. *The Energy Journal,* 25(1), 19-41.

Metcalf, G. E. and K.A. Hassett (1999). Measuring the Energy Savings from Home Improvement Investments: Evidence from Monthly Billing Data. *The Review of Economics and Statistics*, 81(3), 516-528.

US Green Building Council (2011). Green Building Research. Accessed on March 2, 2011 at http://www.usgbc.org/DisplayPage.aspx?CMSPageID=1718.

Von Schrader, C. (2011). I Got Your Deep Energy Savings Right Here, presentation at the Affordable Comfort Inc. Home Performance with Energy Star National Symposium, San Francisco, California, March 30.

TABLES FOR ESSAY 1

Table 1.1 Descriptive Statistics for Gainesville Residential Energy Efficiency Incentives

Rebate Program	Number of participants	Average subsidy	Average engineering estimates of energy savings (MWh per year)
Duct Leakage Repair	2,702	$359	1.3
Air Conditioning Maintenance	3,158	$ 55	0.5
Room Air Conditioner	812	$161	0.2
Super SEER Central Air Conditioner	792	$550	2.4
Attic Insulation	2,033	$361	1.6
Refrigerator Buyback (I)	3,131	$64	1.6
Refrigerator Buyback (II)	544	$ 51	1.2
Pool Pump (I)	797	$307	1.8
Pool Pump (II)	284	$251	1.8
Central Air Conditioner (I)	2,430	$549	1.9
Central Air Conditioner (II)	1,182	$294	0.5
Natural Gas Water Heater Conversion	1,657	$392	2.8
Total	25,558		

Descriptive statistics reported from census sample of rebate participants for each retrofit program. Average rebate amounts and average engineering estimates of energy savings by rebate program are calculated from project-level rebates and energy saving estimates.

Table 1.2 Estimates of the Electricity Savings from Energy Efficiency Investments

Variables	15-SEER Central Air Conditioner	Air Conditioner Maintenance	Attic Insulation	Duct Leakage Repair	Pool Pump	Super-SEER Central Air Conditioner	Gas Water Heater Conversion	Refrigerator Removal	Room Air Conditioner
Treatment Effect	-104***	-17.3**	-63.73***	-97.43***	-189.4***	-139.08***	-99.12**	-79.01***	-4.94
	(15.21)	(7.10)	(11.36)	(8.51)	(21.26)	(8.89)	(39.60)	(9.47)	(19.21)
Treated Houses	433	1,726	995	1,600	502	1,369	64	1,299	268

Each column is a separate model. All models include individual house fixed effects and billing period fixed effects. The dependent variable is household electricity consumption per month (in kilowatt hours). The variable *treatment effect* is a retrofit specific treatment indicator equal to 1 for billing months after retrofit installation for houses participating in that specific program and equal to 0 otherwise. Coefficients represent the monthly electricity use changes after retrofit installation. Numbers in parentheses report standard errors clustered by house. Asterisks denote statistical significance, where ** is significance at the 5% level and *** is significance at the 1% level. All models are estimated with a sample of 2,206,011 monthly electricity observations from 26,453 houses in Gainesville, Florida. All models report a constant term of 1,154.4 with a standard error of 4.04.

Table 1.3 Seasonal Estimates of the Electricity Savings from Energy Efficiency Investments

Variables	15-SEER Central Air Conditioner	Air Conditioner Maintenance	Attic Insulation	Duct Leakage Repair	Pool Pump	Super-SEER Central Air Conditioner	Gas Water Heater Conversion	Refrigerator Removal	Room Air Conditioner
Treatment Effect	-36.81**	-33.00***	-40.06***	-87.13***	-276.6***	-112.69***	-97.91**	-76.36***	31.27
	(16.57)	(7.47)	(11.75)	(8.94)	(22.02)	(9.85)	(47.97)	(10.23)	(20.38)
Treatment Effect * April to September	-128.3***	29.63***	-44.7***	-19.55	164.15***	-49.96***	-2.26	-4.99	-70.41
	(14.10)	(7.01)	(10.58)	(8.03)	(20.03)	(8.10)	(32.52)	(8.91)	(20.38)
Treated Houses	433	1,726	995	1,600	502	1,369	64	1,299	268

Each column is a separate model. All models include individual house fixed effects and billing period fixed effects. The dependent variable is household electricity consumption per month (in kilowatt hours). The variable *treatment effect* is a retrofit specific treatment indicator equal to 1 for billing months after retrofit installation for houses participating in that specific program and equal to 0 otherwise. The variable *April to September* is an indicator variable equal to 1 for billing months during the warm season between April and September, and equal to zero otherwise; the interaction term *(treatment effect * April to September)* represents the additional treatment effect during warmer months. Numbers in parentheses report standard errors clustered by house. Asterisks denote statistical significance, where ** is significance at the 5% level and *** is significance at the 1% level. All models are estimated with a sample of 2,206,011 monthly electricity observations from 26,453 houses in Gainesville, Florida. All models report a constant term of 1,154.4 with a standard error of 4.04.

Table 1.4 Building Size Influence on Estimates of the Electricity Savings from Energy Efficient Investments

Variables	15-SEER Central Air Conditioner	Air Conditioner Maintenance	Attic Insulation	Duct Leakage Repair	Pool Pump	Super-SEER Central Air Conditioner	Gas Water Heater Conversion	Refrigerator Removal	Room Air Conditioner
Treatment Effect	-83.44***	9.99	-34.24***	-59.73***	-93.39**	-100.13***	-83.97*	-58.18***	9.24
	(17.01)	(7.55)	(11.94)	(9.06)	(41.31)	(10.30)	(48.28)	(10.44)	(20.03)
Treatment Effect * > 2,000 Square Feet	-63.05*	-71.37***	-84.35***	-93.62***	-120.4**	-76.86***	-54.22	-60.84***	-100.38
	(34.48)	(15.26)	(26.00)	(18.14)	(47.85)	(17.27)	(82.11)	(21.13)	(61.74)
Treated Houses									
< 2,000 Square Feet	419	1,103	912	1,279	119	900	54	1,033	233
> 2,000 Square Feet	223	674	513	939	515	964	25	585	38

Each column is a separate model. All models include individual house fixed effects and billing period fixed effects. The dependent variable is household electricity consumption per month (in kilowatt hours). The variable *treatment effect* is a retrofit specific treatment indicator equal to 1 for billing months after retrofit installation for houses participating in that specific program and equal to 0 otherwise. The variable *>2,000 square feet* is an indicator variable equal to 1 for large houses with more than 2,000 square feet of conditioned living area, and equal to zero otherwise; the interaction term *(treatment effect * >2,000 square feet)* represents the added energy savings from treatment in large houses. Numbers in parentheses report standard errors clustered by house. Asterisks denote statistical significance, where ** is significance at the 5% level and *** is significance at the 1% level. All models are estimated with a sample of 2,206,011 monthly electricity observations from 26,453 houses in Gainesville, Florida. All models report a constant term of 1,154.4 with a standard error of 4.04.

Table 1.5 Biases in Engineering Estimates of Electricity Savings from Energy Efficient Investments

Retrofit Program	Empirical Estimate of Energy Savings (kWh per month)[1]	Engineering Estimate of Energy Savings (kWh per month)[2]	Engineering Bias[3]
15-SEER Central Air Conditioner	103.97***	41.67	-60%
Air Conditioning Maintenance	17.30**	37.50	117%
Attic Insulation	63.73***	129.17	103%
Duct Leakage Repair	97.43***	110.00	13%
Pool Pump	189.40***	147.25	-22%
Super-SEER Central Air Conditioner	139.08***	167.50	20%
Natural Gas Water Heater Conversion	99.12**	89.45	-10%
Refrigerator Removal	79.01***	126.67	60%
Room Air Conditioner	4.94	19.17	288%

[1]Empirical estimates represent coefficients from Table 1.2.

[2]Engineering estimates are calculated as an average of project-level ex-ante engineering estimates used for program evaluation by Gainesville Regional Utilities.

[3]Engineering bias is calculated as the percentage that engineering estimates deviate from empirical estimates.

Table 1.6 Comparisons of Empirical and Engineering Estimates of Energy Savings with Changes in House Size

	Houses smaller than 2,000 square feet		Houses larger than 2,000 square feet	
Retrofit Program	Empirical estimate of savings (kWh per month)	Engineering estimate of savings (kWh per month)	Empirical estimate of additional savings (kWh per month)	Engineering estimate of additional savings (kWh per month)
15-SEER Central Air Conditioner	-83.44***	-45.65	-63.05*	-0.05
Air Conditioning Maintenance	9.99	-37.50	-71.37***	0.00
Attic Insulation	-34.24***	-129.55	-84.35***	0.00
Duct Leakage Repair	-59.73***	-109.21	-93.62***	-1.18
Pool Pump	-93.39**	-145.67	-120.44**	-1.36
Super-SEER Central Air Conditioner	-100.13***	-166.70	-76.86***	-1.78
Natural Gas Water Heater Conversion	-83.97*	-82.72	-54.22	-6.62
Refrigerator Removal	-58.18***	-129.40	-60.84***	3.57
Room Air Conditioner	9.24	-19.50	-100.38	0.00

Empirical estimates in columns 2 and 4 represent coefficients from Table 1.4. Engineering estimates in columns 3 and 5 are calculated as an average of project-level ex-ante engineering estimates used for program evaluation by Gainesville Regional Utilities. Columns 4 and 5 represent the additional energy savings above estimates reported in columns 2 and 3.

Table 1.7 Cost Effectiveness of Electricity Saved from Gainesville Residential Energy Efficiency Incentives

Retrofit Program	Energy Savings (kWh/year)[1]	Lifetime Energy Savings (kWh)[2]	Average Subsidy Amount ($)	Cost Effectiveness ($/kWh)[3]
15-SEER Central Air Conditioner	1,248	10,116	$ 293	$ 0.03
Air Conditioning Maintenance	208	1,684	$ 55	$ 0.03
Attic Insulation	765	6,200	$ 433	$ 0.07
Duct Leakage Repair	1,169	9,480	$ 368	$ 0.04
Pool Pump	2,273	18,427	$ 283	$ 0.02
Super-SEER Central Air Conditioner	1,669	13,532	$ 556	$ 0.04
Natural Gas Conversion	1,189	9,644	$ 292	$ 0.03
Refrigerator Removal	948	7,687	$ 72	$ 0.01
Room Air Conditioner	59	481	$ 159	$ 0.33

[1]Energy savings estimates are taken from Table 1.2.

[2]Lifetime electricity savings assume a constant 10-year treatment effect with all future electricity savings adjusted with a 5% annual discount rate.

[3]Cost-effectiveness values are calculated from average subsidy amount divided by discounted lifetime energy savings. Calculations assume no free-riders and no costs incurred from program administration or program advertisement campaigns.

ESSAY 2 : **DO ENERGY SAVINGS GROW ON (SHADE) TREES?**

Buildings account for 42 percent of energy use and 38 percent of CO_2 emissions in the United States, making building energy efficiency a key component of broader energy and climate goals (US Green Building Council, 2011). In recent years, state and federal governments have adopted stringent building code regulations and invested billions of dollars into demand side management (DSM) programs that subsidize energy-saving building improvements for homes and businesses. This surge in public investment has spurred a renewed interest in evaluations conventional policies, such as building codes, DSM policies, and behavioral interventions; with studies typically showing these conventional approaches have limited demand-reduction effects and fail cost-benefit tests (e.g. Jacobsen and Kotchen 2012; Davis et al., 2014; Alcott, 2011). Despite the well-established limitations of conventional DSM approaches, empirical research continues to overlook the role of environmental factors known to affect energy consumption. For example, sun exposure, a key environmental factor determining building energy performance, drives numerous engineering decisions about building design, such as building orientation, roof reflectance, heat ventilation, and shade.

Considering the widespread role that environmental factors play in energy efficiency, surprisingly little empirical research considers how the environmental factors, such as green infrastructure, can play a role in energy-conservation. Trees shade buildings during summer months, reducing air-conditioning usage, and can also serve as windbreaks during winter months, reducing heating requirements for buildings. Today, more than 3,400 U.S. cities have tree-ordinances to protect urban tree canopies, which may offer effective DSM policies if shade trees and windbreak trees could substantially improve building energy efficiency (Tree City USA, 2015).

Conventional wisdom suggests that trees have little effect on consumer electricity consumption; however, these studies fail to demonstrate any causal connections. Current empirical research suggests a correlation between tree cover and slightly lower summer energy consumption, but fails to find effects large enough to attract policy attention. Donovan and Butry (2009) use monthly billing data from 460 homes in Sacramento, California, and find that homes with south and west facing trees have lower summertime energy bills. Pandit and Laband (2010) present a similar model for 160 homes in Auburn, Alabama, and find that a 20 percent increase in tree shade reduces summertime electricity bills between 3 and 9 percent, but also substantially increases winter electricity bills.

These studies identify effects based on cross-sectional comparisons of energy use between houses with varying levels of tree shade. Both papers include only basic controls for home characteristics (house square footage and presence of a pool) and rely almost entirely on average early-spring energy to serve as a proxy of baseline energy use without shade. However, these cross-sectional estimators will be biased if tree size is correlated with other variables that affect energy efficiency, such as home age. Pandit and Laband (2010) also require participant permission to read electricity meters, introducing selection bias into the estimates. Endogeneity problems are also likely if homeowners who prefer trees also practice energy conservation.

This essay presents the first causal evidence that trees have a major impact on consumer demand—with large shade trees reducing household electricity use by more than 20 percent. My work contributes to the existing literature on the energy saving potential of urban forests in four ways. First, I use a quasi-experimental design to identify a causal link between tree shade and energy use. Estimates are identified by electricity variation "within" households, reducing omitted variable bias and endogeneity problems of cross-sectional "between" estimators. Second, my study draws from a full census of 30,000 households within the GRU service area, improving the consistency of estimates compared to previous studies with small samples and selection biases. Third, given the context of a city-wide tree ordinance, these estimates have direct policy relevance. Fourth, my estimates suggest that energy savings from tree shade coincide with seasons of peak electricity demand, providing new evidence that tree ordinances may serve as effective demand-side management policies. Results suggest that the energy savings from tree shade are on an order of magnitude greater than other energy-efficiency policy measures, such as a 2006 Florida state law increasing the stringency of building codes (Jacobsen and Kotchen, 2013).

EMPIRICAL SETTING AND DATA COLLECTION

Over the past decade of growing energy demand, Gainesville, Florida, has prioritized demand-side management programs designed to avert or delay major infrastructure upgrades. Since 2005, Gainesville Regional Utilities (GRU) has invested $27.8 million in energy conservation initiatives, including subsidies for energy-efficient retrofits and incentives for solar technology adoption (GRU, 2008). Despite aggressive efforts to slow demand growth, GRU's existing infrastructure will be unable to supply Gainesville's energy needs in the near-future. As a result, GRU commissioned a third-party to build and operate a 100-MW power plant that became operational in 2013.

Under the 30-year contract, GRU must purchase all electricity output from the facility—a cost of $3.1 billion over the contract lifetime. Ultimately, these costs will be passed onto consumers: average electricity prices already increased by 1.9 cents per kWh in 2013, for an increase of 15 percent in one year.[7] The high cost of supply-side solutions has spurred new interest in cost-effective demand-side solutions—specifically, policies that reduce energy demand more cheaply than the marginal cost of increasing energy supply.

In contrast, tree ordinances are not designed for demand-side management; rather the central aim of tree ordinances is to preserve the character, heritage, and beauty of Gainesville. Removing trees on private property has been regulated in Gainesville under a city-wide tree ordinance since 1999.[8] Like tree ordinances in hundreds of other U.S. cities, Gainesville's ordinance requires residents to apply for permits approved by a city arborist prior to the removal of any mature tree within the city limits (Treiman and Gartner, 2004; Kielbaso, 1989; U.S. Conference of Mayors, 2008).[9] Tree ordinance measures are strictly enforced and apply to all large trees on private property.[10] Since the energy savings from trees are largely anecdotal, the Gainesville tree ordinance is not

[7] This calculations compares September 2012 and October 2013 based on GRU electricity rates. Average rates are total utility bills divided by total electricity consumption for a household consuming 700 kWh per month. Monthly bills were $88.36 in 2012 and $101.34 in 2013. A majority of the utility bill increase is due to an additional 2-cent per kWh fuel surcharge directly attributed to capital costs of the new power plant.

[8] Tree removal permitting regulations established under Gainesville City Ordinance Sec. 30-254 (Gainesville Land Development Code, Article VIII). A summary of regulations is available online at http://www.cityofgainesville.org/Portals/0/nod/UF-TreeRemovalRegulations-1.pdf.

[9] The exact number of U.S. cities that protect trees on private property is unknown. Multiple surveys confirm that over 130 cities have private property regulations, no comprehensive survey of all cities with tree ordinances exists. Two surveys suggest that private property restrictions apply to about 13% of U.S. cities (Treiman and Gartner 2004, Kielbaso 1989), while an informal survey of tree ordinances in 135 large cities reports that 68% of tree protection laws cover private properties (U.S. Conference of Mayors, 2008).

[10] During the period of this study, permits were required for any tree over 20-inches in diameter for single-family residences, and 8-inches in diameter for multi-family residences and planned single-family developments. However, regulations began at a 30-inch diameter for four specific tree species: Water Oaks, Laurel Oaks, Sweetgums, and Loblolly Pines.

recognized as part of city energy-efficiency efforts. In fact, like most cities, the Gainesville tree ordinance is void of any energy-conservation criteria to assess eligibility for tree removals or mitigation plantings. Moreover, Gainesville lacks any system to monitor energy savings from tree-related programs, which is information required for all traditional demand-side management programs operated by GRU.

The essay focuses on how changes in residential tree shade translate into changes in actual energy consumption. In particular, this essay focuses on changes in tree shade in a region of the southern United States with a humid tropical climate. The study uses residential utility data for households in the city of Gainesville, Florida. The data were downloaded from Gainesville-Green.com, which is a website designed to encourage energy conservation through the provision of publicly-available information on household energy consumption. The dataset includes monthly billing records for electricity and natural gas consumption for residential households. Monthly billing data spans from 2000 through 2012. Also included in the billing data are residence-specific identifier codes that correspond to unique street addresses. Based on street address, billing data are linked to several other datasets, including datasets of housing characteristics, renovation projects, and tree removals.

Tree Removal Permit Data

The data on tree removal permits are of primary interest because they record which residences have removed trees and, in particular, when the trees were removed. Information on the location and timing of tree removals, combined with billing data, allow a comparison of energy consumption before and after a change in tree cover. Permit data were acquired from the Gainesville city arborist, who enforces ordinances by visiting the residence of each applicant to assess whether trees are eligible for removal. The dataset includes the street address of residents who were issued permits, the number of trees approved for removal by each permit, and the issue date for each permit.

The key characteristic is the issue date because it can be used to determine when a tree removal occurred and, in particular, whether a monthly energy bill is subject to shade conditions before or after the change in tree canopy.

Using permit issue dates to categorize monthly energy bills as occurring pre- or post-tree removal, one must also take account of the fact that permits are issued prior to tree removal and not on the actual date a tree is removed. According to tree ordinances, a tree removal permit is valid for a six-month period after the issue date, after which the permit expires. Thus, for tree-removal residences, monthly energy bills are categorized as pre- or post-tree removal, as follows: a billing-month earlier than the issue date designates a billing record as pre-tree removal; a billing-month occurring six-months after the issue date or later designates a bill as post-tree removal; and a billing-month that occurred less than six-months after the issue date designates a bill as indeterminate because the corresponding tree-cover of the residence is unclear. Thus, all bills occurring within the six-month period following a permit issue date are dropped from the analysis.

While the complete set of permits spans from 1999 through 2012, the analysis only includes permits issued from 2001 through 2011, the period that includes tree-removing residences with a full 12-months of billing data both before and after the permit issue date. Importantly, this essay's empirical strategy is based on a comparison of energy bills before and after a change in tree cover. The best comparison is based on residences with adequate billing information in both tree cover regimes. Thus, for this analysis, treated residents were excluded if they had fewer than 12-monthly billing observations in either the pre-permit period or the post-permit period.[11] Even without this restriction, most treated residences have multiple years of pre-and-post treatment observations for each season. For example, the median treatment date of June 2007 includes 7.5 years (90 months) of billing data in the pre-permit period, and 4.5 years (54 months) of billing data in the post-permit period. Among the final treatment sample, 94-percent of residences have more than 18-months of billing data after the permit issue date, which guarantees at least one observation in each month-of-the-year for both the pre-treatment period and the post-treatment period.

To quantify the size and location of removed trees, an additional dataset is created from aerial imagery. Using two sets of 1-foot resolution aerial imagery taken in 2001 and 2011, a map of tree canopy loss was generated over 7,417 properties, including all tree removal sites and adjacent residences.[12] The resulting map classifies changes in tree cover during the 10-year period between 2001 and 2011, with tree canopy area mapped at a 1-meter resolution. Combining the tree loss map with data on property boundaries, provided by the Alachua County Property Appraiser, the owners of each tree canopy area are then identified.

[11] Results are robust to an alternate treatment definition that requires 24-monthly billing observations in both the periods leading up to and following the permit issue date.
[12] See Appendix 2.1 for additional details about input data and methods used to classify tree canopy loss.

Sun Position and Tree Shade

Additional data are also considered to measure the loss of shade associated with a tree removal. Tree shade can be determined from three metrics: tree size, tree location, and sun position. The first two factors—tree size and location—are quantified by mapping tree loss from aerial imagery. Specifically, a map of tree loss is created at the 1-meter pixel resolution.

In addition to quantifying tree position, a full characterization of shade coverage must also consider information about sun position. Sun position is defined by two angles: aspect angle (north, east, south and west) determines shadow direction; and altitude angle (0 to 90 degrees) determines shadow length. Sun movement describes temporal changes the sun position, which varies based on time-of-day and by season-of-year. Short shadows occur when sun reaches high altitudes, which reach a daily-maximum at solar noon, and an annual-maximum during the summer solstice. Conversely, longer shadows occur at low altitude angles, which occur near sunrise and sunset, and during winter months.

Figure 2.2 is a three-dimensional diagram of seasonal sun paths over Gainesville that conveys temporal changes in sun position relative to a house.[13] During the equinox, marking the first days of spring and autumn, the sun rises due east, sets due west, and arcs across the south reaching a maximum altitude of 60-degrees (the center arch in figure 2.2). On the winter solstice, with a brief 10-hours between sunrise (southeast) and sunset (southwest), the sun traverses the southern horizon at a low-altitude, reaching a

[13] The geographic coordinates at the center of Gainesville, Florida, are used to simulate sun paths and shadow position. Gainesville is located at latitude 29°39.0978′ N and longitude 82°19.4898′ W.

maximum angle of 35-degrees (the right arch in figure 2.2). In contrast, the summer solstice marks a 14-hour day, with sunrise in the northeast, sunrise in the northwest, and a mid-day sun passing overhead at a maximum altitude of 84-degrees (the left arch in figure 2.2).

Figure 2.3 illustrates the shadows of a 50-foot tree at various times of the day and various seasons. In the summer, western trees set back as far as 75 feet provide cooling shade in late afternoon, leading to our prior expectation that west-side trees may reduce summer electricity use from air conditioning. Interestingly, despite providing cooling shade during mid-afternoon, south-side trees may only reduce electricity consumption at distances of less than 25-feet from a house. This follows because south-side trees cast very short shadows during the Gainesville summer, as the mid-afternoon sun is positioned overhead at a high angle in the sky. In the winter, on the other hand, south-side trees provide almost all shading. In fact, due to the low altitude sun position, south-side trees set back as far as 100-feet can provide cooling shade that may increase heating demand for homes. Since most homes use natural gas during the winter, one could expect that south-side trees may increase natural gas used for heating.

Mapping Canopy Loss

Given that the amount of shade a tree provides depends upon its location, it is critical to develop a measure of tree cover loss that encompasses both tree size and tree position around houses. To do so, data from aerial imagery is used to construct a single measure of tree cover change for the decade-epoch spanning 2001-2011. The resulting dataset traces, at a spatial resolution of 1 meter by 1 meter, the patterns of tree cover loss across Gainesville that spans the time period of permit data for tree removals. This section describes how the tree cover change dataset is constructed from raw aerial images.

There are two main challenges in constructing aerial images of tree loss. First, urban regions like Gainesville are composed of tightly interwoven land-uses, which are peppered with low-density tree cover. This makes it difficult to use satellite sensors, like Landsat, which have been used to measure annual tree loss in higher-density tree cover landscapes. Since these satellites have a coarse spatial resolution of 30 meters by 30 meters, satellite-derived images are too blurry to distinguish individual trees, measure canopy size, or identify tree location on residential properties. An alternative is to draw on sensors attached to low-flying aircraft, which are infrequently commissioned to photograph cities at an ultra-fine spatial resolution as detailed as 1 foot by 1 foot. This aerial imagery dramatic increases the spatial resolution and accuracy of classifications, but identifies tree loss over a longer 10-year interval instead of the annual time-interval possible with Landsat.

To generate the data used in this essay, two sets of aerial imagery are used as inputs. Images from 2001, used to characterize initial tree cover before tree removals, have a spatial resolution of 1 foot by 1 foot and include spectral signals for three bands, which are collectively known as the "visible bands" that can be seen with the human eye (blue [459-479 nm], green [545-565 nm], red [620-670 nm]). Images from 2011, used to characterize final tree cover after tree removals, also have a spatial resolution of 1 foot by 1 foot and include spectral signals for four bands, including the three visible bands, and a fourth band of near-infrared signals [841-876 nm]. Images were acquired separately, with the 2001 data purchased from the office of the Alachua County Property Assessor

(ACPA), and the 2011 data was provided by Florida State University. Both data sets were collected during summer months, during leaf-on seasons. Despite attempts to enhance 2001 images with an available band of 1999 near-infrared imagery, the data fusion was unsuccessful because the 1999 data were collected during a leaf-off season. Auxiliary data available on the ACPA website, including GIS building footprints and property boundaries, were used to define a study area and guide classifications.

Combining aerial images across time presents several challenges. Unlike satellites, which collect data over time with the same sensor, aircraft often collect data with different sensors over time, and these instrument changes introduce spectral differences in images across time. Moreover, spectral differences can also occur when images are collected from different altitudes, lens angles, or time-of-day images. Image post-processing corrects most distortions caused by flight patterns, but comparisons between aerial imagery are not consistently reliable at a spatial resolution of 1 foot by 1 foot. To enable comparison of spectral signals across time, both datasets are aggregated from a resolution of 1 foot by 1 to a resolution of 1 meter by 1 meter. Aggregating spatial resolution, in this case by an order of magnitude, has been shown to minimize error in spectral signals and improve the accuracy of change (Wulder et al., 2000). Despite loss of spatial detail, a final resolution of 1 meter by 1 meter permits a suitably precise measure of tree cover area and tree position.

For each 1 meter by 1 meter observation, a total of 129 image inputs is used (43 individual 1 foot by 1 foot single-band spectral observations for each year, plus their calculated differences). This amount of information, which adds 129 dimensions for each 1 meter by 1 meter pixel, is used to estimate a likelihood of tree cover loss for each pixel.

The next step is to take aerial image inputs and implement a computer algorithm to discriminate between tree and non-tree land cover. The purpose of remote sensing is to develop an algorithm that identifies what spectral signatures or set of signatures—in other words, what combinations of spectral and temporal information—best discriminate tree cover and its loss. For example, plants strongly absorb wavelengths of light in the visible spectrum, but also strongly reflect wavelengths in the near-infrared spectrum. Aerial sensors, which measure light reflected from the ground, will record weak visible (red) signals and strong near-infrared signals from green plant cover. One common metric for determining the density of green vegetation on a patch of land is the normalized difference vegetation index (NDVI), which captures differences in the reflectance of near-infrared and visible (red) light spectrum, and is a useful spectral signature for discerning the presence of vegetation.

In practice, however, one can do much better than using the NDVI alone by exploiting additional dimensions of the data. For example, trees have different spectral signatures than other types of green vegetation. These metrics help discriminate between trees and other bright green vegetation—such as lawns, bushes, and gardens—which remain green year-round in Florida. Several methods for change detection can exploit multi-dimensional data, including the popular Classification and Regression Tree; however, such objective methods are most effective when using two images recoded from the same sensor. For classification of aerial imagery taken at different altitudes and angles using different sensors, subjective methods like the pixel-based change detection method and unsupervised classification permit greater flexibility in identifying signals of land use change (see the Appendix for details on selecting the change detection method).

To take maximal advantage of the richness of spectral signals in the imagery data, a statistical learning procedure known as the iterative self-organizing data analysis technique (ISODATA) is used to identify tree loss. This is a change detection model used to determine which spectral signals best correspond to tree cover and its loss (Tou and Gonzalez 1974). The resulting map classifies changes in tree cover during the 10-year period between 2001 and 2011, with the tree canopy area mapped at a 1-meter resolution.

Shade Intensity Variables

After mapping tree canopy loss, some additional data analysis is needed prior to defining variables of treatment intensity. For the identification strategy employed in this essay, areas of tree canopy loss require details about the date of tree removal, and the tree location relative to nearby buildings—details that require auxiliary datasets. To assign a specific date of tree loss, each canopy area is linked to city tree removal permits based on the property address. As a preliminary step, a property address was first assigned to each area of canopy loss based on GIS property boundary data provided by the office of the ACPA. To characterize tree position around nearby buildings, two variables are calculated: the minimum distance between each tree and the house, and the aspect-angle (north, east, south, west) from the house to each tree. As a location reference for calculations, a GIS map of building footprints was downloaded from the ACPA website. Figure 2.3 provides a visual illustration of the three main GIS datasets, including tree canopy loss, property boundaries, and building footprints. Thus, maps of tree loss and building locations characterize tree position using two metrics: distance from house, and aspect angle to house.

To quantify the position of tree loss, each tree removal is grouped into one of sixteen regions around each house, as shown in Figure 2.4. For example, the northern quadrant canopy-loss variables measure the loss area within 25-feet of a house (N-1), loss area within 50-feet of a house (N-1,2), and loss area within 100-feet of a house (N-1,2,3,4). Distance cutoffs are identical for remaining quadrants to the east, south, and west of houses. If a tree canopy spans two or more regions, then the entire canopy area is assigned to the region nearest to the house. For example, in Figure 2.1, the dark green tree is assigned to region N-1. These canopy-loss variables permit flexible tests of how energy savings vary with differences in the canopy area, proximity to residences, and aspect-angle relative to residences.

To simplify interference, sun path information is use to define three levels of change in shade-intensity caused by tree-removals: heavy shade loss, light shade loss, and no shade loss. Figure 2.3 panel (b) illustrates how trees are assigned shade-intensity based on position around a house. Heavy shade trees are positioned within 50 feet of a house along the south, west, or east sides of the house (S-1,2; W-1,2; E-1,2). Light shade trees are set back further than 50 feet from the house on the south, west, or east sides; or within 25 feet of the house on the north side (S-3,4; W-3,4; E-3,4; N-1). Non-shade trees are on the north side of the house set back further than 25 feet from the house (N-2,3,4).

To complement shade-intensity variables, a less-parametric grouping of tree position and size helps to identify marginal treatment effects, such as increasing the area of canopy removed. For each house, twelve canopy-area variables are calculated. These correspond to the area of canopy loss, measured in square meters, contained within each of the sixteen regions around the house.

Matching and Sample Selection

Additional datasets are used to select appropriate control residences through covariate matching methods. Coarsened exact matching is used to identify a single control home for each treatment home based on housing characteristics and baseline energy use patterns (Blackwell et al., 2010). Matching serves to restrict the sample to improve estimate precision. If this ideal is met, then the tree-removal effect can be measured without bias. This is because controls estimate counterfactual outcomes for what the treated energy consumption would have been absent the tree removal treatment.

The goal of covariate matching is to select observable characteristics so that any two residents with the same value for these characteristics will display homogeneous responses to the tree-removal treatment. Ideally, matched treatment and control homes should be "statistical twins" concerning time-varying characteristics related to energy consumption and pre-treatment baseline energy consumption. To achieve this ideal, pre-treatment data on energy consumption are used to match houses with similar energy consumption profiles. Since the earliest tree removals occur in year 2001, baseline electricity and natural gas data from year 2000 are used to match homes with similar levels of summer electricity consumption and winter natural gas consumption. In addition, houses are matched on the increase in electricity consumption between spring and summer, which serves as a proxy for air-conditioning demand.

Time invariant structural characteristics correlated with unobserved time-varying characteristics are also useful matching parameters. The primary motivation for matching in the context of this study is to control for unobserved time-varying characteristics that affect the energy-savings from tree removals. A priori, energy savings from tree shade are presumed to be dominated by climate-related energy consumption. Thus, treatment and control houses should share characteristics that affect climate-related responses in energy consumption. In this study, structural characteristics are used to match houses with similar heating and cooling requirements and construction quality. Data come from the ACPA property sale database, and include information on house age, square footage, number of stories, number of bedrooms, type of air-conditioning system, heating fuel, and type of roof. Building age is correlated with building energy-efficiency and with the age and size of neighborhood trees. Characteristics of house size are correlated with the heating and cooling requirements of a house. Air conditioning (represented by a central air-conditioning dummy) and heating fuel (electricity, gas, or oil) relate to seasonal variation in electricity and natural gas consumption, while type of roof (shingle, tile, wood, metal) may affect heat transfer from solar radiation.

Exact matching is possible for categorical variables, such as heating fuel type and roof type. For continuous variables, such as energy consumption and square footage, exact matching is applied to coarsened data. In this analysis, each continuous variable is disaggregated into 10-strata that correspond to deciles in the distribution of the treated residences. Exact matching is then applied to minimize the multivariate distance between two residences across all coarsened variables.

To select controls based on a coarsened exact matching distance measure, the method of nearest neighbor matching is most common. In particular, 1:1 nearest neighbor matching selects for each tree-removal resident *i* the control resident with the smallest distance from resident *i*. One concern of 1:1 matching is the possible reduction in power from discarding unmatched observations. However, the reduction in power is minimal for two reasons. First, in a two-sample comparison of means, the precision is driven by the smaller group size. Because the treatment group stays the same size, and only the control group decreases in size, the overall power should not be greatly reduced (Ho et al. 2007). Second, power increases when groups are more similar because of higher precision and reduced extrapolation (Snedecor and Cochran, 1980). Given the vast number of control residents available in the sample, matching is conducted without replacement and without imposing calipers, simplifications that also lead to easier interpretation of effects. For estimation, matches are pooled into matched treated and control groups and analyses use the groups as a whole, rather than individual matched pairs.

Estimating effects without bias requires that, given a vector of covariates, the non-treated outcomes are what the treated outcomes would have been had they not been treated (i.e. tree-removal is independent of changes in energy use for "similar" residents). Internal validity of estimates relies on the assumption that, after matching on observable characteristics, tree removal is as-good-as random across the sample. One potential concern is self-selection bias, which could arise if all residents have the option to remove trees, and if those residents electing to remove trees differ in terms of potential energy-use outcomes from other residents.

In this study, tree ordinance laws minimize self-selection concerns because, for most residents, removing trees is not a legal option. In particular, the city arborist only approves permits for trees that are dying, diseased, or damaging property; the arborist denies permits for unnecessary tree removals. The permit vetting process, enforced through site visits, prevents residents from electing to remove trees that are healthy and innocuous. Thus, residents are not in control of whether, or when, they can remove a tree. This follows because the timing of tree death, disease and property damage is unpredictable and determined by natural processes.

In practice, the residents obey tree removal ordinances with near-perfect compliance, as regulations are well-known and permit violations result in expensive penalties. For example, removal of a non-permitted tree with a 30-inch diameter trunk would incur a penalty of $2,500—a cost that is an order of magnitude greater than the same tree removal with an approved permit.[14] To further protect healthy tree canopies, ordinances also prohibit actions that degrade tree health, such as over-pruning or tree-top removal.

EMPIRICAL ANALYSIS

City permits for tree removals, combined with Gainesville data on residential characteristics and utility consumption, provide an opportunity to examine the effect of tree shade on actual electricity and natural gas consumption. This section describes the empirical strategy and results. The first analysis applies a difference-in-difference model to compare changes in energy consumption between residences with and without tree removals. The second analysis tests whether or not the effect of tree removals is greater for trees positioned to provide building shade. The third analysis applies a continuous

[14] Tree ordinance rules mandate that an eligible tree removed without a permit must be replaced on an inch-by-inch basis with 3-inch diameter nursery trees, which cost approximately $250 apiece. Therefore, the example of a 30-inch diameter tree, if removed without a permit, would require on-site mitigation planting of 10 trees with 3-inch diameters, at a total cost of $2,500. In contrast, removal of any size tree with a valid permit only requires on-site planting of two trees with 1.5-inch diameter, at a total cost of approximately $250.

difference-in-differences model to examine how the effect of tree removals increases with tree size.

Because tree shade only affects energy consumption related to space cooling or space heating, the effect of tree removal, if it exists, will be greater during seasons when cooling and heating make up a larger share of a household's energy demand. For electricity, the effect of tree removals should be greatest during summer months when electricity demand for air conditioning is at its peak, and when solar radiation is most intense. For natural gas, the effect of the code change is expected to be greatest in winter months when demand for natural gas-based heating is at its peak. Given the tropical climate of Gainesville, coupled with longer, more intense sun exposure during summer months, tree shade should have a greater effect on summer electricity demand, and a smaller effect on winter natural gas demand.

Tree Removal Treatment

The average effect of tree removal on energy consumption can be estimated without using information about tree size and location relative to a house. I begin with this simple approach to establish a baseline treatment effect. Using only tree removal permits and billing data, I estimate difference-in-differences models for both electricity and natural gas of the form

$$Y_{it} = \tau w_{it} + \lambda_t + c_i + \varepsilon_{it} \tag{1}$$

where the dependent variable Y_{it} is either monthly electricity consumption (kWh) or monthly natural gas consumption (therms), depending on the model; i indexes residences; t indexes the month-year of the billing record; w_{it} is an indicator variable equal to one for billing months after a tree removal for those residences that remove a tree; λ_t is a month-year specific intercept that controls from month-to-month effects common to all residences, such as weather fluctuations or changes in the price of electricity or natural gas; c_i is a residence-specific intercept; and ε_{it} is assumed to be a normally distributed error term.

Equation (1) is estimated with a fixed-effects estimator and clustered standard errors at the residence level. The coefficients of interest, contained in the vector τ, are difference-in-differences estimates of how changes in energy consumption differ between residents with tree removals and residents without tree removals. If trees removals do, in fact, affect energy consumption, then, around the timing of tree removal, residences should experience changes in energy consumption that differ from changes experienced by other residences. With respect to electricity, if tree shade has the expected effect, residences should have increased electricity consumption after a tree removal. This follows because air conditioning demand, the primary determinant of residential electricity usage, should increase as additional summer sun exposure heats homes after a tree removal. With respect to natural gas, after-tree-removal residences should have decreased natural gas consumption. This follows because heating demand, the primary residential end-use of natural gas, should decrease as additional winter sun exposure naturally warms homes after a tree removal.

Table 2.1 reports the estimates of specification (1) for electricity and natural gas (columns 1 and 2). Focusing first on electricity, results suggest that after a tree removal, households consume approximately 29 kWh per month more than households without tree removals. The result is statistically significant at the 1-percent level. In terms of percentage difference, the estimates suggest that the average tree removal results in a 2.5-percent increase in residential electricity consumption. Focusing instead on the electricity saved by trees, these findings suggest that a residential tree in Gainesville reduces household electricity consumption by almost 2.5-percent, assuming that trees targeted for removal provide shade that is similar to other residential trees.

As expected, results for natural gas have a negative coefficient, consistent with the hypothesis that winter heating demand decreases after a loss of tree shade. Although the negative sign is expected, the coefficient estimate for tree removals is not different from zero at any conventional level of statistical significance. Furthermore, the natural gas effect is smaller in magnitude than the electricity effect, as the natural gas estimate suggests that tree removal households consume 0.5 therms per month less than households without tree removals, a change in energy equivalent to about 15 kWh. Hence, a comparison of electricity and natural gas coefficients is consistent with the expectation that, in Gainesville, tree shade has a greater effect on summer electricity demand than winter natural gas demand.

Tree Shade Treatment

While the models presented in Table 2.1 provide an estimate of the energy effects for average tree removal, the effects may differ in important ways depending on tree size and tree location. In fact, the estimates from specification (1) should be interpreted as a lower-bound of the true energy savings from tree shade. This follows because tree removal is an imperfect proxy for change in tree shade, which contains measurement error. For example, some trees are not originally positioned to shade houses, and these no-shade tree removals would not affect energy use. Such measurement error in the treatment should, in theory, attenuate estimates towards zero and under-predict the true effect of shade on energy use. The magnitude of this bias will depend upon whether or not the average tree was positioned to shade a home prior to removal. To address these concerns subsequent models account for heterogeneity in shade changes across tree removals to achieve more precise estimates of the energy savings from tree shade.

As discussed previously, tree removal permits are combined with aerial data on tree cover change and building location to determine how much shade a tree provided the residence before a tree removal. For simplicity, based on tree position relative to buildings, tree removals are assigned to one of three groups: heavy-shade, light-shade, or no-shade. With these definitions of shade-intensity, I estimate expanded versions of specification (1) for electricity and natural gas as follows:

$$Y_{it} = w_{it}[Heavy_i,\ Light_i,\ NoShade_i]\tau + \lambda_t + c_i + \varepsilon_{it} \qquad (2)$$

where the indicator variable for billing months after a property-specific tree removal is interacted with each of the shade-intensity variables. Following the same estimation strategy, equation (2) is estimated with a fixed-effects estimator, and standard errors are clustered at the residence level.

The parameters of interest, contained in the vector τ, are shade-intensity treatment effect coefficients that represent the average monthly energy change caused by the loss of shade from a tree removal. Earlier results and intuition can guide expectations about the direction and magnitude of energy across shade-intensities. With respect to electricity, if the loss of tree shade is a key driver of tree removal effects, one would expect greater increases in electricity consumption for residences that remove trees providing heavy shade on a house. This follows because air conditioning, which is used more intensively with higher outdoor temperatures, should increase demand as residences experience a greater loss of house shade. With respect to natural gas, previous results are not statistically significant and provide limited guidance; however, intuition suggests that natural gas heating, which is used less intensively with higher indoor temperatures, should decrease in demand as houses experience a greater loss of house shade. This follows because natural gas heating, which is used more intensively with lower indoor temperatures, should decrease in demand with a greater loss of house shade; however, because sun exposure is less intense during winter months, and some trees drop foliage, shade loss should affect natural gas consumption less than summer electricity consumption.

The first column of Table 2.2 reports the electricity estimates of specification (2). Focusing on the interacted shade intensity variables, which identify the amount of shade loss experienced by residences, coefficients have the expected signs and expected pattern of effects that increase with shade intensity. Electricity consumption is increasing with the intensity of shade lost, and these results are consistent with electricity being the primary energy source for cooling of Florida residences.

While the sign and trend of shade intensity coefficients are predictable, the magnitude of effects is unexpected. Surprisingly, based on the interaction with $Heavy$ shade, estimates suggest that electricity consumption increases dramatically after residences remove heavy-shade trees. In particular, residences increase electricity consumption by an average of 203 kWh per month after the removal of a heavy-shade tree, a result that is significant at the 1-percent level. In terms of percentage difference, the estimates suggest that removal of a heavy-shade tree results in a 17.7-percent increase in residential electricity consumption. The coefficient for heavy shade tree removals represents the potential role of specific shade trees in regulating residential energy consumption, rather than the average role of trees estimated in specification (1). By comparison, the coefficient for light shade tree removals suggests a more modest increase in electricity consumption, averaging 48 kWh per month, or an increase of 4.2-percent, a result of an unassuming magnitude that is also significant at the 1-percent level.

The coefficient estimate for no-shade tree removals provides a falsification test, which can address concerns about potential self-selection of tree removing residents. My analysis shows little evidence of observable differences between residents with tree removals compared to residences without tree removals.

It is theoretically possible, however, that an increase in electricity consumption after a tree removal is driven by a trend in energy usage that is related to the self-selection of a resident choosing to remove a tree. For example, if residences that remove private trees were experiencing faster growth in electricity demand, this could lead to a spurious positive "effect" from tree removals. On the other hand, if electricity increases occur through the causal channel of shade change, then one would expect a zero effect from tree removals without any change in shade. Trees that never shaded a home, or no-shade trees, serve as a test case of tree removals without any change in shade on the house. Results of specification (2) lack any evidence of a spurious effect for electricity usage, since the coefficient for no-shade tree removals is not different from zero at any relevant level of statistical significance, nor of a magnitude similar to cases of trees providing some shade for the house.

Given the identifying assumption that households removing no-shade trees are no different than households removing shade trees, the estimates presented in table 2.2 represent the effects of tree shade on energy use, rather than other confounding factors related to self-selection into treatment.. In addition, since Gainesville arborists reject permits for healthy, non-hazardous trees, homeowners are often removing diseased or damaged trees, rather than removing trees because they attach a low value to trees.

Turning to the natural gas results, the interacted shade-intensity variables in column (2) reveal a nuanced relationship between trees and natural gas consumption. The largest source of variability in natural case consumption, among homes with access to natural gas, would be due to its use in heating during winter months. As expected, the coefficient on heavy-shade tree removals has a negative sign, suggesting that a large increase in sun exposure following a tree removal leads to lower natural gas consumption for heating. While the sign of effects is predictable, as in the case of electricity results, the large magnitude of effects on natural gas is also somewhat surprising. In particular, residences reduce natural gas consumption by an average of about 5 therms per month after the removal of a heavy shade tree, which amounts to a 14-percent decrease in natural gas consumption. Further, this result that is significant at the 1-percent level.

In contrast, and contrary to expectations, light-tree shade removals have the exact opposite effect of heavy tree shade removals. Unlike heavy-shade results, the coefficient for light-shade removals suggests that residences increase natural gas consumption by 4.7 therms per month after the removal of a light shade tree. This result is significant only at the 5-percent level.

One possible explanation for a positive coefficient is straightforward: the effect of light-shade tree removals is unrelated to shade, but, instead, operates through another causal channel, which is the well-known service trees provide as windbreaks, or barriers that insulate houses from cold winter winds. Windbreak trees reduce wind-chill most effectively when set back from buildings by 50 to 100 feet, a distance that corresponds closely with the outer-ring buffer zones used in this study to define light-shade trees (see Figure 2.4).

Although windbreaks are most common in colder climates, natural gas results provide compelling evidence that trees set back from properties, or de-facto windbreaks, can provide energy-efficiency services during winter. Focusing now on the falsification test, the coefficient for non-shade tree removals is not statistically significant at any relevant level for natural gas, a finding that is consistent with the falsification test for electricity. Overall, natural gas results suggest that trees have a complex effect on winter energy consumption, with opposing effects depending on location: trees providing shade increase natural gas consumption, while trees serving as windbreaks decrease natural gas consumption.

Continuous Tree Canopy Loss Treatment

The analysis thus far builds a case that tree removals increase electricity consumption, due to change in shade, and also affect natural gas consumption, both due to changes in shade and windbreaks. The empirical strategy identifies effects based on both the timing of tree removals, from billing data, and the location of tree removals from satellite data used to classify shade density. Although proxy metrics for shade density refine inferences from models, nonparametric measures of tree location may help identify more specific information about how effects change with tree orientation and proximity to houses. In addition, while specification (2) provides estimates of average energy effects based on shade density, as defined by tree location, the effects may also differ in important ways depending on the size and number of trees removed. To test how tree size influences energy effects, and to further examine tree location affects, subsequent models include continuous variables of the size of trees removed, defined in terms of tree canopy area.

Two groups of four tree canopy variables are created. The first four variables record the area of tree cover removed in each quadrant within a 25 foot buffer zone around a house, with each variable representing a section to the north, east, south, or west. The second group of four variables records the area of tree cover removed in a ring-shaped buffer zone containing the area from 50 to 100 feet around the house, again divided into quadrants to the north, south, east and west. These two groups of continuous variables of canopy area are estimated separately for both electricity and natural gas, following the form

$$Y_{it} = w_{it}[North,\ East,\ West, South]\theta + \lambda_t + c_i + \varepsilon_{it} \qquad (3)$$

where the indicator variable for post-tree-removal is interacted with the continuous variables of tree canopy area within each of the directional quadrants for either the 0-to-25 foot buffer or the 50-to-100 foot buffer, depending on the model. Consistent with previous models, equation (3) is estimated using a two-way fixed effects estimator with standard errors clustered at the residence level.

The parameters of interest, contained in the vector θ, represent the marginal effect on energy consumption from increasing the canopy area of a tree removal, with separate estimates for each of the eight areas defined by orientation and proximity to a house. Based on previous shade-density results, and intuition about seasonal changes in sun angle, expected signs of coefficients for specific directional quadrants are straightforward. For electricity, I expect coefficients with a positive sign, representing higher consumption for increases in the canopy area of tree removals, especially for trees located along the west-side of houses. This is because trees on the west provide shade from the afternoon sun during the hottest hours of the day. Furthermore, for electricity, I expect tree cover within 0-to-25-foot buffer areas to have larger effects than tree cover for 50-to-100-foot buffer areas, since the summer sun is positioned high in the sky and casts short shadows.[15]

For natural gas, I expect coefficients with a negative sign, representing lower consumption for increases in the canopy area of tree removals, especially for trees along the east or south sides of houses. This is because sun from these directions is most relevant for winter heating: the eastern sun in the morning warms houses during in the coldest hours of day; and southern sun extends throughout most of the day during winter months. In addition, for natural gas-heating, tree cover in 0-to-25-foot and 50-to-100-foot buffer areas can be expected to have similar effects because winter sun is positioned low in the sky and casts long shadows. However, the potential energy effect of windbreak trees, if it exists, makes expectations about the sign and magnitude of natural gas effects ambiguous.

[15] Additional models were also used to investigate possible seasonal effects from tree removal. However, models including a treatment effects indicator interacted with an indicator variable equal to one for the months of April to November yielded estimates with economic inferences similar to the continuous treatment estimates presented in this section. Thus, they are not included here.

The first two columns of Table 2.3 report the electricity estimates of specification (3), including results for tree canopy area within 0-to-25-feet (column 1) and 50-to-100-feet (column 2) of houses. For electricity, coefficients have the expected signs for both buffer areas. Focusing first on west-side tree removals, coefficients are positive, suggesting that electricity increases with the removal of larger tree canopy to the west. These results are consistent with the expectation that loss of shade from the west, which occurs during the hottest hours of the afternoon, raises inside temperatures, and increases demand for electricity to cool houses. In addition, comparison between west coefficients reported in column (1) and column (2) suggest that electricity demand increases most from the loss of tree canopy in close proximity to houses, a finding consistent with effects caused by short shadows during summertime.

In particular, marginal effects suggest that removal of an additional 10-square-meters of tree canopy area on the west side of homes increases electricity usage by about 6 kWh per month for trees less than 25-feet from houses, and by 4.7 kWh per month for trees set back 50-to-100 feet from houses, results that are statistically significant at the 1-percent level and 10-percent level, respectively. In contrast, all remaining coefficients for continuous canopy area to the north, east, and south are found to be no different from zero at any conventional level of statistical significance.

Turning to natural gas results, reported in columns (3) and (4) of Table 2.3, all coefficients have the expected signs. Focusing first on south-side tree removals, coefficients are negative, suggesting that loss of shade from the south, which extends for most of the day during winter, increases inside temperatures, and reduces demand for natural gas-based heating. In addition, comparison between south coefficients reported in column (3) and column (4) suggests that natural gas demand decreases by similar amounts for the loss of tree canopy area anywhere within 100 feet of a house, a finding consistent with long shadows from the south during wintertime. In particular, marginal effects suggest that increasing the area of tree canopy removal by 10-square-meters along the south-side of homes decreases natural gas demand by 0.13 therms per month for trees less than 25-feet from houses, and by 0.07 therms for trees set back 50-to-100 feet from houses. Focusing on the east, the coefficient for canopy area is positive, as expected for morning sun that warms homes during the coolest parts of the day; however, these results are not different from zero at any conventional level of statistical significance. In contrast coefficients for canopy area to the north and west are negative, perhaps that appear to reflect windbreak effects that insulate houses from cold winter winds. However, again, results are not statistically significant.

In sum, canopy area variables provide evidence that electricity and natural gas effects increase with tree size. Results also confirm that previous tree shade results are robust to nonparametric definitions of tree shade, which verify that tree removal effects for electricity are, in fact, identified in the west, where summer sun is most intense, while tree removal effects for natural gas are identified in the south, where the sun is positioned for most of the winter. Furthermore, the magnitude of marginal effects, when aggregated to reflect the size of a large tree canopy of 350-square-meters are consistent with the magnitudes of electricity and natural gas effects estimated for heavy-shade variables with specification (2).[16]

---

[16] The most common tree species removed in Gainesville is the live oak, which has a span of 25-meters at full maturity, and a canopy area of approximately 450-square-meters.

CONCLUSIONS

As ecologically-conscious agencies search for effective policies to reduce energy consumption, there is, of course, the hope for a solution arising from nature itself. Intuitively, policymakers wishfully think of shade trees. However, to this point, these possibilities have remained in the realm of idealistic speculation. Taking a quasi-experimental approach to estimate accurate causal impacts, this study establishes that shade trees are, in fact, quite an effective means of reducing energy consumption. Not only are shade trees effective, but the magnitude of effects is surprisingly substantial. Therefore, a policy that both promotes tree shade canopies and also protects existing shade trees can be a powerful demand side management tool. Further, this study compiles significant evidence that can help homeowners locate trees strategically to lower both summer electricity bills and winter natural gas bills. Such recommendations could become an important element of consumer education to reduce future energy use.

APPENDIX FOR ESSAY 2: CHANGE CLASSIFICATION METHODS[17]

Change classification occurs in two stages: change detection and change classification. First, tree cover change between 2001 and 2011 was identified using a change detection model found under the image interpreter tool in Erdas Imagine, version 9.3 (ERDAS, 1999). The pixel-by-pixel algorithm analyzes differences between spectral values of imagery at each time period. The probability of change is determined by a maximum likelihood framework. Probability thresholds used to characterize change are determined subjectively by an iterative process of visual inspection of output in areas of known change. The model is revised to balance errors of omission and errors of commission. The output of the change detection model is a binary indicator map indicating whether or not change occurred for a given pixel location.

Vegetation is best characterized by red and near-infrared (NIR) wavelengths on the electromagnetic spectrum. Input data for the change detection include an NIR band from 2011 aerial imagery and a red band from 2001 imagery. An NIR band is not available for the 2001 imagery, which requires change detection across different spectral bands.

Second, classification of changed pixels into thematic groups begins after change detection has been finalized. Unsupervised classifications were conducted using the 'feature analyst' and 'ISODATA clump' tools in Erdas Imagine. The change classifications fall into three categories: trees, impervious surface, and open space. The 2011 four-band imagery was used to create baseline classifications that include additional

[17] I thank Binesh Maharjan, remote sensing specialist at the Global Ecosystem Center, who conducted this tree cover change analysis using Erdas IMAGINE (version 9.3) software.

categories of bare earth and water. The 2001 classifications were established by reclassifying detected changes from the 2011 baseline. Particular attention was provided to mapping changes in the tree cover class. The 'feature analyst' tool was used for initial classifications; the 'ISODATA clump' tool was used to revise classifications and create spatial cohesion by combining adjacent similar classified areas.

Other input data and classification methods were also attempted but provided problematic results. Classifications based on an NIR band from 1999 imagery proved problematic because of spectral differences between the 1999 and 2011 NIR sensor and differences in the season and time-of-day of imagery. A fusion between 1999 NIR and 2001 red imagery created similar problems. A more objective Classification and Regression Tree (CART) model also produced unreliable results due to spectral differences between 2001 and 2011 imagery. Although CART methods are typically preferred, they are most effective when using two sets of data from the same sensor. For classification of aerial imagery taken at different altitudes and angles using different sensors, subjective methods like the pixel-based change detection model and unsupervised (ISODATA) classification, and feature extraction methods, permit greater flexibility in defining the criterion for land use changes.

REFERENCES FOR ESSAY 2

Allcott, H. (2011), Social Norms and Energy Conservation. *Journal of Public Economics*, 95(9-10), 1982-1095.

Blackwell, M., S.M. Iacus, G. King, and G. Porro. (2009). cem: Coarsened exact matching in Stata. *Stata Journal*, *9*(4), 524-546.

Davis, L.W., A. Fuchs, and P. Gertler. (2014). Cash for Coolers: Evaluating a Large-scale Appliance Replacement Program in Mexico. *American Economic Journal: Economic Policy*, 6(4), 207-238.

Donovan, G.H., and D.T. Butry. (2009). The Value of Shade: Estimating the Effect of Urban Trees on Summertime Electricity Use. *Energy and Buildings,* 14, 662-668.

ERDAS, Inc. (1999). Erdas IMAGINE Field Guide. Fifth Edition, Atlanta, Georgia. Accessed May 4, 2015. Availible online at: http://www.gis.usu.edu/manuals/labbook/erdas/manuals/FieldGuide.pdf

Gainesville Regional Utilities (2014). Proposed Annual Operating and Capital Budget: Budget Year 2014. Gainesville, Florida. Accessed May 4, 2015. Availible online at: https://www.gru.com/Portals/0/PDF%20Forms/FY2014Budget/Budget%20Letter.pdf

Jacobsen, G. D., and M.J. Kotchen. (2013). Are Building Codes Effective at Saving Energy? Evidence from Residential Billing Data in Florida. *Review of Economics and Statistics*, 95(1), 34-49.

Kielbaso, J. (1990) Trends and Issues in City Forests. *Journal of Arboriculture,* 16(3), 69–76.

Pandit, R., and D. Laband. (2010). Energy Savings from Tree Shade. *Ecological Economics*. 69, 1324-1329.

Tou, J. T., and R.C. Gonzalez. (1974). Pattern Recognition Principles. Addison-Wesley Publishing Co., Inc. Reading, MA., 395 p.

Treiman, T., and J. Gartner. (2004). Community Forestry in Missouri, U.S.: Attitudes and Knowledge of Local Officials. *Journal of Arboriculture,* 30(4), 243-250.

United States Conference of Mayors (2008). Protecting and Developing the Urban Tree Canopy: A survey of 135 Cities. Report prepared by City Policy Associates, Washington, D.C.

Wulder, M., O. Niemann, and D. Goodenough. (2000). Local Maximum Filtering for the Extraction of Tree Locations and Basal Area from High Spatial Resolution Imagery, *Remote Sensing of Environment*, 73(1), 103-114

FIGURES FOR ESSAY 2

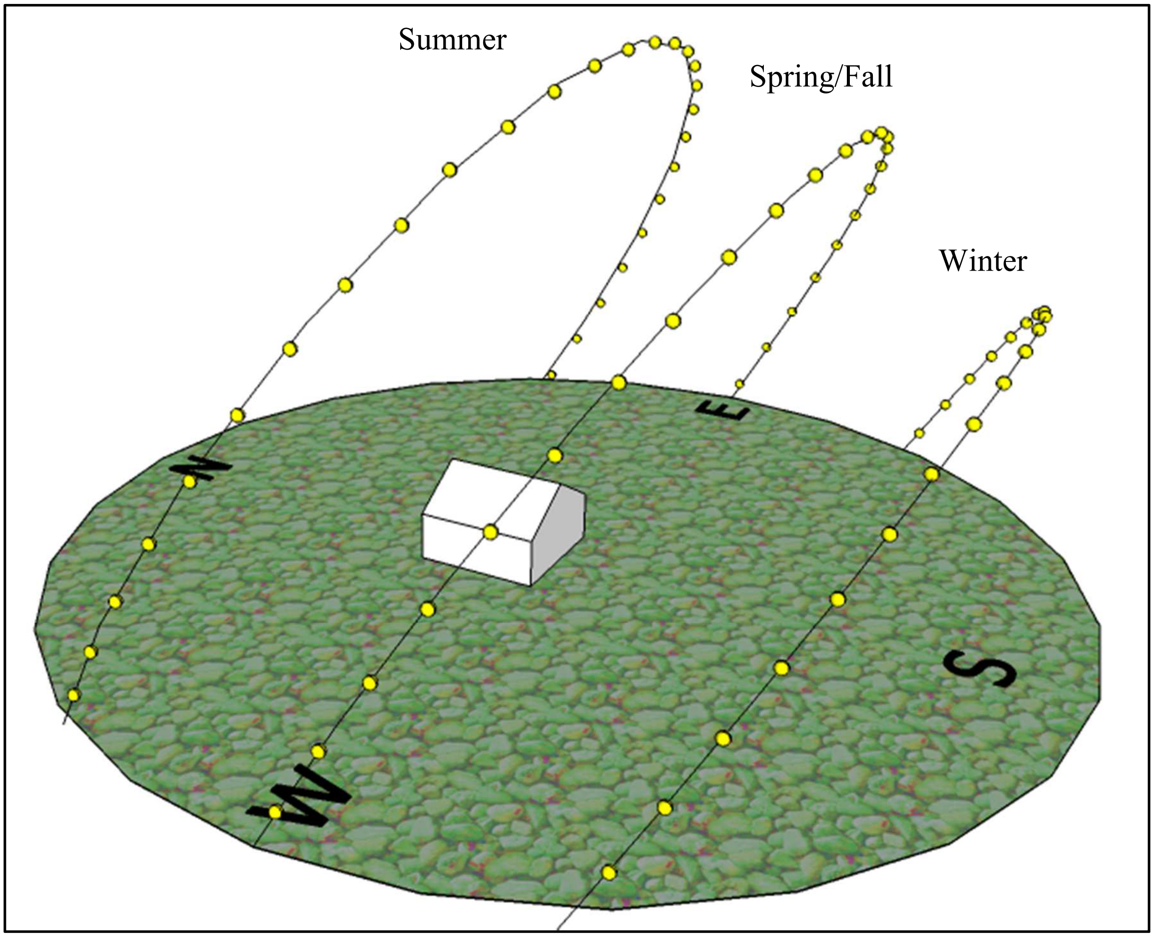

Figure 2.1 Sun path on the summer solstice (June 21), the spring and fall equinox (March 21 and September 21), and the winter solstice (December 21). Each yellow sphere marks the position of the sun in 30-minute increments between sunrise and sunset. From the house, athmuz angle measures sun direction (north, east, and south, west) around circular plane; altitude angle measured as the degree angle from house foundation to sun.

(a) Summer (b) Winter

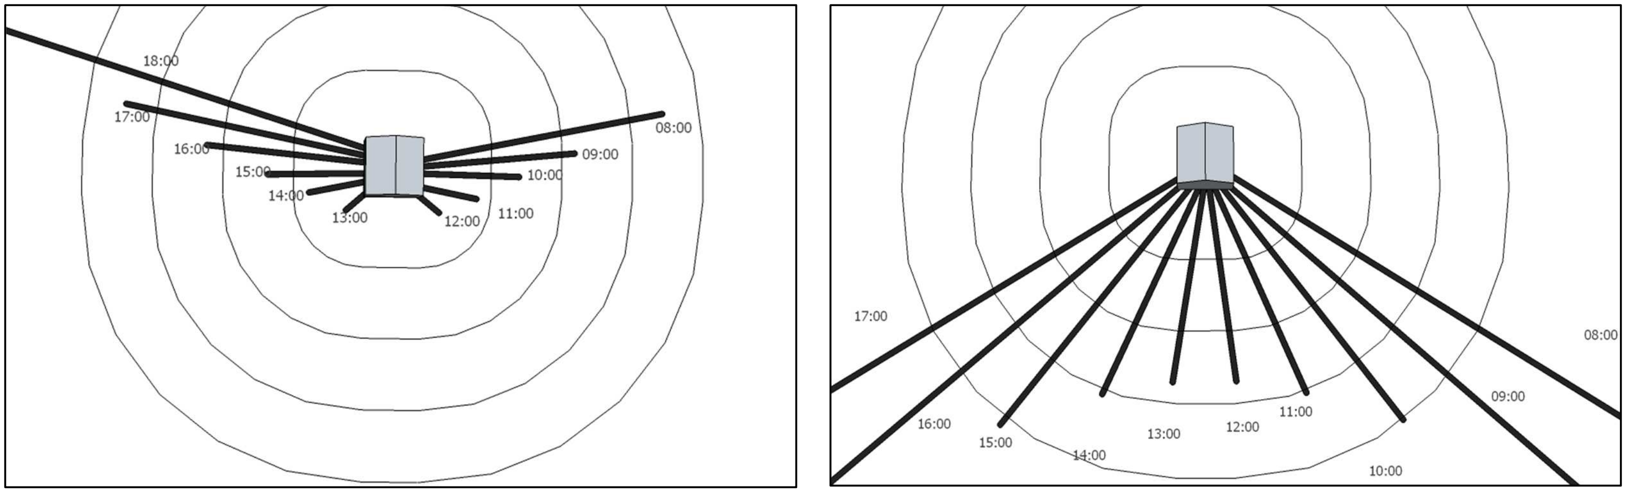

Figure 2.2 Hourly shadows cast by a 50 foot tree located at the maximum distance to shade the roof of a two-story house (15 foot roof) in Gainesville, Florida. Each black line represents the length and direction of a shadow on the hour during 8am-6pm. Concentric circles represent areas within 25ft, 50ft, 75ft, and 100ft buffers surrounding a square house. Panel (a) conveys the shadows on June 21, 2010 (summer solstice). Panel (b) conveys the shadows on December 21, 2010 (winter solstice).

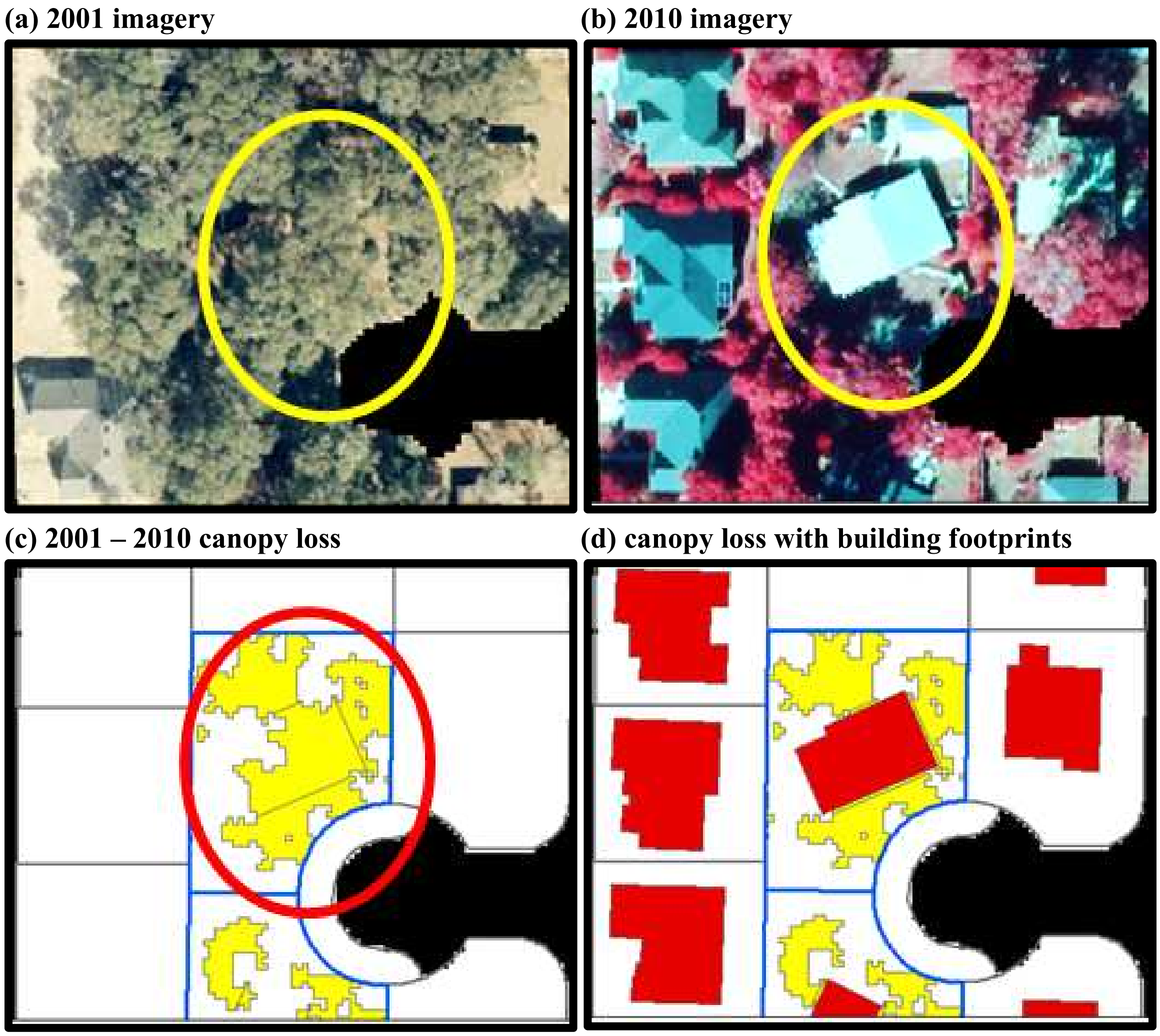

Figure 2.3 Visual illustrations of inputs and outputs used to create a canopy loss map, as well as auxiliary cadastral data used to define shade-intensity variables. Panels (a)-(d) cover identical spatial extent in a Gainesville neighborhood to illustrate the remote sensing data inputs. Panel (a) shows 2001 color imagery input (1-foot resolution) and panel (b) shows 2010 color infra-red imagery input (1-foot resolution). Panel (c) shows the tree canopy loss map (1-meter resolution) for two properties with tree removal permits (blue polygons). Yellow polygons represent areas of canopy loss during 2001-2010, blue polygons represent property boundaries for residents with tree removals, and gray polygons are neighboring property boundaries. Panel (d) adds cadastral data of building footprints used to determine the position and shade potential of canopy loss on each property.

(a)

(b)

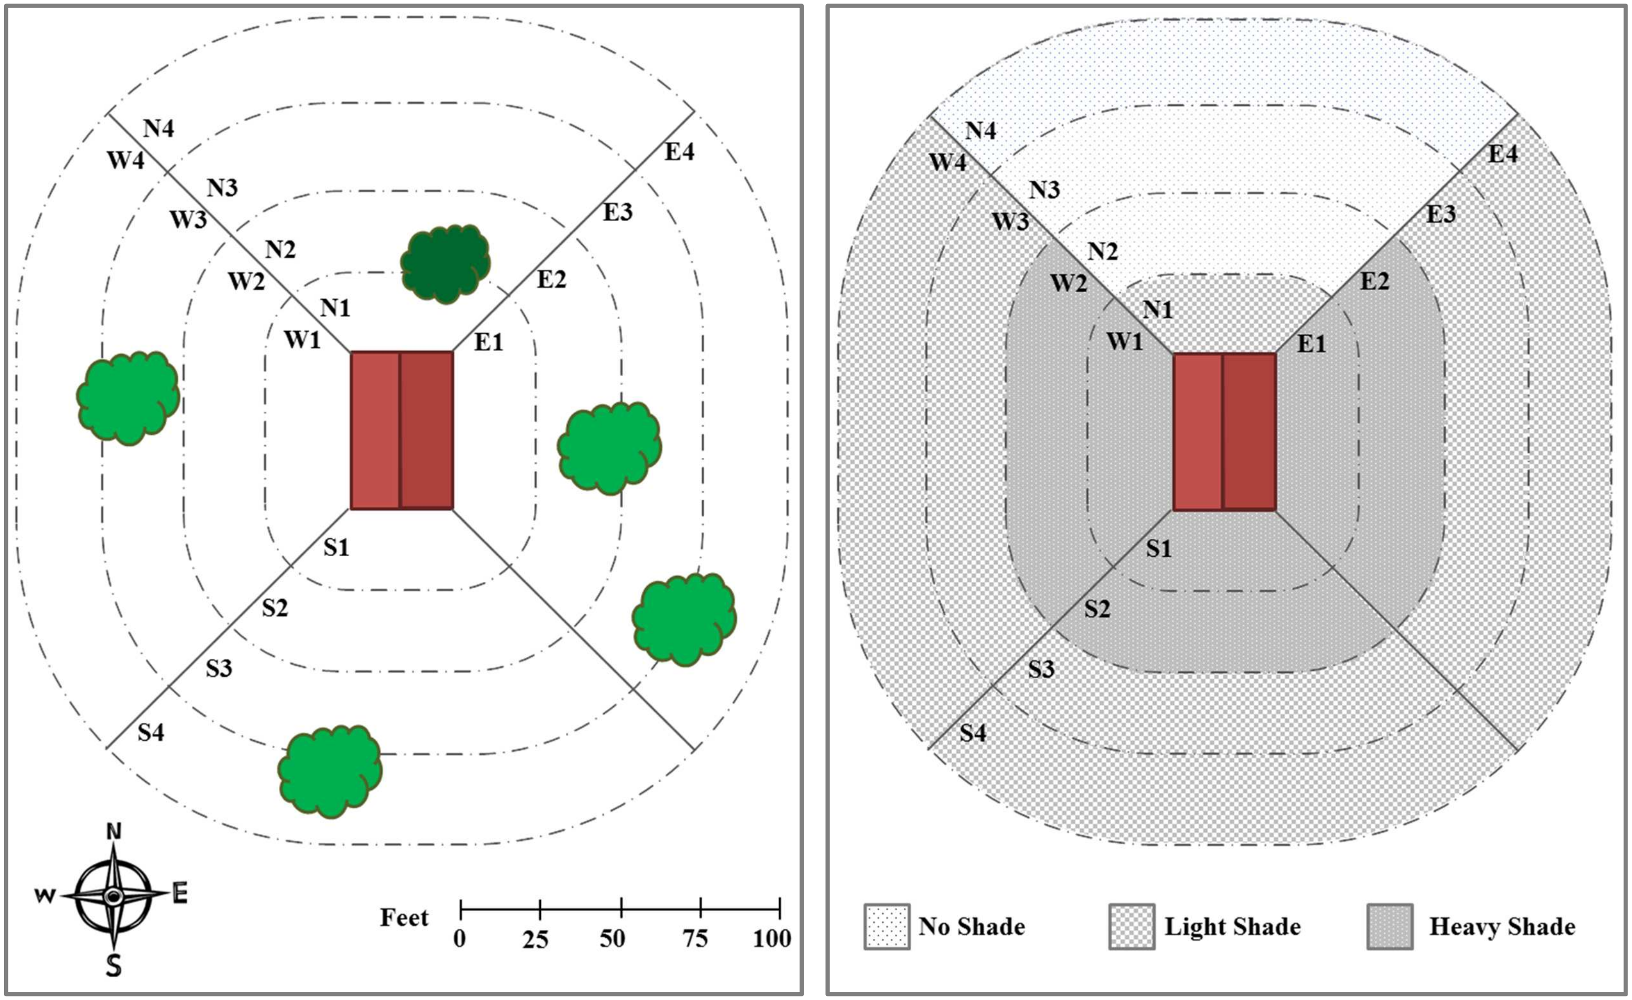

Figure 2.4 Sixteen concentric zones used to quantify tree position and shade-intensity. Panel (a) illustrates the four directional quadrants (north, east, south, west) and four distance buffers (25ft, 50ft, 75ft, 100ft) that define the sixteen zones of tree position for each house. Panel (b) shows the aggregated zones used to define three levels of shade intensity: no shade (white), light shade (checkered), and heavy shade (gray).

TABLES FOR ESSAY 2

Table 2.1: Effects of Tree Removal on Energy Consumption

Variables	kWh	therms
Tree Removal Effect	29.14 ***	-0.54
	(7.25)	(0.53)
Constant	306.14 ***	75.11 ***
	(33.79)	(1.30)
Number of utility bills	174,268	127,906
Number of residences	1,338	970
R^2	0.267	0.581

All regressions include fixed effects for residence and fixed effects for billing month. Dependent variable is monthly electricity consumption (kWh per month) for model (1) and monthly natural gas consumption (terms per month) for model (2). The independent variable is an indicator for utility bills occurring after a tree removal among residences that remove a tree. All models include standard errors clustered by residence. Asterisks denote statistical significance at levels of 1-percent (***), 5-percent (**), and 10-percent (*).

Table 2.2: Effects of Shade Intensity on Energy Consumption

Variables	kWh	therms
Removal * Heavy Shade	203.65 ***	-4.98 ***
	(25.19)	(1.87)
Removal * Light Shade	48.15 ***	4.69 **
	(9.59)	(1.87)
Removal * No Shade	-54.28	-1.84
	(55.45)	(1.64)
Constant	307.32	8.89
	(33.86)	(1.78)
Number of utility bills	174,268	126,936
Number of residences	1,338	970
R^2	0.267	0.581

All regressions include fixed effects for residence and fixed effects for billing month. The dependent variable is monthly electricity consumption (kWh per month) for model (1) and monthly natural gas consumption (terms per month) for model (2). Independent variables are tree removal indicators interacted with a measure for intensity of shade provided by removed trees. All models include standard errors clustered by residence. Asterisks denote statistical significance at levels of 1-percent (***), 5-percent (**), and 10-percent (*).

Table 2.3: Continuous Effects of Tree Canopy Loss on Energy Consumption

Variables	25 feet setback kWh	100 feet setback kWh	25 feet setback therms	100 feet setback therms
Removal * Area North (m²)	0.162	0.142	0.023	0.021
	(0.288)	(0.268)	(0.025)	(0.024)
Removal * Area East (m²)	-0.077	-0.301	-0.014	-0.013
	(0.249)	(0.198)	(0.018)	(0.013)
Removal * Area South (m²)	0.110	-0.044	-0.013 *	-0.007 *
	(0.167)	(0.132)	(0.007)	(0.004)
Removal * Area West (m²)	0.594 ***	0.467 **	0.005	0.001
	(0.195)	(0.227)	(0.007)	(0.009)
Constant	989.079 ***	989.158 ***	74.591 ***	75.109 ***
	(14.976)	(14.973)	(0.986)	(1.300)
Number of utility bills	175,606	175,606	127,906	127,906
Number of residences	1,338	1,338	970	970
R^2	0.267	0.267	0.547	0.581

All regressions include fixed effects for residence and fixed effects for billing month. The dependent variable is monthly electricity consumption (kWh per month) for models (1) and (2), and monthly natural gas consumption (terms per month) for models (3) and (4). Independent variables are tree removal indicators interacted with square-meters of canopy area removed within each directional quadrant. Models (1) and (3) use canopy area within 25 feet of a residence; models (2) and (4) use canopy area setback by 50 feet to 100 feet from a residence. All models include standard errors clustered by residence. Asterisks denote statistical significance at levels of 1-percent (***), 5-percent (**), and 10-percent (*)

ESSAY 3 : **PROTECTING FORESTS: IS IT THE POLICIES OR THE ACTORS THAT COUNT?**

Deforestation accounts for 12-20% of global carbon dioxide emissions to the atmosphere—the second largest source after fossil-fuel combustion. Protected Areas (parks) have long served as the primary policy tool to slow deforestation and remain a major component of forest conservation policies. One such model is the Amazon Region Protected Areas (ARPA) Program, a program that doubled park coverage in the Brazilian Amazon from 2002 to 2012 to create the largest system of parks in the world. During this time, Brazil achieved a 70% reduction in annual deforestation rates (Nepstad et al. 2014). Brazil's transformation has been spectacular; in 2005 Brazil was ranked as the 3rd largest emitter of carbon emissions, surpassed only by China and the United States whereas today it stands as the global leader in carbon reduction (FAO/World Bank 2005). However, how much of a role ARPA actually played in achieving these historic reductions is unclear in light of other new policies also addressing deforestation (Nepstad et al. 2014).

Given the scale of investment in parks and the potential role they play in climate policy, surprisingly little is known about how much carbon emissions are avoided by the implementation parks. As industrialized countries are poised to invest in Reduced Emissions from Deforestation and Degradation (REDD) as an integral part of climate policy, there is a timely need for comprehensive evaluations of park effectiveness. Of particular importance is understanding when parks may actually *increase* carbon emissions as a result of policy trends that promote resource extraction, and the concomitant increase in carbon emissions within parks.

Complexities of empirical park evaluation are rooted in the non-random siting of parks, leading to biased estimates from simple comparisons of protected and unprotected areas. Evaluations have also been plagued by a lack of consistent time series data on deforestation and the carbon emissions associated with forest loss. The best existing studies apply matching techniques to target control locations similar to existing parks to correct for non-random site selections. The overwhelming majority of these studies suggest that parks are broadly effective at reducing deforestation. However, few studies consider the decision process underlying site selection. As a result, studies that omit key variables associated with site selection may cause considerable bias in estimates of park effectiveness (Andam et al., 2008). Moreover, traditional use of forest cover as a proxy for ecosystem services, such as carbon sequestration, limit the policy implications of park evaluations (Ferraro et al., Busch et al. 2015).

Cross-sectional studies, which provide most of the available empirical evidence, are vulnerable to "hidden bias" due to omitted variables from matching estimates (Andam et al. 2008). The best empirical studies control for physical characteristics, but ignore the critical role that human inhabitants play in affecting deforestation and park selection. The potential bias from omitting social variables may invalidate central findings in the literature, such as the large deforestation-reducing effect of parks estimated for indigenous groups. Application of panel techniques are needed to control for all time-constant confounding variables, both observed and unobserved, in order to mitigate the problem of "hidden bias" found in existing cross-sectional studies.

Currently, only one study applies panel techniques to compare changes in deforestation rates inside and outside groups of parks. In this study, Blankespoor et al. (2014) claim that Latin American parks created after 2002 have no effect on deforestation rates. Despite the use of panel techniques, however, by restricting their sample to a pairwise comparison of buffer zones around park boundaries, the study ignores the effects of park interiors, and therefore is likely to underestimate the overall effectiveness of park units (Soares-Filho et al. 2010). Furthermore, similar to most literature, this study uses changes in forest cover as a proxy for avoided carbon emission effects.

Only one study estimates emissions-based treatment effects, which uses matching techniques and attributes large-amounts of avoided emissions to Brazilian parks established during 2002-2008 (Ferraro et al., 2015). However, this study uses data on carbon stocks for a single-year, a data limitation that requires unusually strong assumptions to infer the effect of parks on carbon flows.

This essay applies difference-in-difference methods to estimate the avoided carbon emissions from parks in the Amazon basin. The analysis focuses on parks established since 2001 and uses a spatially explicit panel dataset of annual deforestation data from 2000 to 2012, combined with estimates of initial carbon stocks in year 2000, and data on the timing and coverage of parks. This study illuminates the differing objectives of the state, federal, and indigenous agents that designate and administer parks, and empirically tests for consequent heterogeneous effects consistent with varying incentives faced by each agent. This study also highlights the importance of adjusting forest cover estimates for spatial variation in carbon stocks and identifies possible sources of omitted variable bias from the matching estimates is existing literature in the context of park evaluation.

My results challenge conventional claims about park effectiveness, suggesting significant heterogeneity in carbon outcomes driven by contrasting objectives of state and federal governments. Fixed effects estimates show federal parks achieve 18 tons of avoided carbon emissions per km^2 per year. However, the full effects of park expansion are offset by adverse selection of parks by state governments that led to 10 tons of addition carbon emissions per km^2 per year.   Results raise concerns that the ARPA program fails to account for differing incentives facing state and federal governments when balancing tradeoffs between environmental and economic objectives. This study provides a methodological framework for future park studies evaluating social and economic outcomes to better understand the effects of parks on local communities.

This study is structured as follows. The next section discusses related work and how government agents select parks. It foreshadows the regressions to be estimated and motivates the expected signs of various parameters. The subsequent section describes the data and variables used in the analysis, including descriptive statistics. Then the main section develops the econometric evidence regarding determinants of park effectiveness, including sources of heterogeneity and the temporal dynamics of effects. Results are compared to existing cross-sectional estimates. The final section concludes with a discussion of policy implications.

DEFINING AND MEASURING PARK EFFECTIVENESS

Park Location and Management

To correct for non-random site selections, the best existing studies apply matching techniques to identify control locations similar to designated parks based on observed physical characteristics. However, these studies omit key socio-demographic variables associated with site selection that may cause considerable bias in estimates of park effectiveness (Andam et al., 2008; Pfaff et al., 2014). For example, parks are regularly placed around indigenous populations, who reduce deforestation by preventing access to outside groups. Not surprisingly, studies overwhelmingly find that indigenous parks "cause" avoided deforestation because, when "control" areas lack any indigenous populations, the effect of park designation is confounded with the effect of indigenous groups on deforestation rates.

Few studies consider the decision process underlying site selection. Previous studies focus on two park management categories: either strict parks that allow no human activities, or multi-use parks that allow some economic activities. In the Amazon, studies present mixed results about park management: some claim strict parks are most effective while others claim multi-use parks are most effective (e.g. Soares-Filho et al., 2010; Pfaff et al., 2014). The most sophisticated studies indicate that park location matters most rather than management (Nolte et al., 2013b; Ferraro et al., 2013). The mixed results suggest that location and management considerations, examined in isolation, fail to explain the key factors affecting park effectiveness.

Government Actors and Objectives

This study examines determinants of park effectiveness using a new framework where government actors select park location and management as a joint-decision. Each government actor is considered to have exogenous objectives that strategically select parks to optimize along environmental, cultural, or economic outcomes.

Parks can be grouped in two ways. First, they can be grouped by the level of government responsible for park management, typically state government or federal government. Second, parks differ by strictness of restrictions on human activity, which is generally divided into two categories: strict protection or multiple-use protection. In total, combining government and strictness levels, parks can fall into one of four main groups: state-strict, state-multiple-use, federal-strict, and federal-multiple-use. The strictness of a park designation is only one factor that may influence the avoided deforestation from parks.

Government actors make decisions on park location, strictness of restrictions on human activity, and ultimately plan for park management, law enforcement, and infrastructure development. For parks in the Amazon basin, environmental protection often occurs at the expense of local communities whose livelihoods are based upon resource extraction, timber harvesting, cattle ranching, and agriculture. State and federal governments, the park decision-makers, face different political pressures and economic incentives that influence strategic park selection to favor either environmental objectives or local-economic growth objectives more strongly. Considering alternative possible incentives of government agents can help researchers reveal and differentiate an agent's objectives when targeting new parks. Furthermore, understanding agent interactions, for example, when one agent's behavior changes the incentives of another agent, will enable more comprehensive evaluations of program effects under alternative objectives, including indirect (unintended) effects on secondary actors.

Brazil: A Case of Conflicting Government Objectives

This section characterizes the relevant incentives of federal, state, and indigenous agents, with the specific aim of revealing agents' objectives when selecting new parks. To provide context for real-world incentives faced by actors, this section focuses on park-selecting agents in Brazil during 2000-2010.

Since 2000, the federal government has aligned policies to increasingly target environmental objectives in the Amazon. In 2002, strong international pressure to address agro-industrial deforestation prompted the federal government to initiate the Amazon Region Protected Areas Program (ARPA) as an effort to finance and coordinate a rapid expansion of the park system targeting areas with high development pressure. In 2004, a remote sensing monitoring system, the Detection of Deforestation in Real Time (DETER), increased law enforcement capacity leading to a string of unprecedented high-profile arrests for illegal deforestation. Federal measures to protect Amazon forests were politically popular in the coastal regions of Brazil, which makes up over 80% of the population (Fernside 2002, Nepstad 2014).

In 2008, the federal government strategically positioned Brazil as a climate change leader in order to attract investment and redirect attention from economic growth agendas. During the 2008 climate negotiation at the United Nations, the Brazilian government committed to a 2020 target of reducing deforestation by 80% from baseline rates during 1994-2005. In a sign of support, Norway pledged $1 billion performance-based funds to the Amazon Fund, a Brazilian government fund devoted to achieve deforestation targets. The federal government continues to position Brazil as a strategic recipient of international investments. For example, the 2012 Brazilian Forest Code (NFC) established a geospatial registry of properties, setting the groundwork to assess compliance of individual property owners in future REDD payment programs.

State governments, on the other hand, have stronger incentives to establish parks for economic development. Amazon state governments depend on agricultural expansion, with 90-percent of state revenue collected from value added taxes on municipal exports of beef and agricultural products (Hochstetler and Keck 2007). Further, corruption within state government agencies has been endemic. Some state officials are complicit or collude with illegal deforestation activities (McAllister 2008a, Hochstetleer and Keck 2007).

Regarding park designation, state agents may often respond to federal programs, which is an example of agent interactions with wide-reaching implications for program evaluation. Public opinion has been marked with antagonism against federal intervention limiting land use for environmental conservation at the expense of local populations. Prior to 2005, state government involvement in establishing parks may thus target economic development rather than carbon reduction, and may serve the purpose of establishing state jurisdiction of parklands to avoid seizure of land for federal parks. Furthermore, in anticipation of 2006 Forest Concessions Laws, state governments have used ARPA funds to build infrastructure and develop “sustainable” management plans as a basis for attracting investment, generating employment, and generating income from timber harvesting and Brazil nut collection (Impa, Bandiera et al., 2010).

On the other hand, in 2008, the Critical Counties program suspended access to agricultural credit for all farmers in 36 counties placed on a deforestation "blacklist" and further developed a list of specific properties "embargoed" from any type of government loan. Thus, following 2008, state incentives to create parks may have shifted towards environmental outcomes to remove municipalities from blacklists and to access revenue from the Amazon Fund.

Descriptive statistics of parks selected by federal and state governments demonstrate systematic differences in strictness of protection and baseline deforestation pressures. Table 3.1 shows total area of park designations by federal and state governments, aggregated by strictness of protection and 3-year presidential terms. Trends show a considerable increase in federal park designations in both strict and sustainable use areas during 2003-2006, followed by a large drop in federal park designations during 2007-2010 after low deforestation rates eased international pressure. State park designations dramatically increased favoring multi-use parks to strict parks by 3-to-1 during 2003-2006, followed by a drop in state designations after threats of federal park designations subsided during 2007-2010.

Table 3.2 shows the 2001 baseline deforestation rate in areas later targeted for protected status during 2002-2009. The proportion of parks across different designations is divided into three levels of human pressure using baseline deforestation rates: no pressure (< 0.05%), moderate pressure (0.05-0.1%), and high pressure (>0.1%). During 2002-2009, states located 84 percent of strict parks in areas with no human pressure, whereas only 42 percent of federal strict parks were placed in areas with no human pressure. Similarly, 71 percent of state multi-use parks were located in areas with no human pressure, compared to 43 percent of state multi-use parks.

As this brief review illustrates, prior matching studies that fail to differentiate the incentives for park formation likely yield very incomplete information for policy-making purposes. Federal governments apparently seek strict protection in areas of high deforestation pressure in response to international pressure to curtail high rates of deforestation. Conversely, state governments apparently seek multiple-use protection in areas of low deforestation pressure to promote economic development while preempting federal seizure of land for conservation purposes. The purpose of this study is to use available data to investigate the effects of this differentiation in incentives on the effectiveness of parks for carbon reduction. Results indeed demonstrate a dramatic difference that can be attributed to differing incentives of federal and state governments.

Measuring Park Effectiveness: Carbon Emissions Versus Deforestation

By using carbon data observed immediately prior to the period of deforestation data, treatment effect estimates for avoided carbon emissions should be unbiased. In fact, this is the first study to estimate emissions-based treatment effects using directly observed carbon stocks in the pre-treatment period. The two existing studies that directly estimate emissions-based treatment effects both use carbon data estimated during the post-treatment period after the implementation of conservation policies (Ferraro et al., 2015; Busch et al., 2015). These studies must interpolate unobserved, pre-treatment carbon densities by imposing assumptions that average carbon density remains constant over time, so that carbon measures of survived forests can be used to predict those of cleared forests. In the Amazon, this assumption is likely to fail because carbon-dense hardwoods, such as mahogany, are selectively targeted by loggers due to their high-value on international markets.

Emissions-based treatment effects estimated in this study provide new insights about the heterogeneous ecosystem-service benefits across park types. Unlike Ferraro et al. 2015, which estimates the cumulative effect for parks in Brazil, the analysis here tests separately for heterogeneity in avoided emissions across different park types. Heterogeneous effects across different park types may exist even when deforestation rate effects are similar. For example, some types of parks may be intentionally sited on land with unobserved unifying ecological characteristics that affect carbon densities as well as the siting decision. Or, due to differences in management, some types of parks may prevent deforestation in areas of the park with the highest-carbon density (primary forests), while other park types only have effects in areas of the park with the lowest-carbon density (secondary forests). Hence, strict parks may, at the same time, generate the least avoided deforestation and also generate the most avoided carbon-emissions, assuming strict parks have the highest carbon densities. Thus, impacts estimated in terms of avoided deforestation may be inadequate to compare policies designed to generate avoided-emissions. As a further example, if multi-use parks permit selective logging of high-value tree species with high carbon densities, there is potential for multi-use parks to accelerate carbon emissions. In such cases, studies that estimate average park effects will average across park effects with opposing signs.

EMPIRICAL SETTING AND DATA COLLECTION

Study Area

The study area is the humid tropics biome in the Amazon region of South American, where more than 40 percent of the world's deforestation emissions are located (Harris et al., 2012). Rainforests cover 5,500,000 km^2 spread across nine countries in the Amazon region. Brazil contains the lions-share of forests, with 63% of the rainforest area, followed by Peru with 17%, Columbia with 12%, Bolivia with 5%, and minor amounts in Venezuela, Ecuador, Guyana, Suriname, and French Guyana.

Figure 3.1 illustrates the study area including country boundaries (yellow lines), protected area coverage (light green areas), annual forest loss data from 2000-2012 (rainbow areas denoting year-of-forest-loss), and year 2000 forest cover density (greyscale ranging from 100% forest in black, to 0% forest in white). Figure 3.1 includes four panels that present data at varying spatial scales. Panel (a) shows the entire study area, with all 9 countries and shows that parks cover 36% of the Amazon. Despite the attention given to forest loss, more than 80% of the Amazon rainforest remains. During 2000-2012 deforestation affected only about 2% of the Amazon rainforests. Panel (b) zooms-in on an area with high-levels of deforestation since 2000, providing some visual evidence that deforestation in Brazil is higher than nearby areas within Peru and Bolivia. Panel (c) further zooms-in on the same region and illustrates that pre-2000 deforestation (white) around the inter-oceanic highway represents a frontier for progressive deforestation since 2000 (shown in colors). Panel (d) zooms-in on an adjacent area showing fish-bone patterns of deforestation on the Brazil-side during 2000-2005, and small hot-spots of deforestation in Bolivia after 2006.

Protected Areas Data

Protected area boundaries come from the World Database on Protected Areas (WDPA) created by the United Nations Environment Program and the International Union for the Conservation of Nature (IUCN). For each protected area (PA), the database includes information on the year of establishment, IUCN management category, the government agency with decision-making authority, and government agency with management responsibility. The WDPA sets a consistent standard for park information across all countries and represents a complete and accurate dataset when cross-validated with country-level government databases in the study area (UNEP-WCMC, 2007).

Forest Loss Data

The panel data techniques applied in this essay require longitudinal data comprising repeated observations on units of analysis over time. Until recently, consistent panel data on deforestation rates over did not exist. Although some studies have attempted to compare various forest cover maps across time to generate epochal forest loss products, such "panel" data contain considerable measurement error because static forest cover maps vary in quality across time (e.g. Sims 2010). Recently, consistent datasets of annual deforestation have been created by using satellite images recorded by the same sensor at regular intervals over time.

To characterize deforestation rates over time, this study uses the Hansen et al. (2012) dataset produced at a 30-meter spatial resolution with annual time-steps based on Landsat satellite imagery. It spans years 2000-2012.

Despite the seemingly coarse resolution of park-level data, econometric theory provides unambiguous guidance to aggregate data up to the level of variation in policies

(Angrist and Pinsche, 2007). Unlike analyses that select subsamples of 30-m pixels, which reduces efficiency to relieve computational problems, this analysis captures the full-variation in the 8 billion pixels that comprise the 12-year 30-m panel dataset without diminishing the statistical power to identify effects in a difference-in-difference model. In fact, after weighting observations by land area, PA-level analyses produce unbiased estimates identical to a panel model of each individual pixel (Angrist and Pinsche, 2007). As an additional benefit, coarser spatial resolution data mitigates spatial correlation and diminishes measurement error effects of possible spatial misalignments between datasets.

Carbon Biomass Data

Annual carbon emissions from each park are derived by combining annual forest loss products with existing data on forest carbon stocks prior to park designation. As an input, spatially explicit carbon density data provided by Saatchi et al. (2011) are used to estimate forest carbon stocks in 2000 at 1-km spatial resolution. These data contain estimates for above-ground biomass, below-ground biomass and carbon density.

Carbon loss variables are calculated in three steps.[18] First, to resolve differences in the spatial-resolution, data on annual forest loss are aggregated from 30-m to 1-km resolution by calculating the 30-m forested pixels and forest loss pixels as a percentage of each 1-km grid. Second, for each 1-km cell, forest carbon stocks in 2000 and annual forest carbon loss were calculated by multiplying carbon density with the percent forest cover and percent annual forest loss, respectively. Third, the 1-km grids over the study

[18] I thank Xiaopeng Song, who performed the calculations to generate annual carbon loss variables.

area were aggregated to derive PA-level estimates of forest carbon loss for each year between 2000 and 2012.

Land Characteristics Data

Ecoregions are classified by the World Wide Fund for Nature (Olson et al., 2001). Agricultural suitability is derived from the Global Agro-Ecological Zones dataset provided by the International Institute for Applied Systems Analysis. For each 1-km^2 grid cell covering forest lands, this index ranks agricultural potential on a scale of 1 (very high potential) to 7 (very marginal potential) with an additional category of 8 for land not suitable for agriculture (Fischer et al., 2002).

Physical landscape characteristics include slope, elevation, and accessibility. Slope is measured in degrees from horizontal for 1-km^2 grid cells (Fischer et al., 2002). Elevation is measured in meters and is derived at a 90-m spatial resolution by the Shuttle Radar Topography Mission and published online by the Consortium for Spatial Information (Jarvis et al., 2008). Accessibility measures travel time to major cities along transportation networks (e.g. roads, navigable rivers) and off-transport networks (walking speed given environmental factors). These measures were created at a100-m scale by the World Bank and published online by the European Union Joint Research Centre (Uchida and Nelson, 2009). Climate variables for monthly mean precipitation in millimeters and monthly mean temperature in degrees Celsius are provided for 1-km^2 grid cells by WorldClim Global Climate Data (Hijmans et al., 2005).

Strictness of Protection Variables

The variable for "strictness" of protection in this study is defined using the six IUCN management categories adopted for national protected areas. These standardized classifications identify management types based on the restrictions on human activity, loosely ranked from most strict (category I) to least strict (category VI). In cases where park sites are nationally recognized as protected but have not been assigned an IUCN category, data are not. Some indigenous areas that are considered international parks that are listed as "Non-Applicable" under management categories, as their primary purpose is not intended for conservation or sustainable development.

This study follows conventions in the literature (e.g., Nelson and Chomitz, 2011) that define four classes of park strictness: strict protection (IUCN categories I-IV), multi-use protection (IUCN categories V-VI), and not-reported (no IUCN category).[19] The fourth class of indigenous protection cannot be defined using IUCN categories; instead indigenous parks are defined from governance variables detailed in the following section.

Government Ownership Variables

The government agent variable classifies park ownership across four government types: federal, state, indigenous, and non-government entities. These agent-types are derived from WDPA information on governance that describes the decision-making structure of the PA, or the government agent(s) responsible for selecting, declaring, and administering park lands. First, federal ownership applies to governance identifying "federal or national agencies" in charge of parks. Second, state ownership applies where

[19] Strict parks include state and national biological stations, biological reserves, and national and state natural areas. Multi-use parks include state and national forests, extractive reserves, and sustainable development reserves.

"sub-national" agencies are in charge. Third, indigenous ownership includes lands identified as "declared and run by indigenous people." Fourth, non-government ownership includes small areas of lands declared and run by individual landowners, non-profit organizations, for profit organizations, or with unreported ownership.

Additional rules are applied to classify federal and state ownership for the 3.5% of parks where the governance type is ambiguous.[20] In cases where the governance type is reported as "government delegated management" or "collaborative or joint management" parks are assigned to either federal or state governance depending upon which level of government that officially declared the PA. For example, federal ownership would be applied when a federal agency designates a PA, even in cases where management tasks are later delegated to local state agencies. This is because the designating government selects the park siting and strictness of protection, the two permanent characteristics affecting selection bias, and also sets the basis for cooperation or antagonism with human inhabitants. Recent research suggests that management activities have little impact, if any, on deforestation outcomes—a finding that validates the importance of initial selection decisions (Nolte et al., 2013a).

Unlike IUCN categories used to define strictness of protection, the WDPA variable for governance clearly defines indigenous parks. Hence, for this analysis, indigenous parks are always defined from the WDPA governance variable. Thus, the sample of indigenous parks is identically labeled as 'Indigenous' in both the government ownership variable and the strictness of protection variable. In contrast, parks owned by

[20] Ambiguous governance categories make up 3.5% of individual parks but represent only 0.35% of total land area in parks.

state and federal governments can take on strictness of protection values of either multi-use parks or strictly protected parks.

Indigenous parks represent a special category of protected area found predominately in Latin America. Furthermore, within the Amazon, Brazil is the only country to designate new indigenous parks during the past 20 years. Although formally established by Brazil's federal government, indigenous parks are governed through different regulator frameworks than other parks and grant considerable autonomy and land-rights to indigenous inhabitants. Indigenous lands are subject to restrictions on development and resource use that are devised through a joint planning process involving federal, state, and local governments and indigenous communities.

EMPIRICAL METHODS

Selection of the Sample

In order to study the impact of declaring a piece of land a PA, one needs to have the deforestation data of that land both before it was declared a park and after it was declared a PA. Such data is available for parks that were created between 2001 and 2009 in four countries in the Amazon: Brazil, Peru, Bolivia and Columbia. Of the park land in this data, 83% is located in Brazil, 14% is located in Peru with the remaining 3% located in Bolivia and Columbia.

However, it is not enough to look at deforestation rates before and after land becomes a PA. There are other factors, such as changes in climate, accessibility, and economic demand that also affect deforestation rates over time and occur regardless of whether the land is protected or not. In order to control for changes in deforestation rates that are independent of the land becoming protected, the deforestation rates of lands that were already parks at the start of the time period of this study can be compared with the deforestation rates of lands that transitioned into parks during the course of the study's time period. Accordingly, the lands in the comparison group in this study are lands designated as parks before the time period of this study began, i.e., before 2001. These are compared to the lands transitioning to parks between 2001 and 2009.

One might also suggest doing a comparison as well with lands that were never designated protected areas throughout this study's time period to see what changes in deforestation rates they experienced. However, problems of consistency exist across those lands that make them unsuitable for this study. Those lands may not have been selected for parks because they were never good candidates for park designation. For example, some lands are large, sprawling, urban areas. Other the lands might be good candidates, but for various reasons, such as a variety of political factors, they have not been chosen (yet) for park designation. This may apply to remote jungle areas. Without detailed knowledge about the selection process and the data to replicate it, identifying which among the undesignated lands are potentially good park candidates for comparison purposes has the risk of being too speculative for this study.

In summary, the panel data models used in this essay are estimated with two groups of lands: one group consists of lands that are transitioning to parks between 2001 and 2009 and the other group consists of lands already transitioned to parks before 2001. The causal effect of parks is estimated for those lands that were undesignated as of 2001, and transitioned to parks at some point during the period of 2001 and 2009. To control for other factors affecting deforestation over the 2001-2009 time period, the impact of park designation on transitioning lands is examined using lands that became parks before 2001, which remain protected throughout the study time period, for comparison purposes.[21] Figure 3.2 presents a map showing the spatial configurations of pre-2001 and post-2001 parks.

Collectively, the parks in this sample cover a large fraction of Amazon rainforest in each country, including 45% of rainforest in Brazil, 22% of rainforest in Peru, 19% of rainforest in Bolivia, and 10% of rainforest in Colombia. More broadly, the sample of parks established during 1985-2009 covered 34% of the Amazon rainforest in 2009, while parks established before 1985 covered only 3% of rainforest, and unprotected lands covered the remaining 64% of rainforest.

Unit of Analysis

Ideally, the unit of analysis should correspond to the decision making process. Following this principal, the unit of analysis in this study is the individual park, which constitutes the spatial scale at which a piece of land becomes protected at a point in time.

This is, in fact, the first study to apply a park-level unit of analysis. By comparison, previous studies typically use pixels, aggregated cells, or administrative divisions as the unit of analysis. This limitation stems from a reliance on cross-sectional analysis, as well-defined "unprotected areas" analogous to parks simply do not exist (Blackman, 2013). In this study, panel data are used to circumvent this problem by exploiting variation in the timing of park designation.

[21] Specifically, the comparison group includes parks established during 1985-2000, which includes 85% of all Amazon park lands created prior to the study time period. The remaining 15% of parks established before 1985 are excluded from the sample, as the selection process for older parks is likely to differ from the selection process of more recent parks considered in this study. Among the Amazon parks created between 1985 and 2000, 79% of park land is located in Brazil, 8% is located in Columbia, 6% is located in Peru, 4% is located in Bolivia, and the remaining 3% is located in Guyana and Venezuela.

There were 610 parks designated between 1985 and 2009 in the Amazon tropical forest biome that contained at least 1 km^2 of tropical forest. All datasets were first aggregated to 5,738,163 1-km^2 grid cells that collectively covered the Amazon humid tropics biome. Next, the 1,963,594 1-km^2 grid cells contained within the park sample were aggregated to the PA-level. The final sample includes 7,320 PA-year spatial-temporal units of observation for the 12 year panel of 610 parks.[22]

The unit of observation for this study is a park-year, and the dependent variable is either annual deforestation rate (percent forest loss per year), or carbon loss per year, depending on the model. To construct the variables from binary forest loss data at a 30m or 250m resolution, data are aggregated to a 1km^2 unit of location with a continuous measure of deforestation in each year from 2000 to 2012. For the sample of parks established between 1985 and 2009, indicator variables for each type of park designation are equal to one if a protection status is present in a given location and year and zero otherwise. Due to uncertainty in the date of forest loss within a given year, I exclude the dependent variable in the year of designation.

[22] Since the empirical objective is to estimate average "treatment" effect of park designation on carbon emissions (and forest loss) in absolute terms, estimates place more weight on larger parks to allow inference of area-based treatment effects in terms of percentage based metrics.

Empirical Model

In the case of a several types of park designations, the quasi-experiment entails estimating the following two-way fixed effects model:

$$Y_{\{it\}} = \theta_{\{i\}} + \varphi\tau_{\{t\}} + \gamma_{\{j\}}\, AT_{\{ijt\}} + \varepsilon_{\{ijt\}} \qquad (3.1)$$

where $Y_{\{it\}}$ is either the deforestation rate or carbon loss at park *i* in year *t*; $\theta_{\{i\}}$ are park-specific indicator variables that capture mean differences in deforestation across parks due to differences in factors that remain constant over time, such as landscape characteristics or stable human settlements; $\tau_{\{t\}}$ are year-specific indicator variables capturing broad trends that affect deforestation over time; and $AT_{\{ijt\}}$ is an indicator variable equal to 1 for years after a park of type j is designated and equal to 0 otherwise where type j indicates either government ownership or strictness of protection, and ε_{ijt} is an error term that represents unmeasured factors.

The coefficient $\gamma_{\{j\}}$ isolates a separate effect of designation for each park type (i.e. average treatment effect); while coefficient φ separates counterfactual deforestation trends.The model therefore estimates the relative effectiveness designation status across park types. Because all parks are designated by year 2009, effects are identified for parks established after 2001 that experience a change in protection status during the time interval of available deforestation data. Additional models include state-by-year effects that absorb state-specific shocks, as well as park-specific linear and quadratic time trends that allow park deforestation rates to trend nonlinearly. Alternate specifications include spatial leads and lags, which test for temporal trends in effects.

EMPIRICAL ANALYSIS

Baseline Difference-in-Difference Estimates

Initial estimates of equation (3.1) are contained in table 3.1. Each column resents a regression of the rate of deforestation on park and time indicator variables, state-by-year effects (columns 2, 3, 5 and 6), linear and quadratic park time trends (columns 3 and 6), and indicator variables for the four types of park designation, which are equal to one if a protection status is present in a given park and year and zero otherwise. The first three columns contain simple estimates for the effect of park designation on deforestation rates. The coefficient of -0.039 in column 1 indicates that after removing mean park deforestation levels and common year effects, the deforestation rate fell by approximately 0.04 percentage points more in parks with a designated protection status than in candidate areas without any park designation. However, this effect is estimated imprecisely.

The second column adds 180 state-by-year effects to the model, decreasing the point estimate considerably to -0.027 percentage points and reducing the standard error slightly after controlling for yearly shocks in each of the 15 states. The third column adds 610 park-specific linear time trends and 610 quadratic time trends to the model, decreasing the point estimate slightly for the effect of park designation.

Heterogeneity across Government Actors

The last three columns of table 3.1 estimate separate effects of park designation across each of four levels government: federal, state, indigenous, and non-government entities. The point estimate for a federal park designation is strongly negative at -0.07 percentage points with park and year effects, and is only minimally affected by the inclusion of state-by-year fixed effects. Adding linear and quadratic park time trends to the model increases the magnitude of the federal park point estimate considerably to -0.12 percentage points, and remains significant despite the inclusion of more than 2,000 covariates. Thus, federal parks appear to generate avoided deforestation, a result that is robust to model specification.

Unlike federal park designations, point estimates of indigenous and state parks are sensitive to the addition of covariates. Not surprisingly, indigenous parks appear initially to contribute to lower deforestation rates, a result that echoes findings in previous cross-sectional studies. However, the point estimate for indigenous parks approaches zero and becomes insignificant after adding state-by-year effects, and remains insignificant after the inclusion of park time trends. It appears that indigenous parks were established in areas where deforestation rates were already decreasing and are thus not significantly additive.

Upon closer inspection, sensitivity of estimates for indigenous areas to state-specific shocks is not surprising, as deforestation drivers vary widely across space. Since 2001, most indigenous areas were located in three remote states of the northern Amazon: Amazonas, Para, and Roraima. In contrast, during 1985-2000, a large number of indigenous areas were also located in two highly populated states inside the arc of deforestation: Mato Grosso and Rodonia. Since indigenous parks are designated in different regions over different time periods, the omission of state-by-year indicator variables (column 4) confounds the effect of designating indigenous parks with any regional variation in deforestation trends. The indigenous coefficient diminishes in magnitude and significance (column 5) after controlling for state-by-year effects.

State parks, on the other hand, have a quite surprising evolution of point estimates. While state park point estimates are negative and insignificant with few controls, the point estimate reverses signs to become positive and significant once park time trends are included. It appears that state parks prop up elevated deforestation rates in areas where deforestation rates were already starting to decrease. These results support the hypothesis that state agents create parks to advance economic objectives that, on average, accelerate deforestation rates. The addition of linear and quadratic time trends control for park-specific factors affecting deforestation trends and reflect state actor's selection of park sites with declining deforestation pressure. Estimates suggest that the state policy intervention of establishing a park slowed this decline in deforestation pressure, creating a positive impact on deforestation rates. State actors may increase deforestation directly, through issuance of legal logging concessions, or indirectly by building infrastructure or supporting local livelihoods that may attract illegal deforestation or land speculation by poor migrant groups.

The remaining group of non-government owned parks, which are designated by non-profit organizations, private land owners, or undefined ownership yield point estimates that are insignificant across all models. The reasons for the overall low impact of non-government owned parks may vary across constituent groups. In Brazil, large private landowners often declare forested areas as parks in order to avoid paying property taxes; however, landowners may set aside only the marginal forest lands that were never worth clearing. Meanwhile, non-profit and for profit groups may have limited resources to enforce park regulations. In addition, a large portion of the undefined category is in Peruvian parks designated as "buffer zones" or rings of land that surround existing parks

These buffers may serve to limit access to interior parks rather than restricting deforestation within the buffers themselves.

Tests of joint significance confirm the existence of treatment effects and the relevance of covariates. Specifically, using the model presented in column 6, an F-test of the hypothesis that all four government treatment coefficients are jointly equal to zero is rejected at a 5-percent level, with a p-value of 0.04. In addition, an F-test of the hypothesis that state-by-year coefficients and park time trends are jointly zero is strongly rejected at the 1-percent level. Further F-tests of joint significance separately applied to state-by-year coefficients, linear park time trends, and both linear and quadratic park time trends are all rejected at the 1-percent level. Hence I employ linear and quadratic park trends and state-by-year covariates in all remaining specifications.

Controlling for Strictness of Protection

The results in table 3.1 suggest that parks declared by federal governments are the only one of the four types of government parks to reduce deforestation. It is possible, however, that the difference is explained not by the role of the federal government itself, but other facets of the park designation that are correlated with federal parks. Measures of the "strictness" of protection, which are commonly used to test for heterogeneous effects of park designation, deserve particular attention.

To examine this issue, I begin with the specification from column 6 of table 3.1 (containing quadratic time trends and state-by-year effects) as column 1 of table 3.2. Table 3.2 column 2 presents estimates from a similar model using explanatory variables measuring strictness of protection, and column 3 presents model estimates including strictness of protection variables and government ownership variables. The second column of table 3.2 shows that of three common groupings of strictness—strict, multi-use, and indigenous—none have an effect on deforestation that is different from zero at relevant levels of statistical significance. The third column of table 3.2 shows that point estimates remain insignificant with the inclusion of controls for type of government.[23] Consistent with ambiguity in the existing literature, these results provide little evidence that strictness of protection is a key determinant of avoided emissions from parks (e.g. Ferraro et al., 2013).

Furthermore, point estimates of effects across government type are robust to models that control for the strictness of park designation. For federal parks, while adding controls for strictness to the model reduces the point estimate somewhat, from -0.12 percentage points (column 1) down to -0.09 points (column 3), the effect remains large and statistically significant. For state parks, including controls for strictness increases the point estimate considerably, from 0.046 percentage points to 0.074 points, with a slight increase in significance. Thus, the type of government ownership appears to be much more important than the strictness status of protection.

Comparing Panel and Cross Sectional Estimates

The fixed effects estimates in table 3.2 differ from cross-sectional matching estimators in existing literature. In particular, studies based on matching estimators consistently find indigenous parks to reduce deforestation, whereas fixed effects results show little evidence of effects from the establishment of indigenous parks (e.g. Nelson and Chomitz 2011). One possible reason could be differences in methodology: if

[23] Indigenous parks are defined in the same way for both the government-ownership variable and the strictness-of-protection variable. Hence, to avoid mulicolinearity, only one "indigenous" indicator is included in specifications that include treatment variables for both government ownership and strictness of protection.

matching omits important determinates of deforestation on indigenous lands then matching estimators will be biased, while any bias from time-invariant unmeasured determinants is purged from fixed effects estimators. It is also possible, however, that disagreement with conventional wisdom stems from the use of different data on deforestation, rather than differences in methods.

To examine the role of methodology, I repeat the specifications in table 3.2 using cross-sectional estimators on the data—namely effects by government type, categories of strictness, and simultaneous estimation of government and strictness measures. As presented in table 3.3, each model is estimated with a naïve specification, a specification that adds controls, and with a sample restricted to a matched control group. First, naïve models (columns 1, 4, and 7) include year effects and park indicator variables equal to one if a protection status is present in a given park and year, and equal to zero for always-unprotected areas. It is worth noting that parks appear in the data only for years following formal park designation. Using always-unprotected areas as an alternative control group, as in previous matching studies, allows the most direct comparison of the cross-sectional approaches used in other studies to the panel approach introduced in this study.

Next, in columns 2, 5, and 7, I add state-by-year effects and detailed geographic variables related to deforestation, including: initial tree cover density, elevation, slope, accessibility (travel time to major cities), an agricultural suitability index (ranging [0-10]), average monthly precipitation, and average annual temperature. To employ a simple matching model, Columns 3, 6, and 9 use propensity score matching at the 1-km^2 pixel level using variables for state and travel time to major cities.

Estimates of effects by government ownership are presented in columns 1-3 of table 3.3, and estimates of effects by strictness of protection are presented in columns 4-6 of table 3.3. The first and fourth columns of table 3.3 show that naïve model estimates are significant and negative for all government types and all strictness categories, with estimates in the range of -0.18 and -0.31 percentage points. These estimates, that compare parks to unprotected lands, however, contain bias due to the non-random selection of park locations. The second and fifth columns of table 3.3 show that the addition of state-by-year effects and geographic covariates has minimal effects on federal, state, strict, and multi-use designations. In contrast, point estimates for indigenous designations jump considerably from -0.31 to -0.55 percentage points in both models, suggesting indigenous park cross-sectional estimates are sensitive to model specification.

Figure 3.3 graphically displays naïve (column 1) and matching (column 3) point estimates alongside fixed-effects estimates (table 4, column 1) for federal, state, and indigenous parks. For federal parks, matching model estimates are almost identical to fixed-effects estimates, suggesting that matching methods eliminate bias present in naïve cross-sectional specifications. For state parks, matching model estimates are insignificantly negative, suggesting that matching reduces some cross-sectional bias but does not lead to the significant positive estimate of the fixed-effect model. For indigenous parks, matching estimates maintain a strong negative point estimate of -0.29, almost identical to the naïve estimate with minimal bias corrections towards the insignificant near-zero estimates of the fixed effects model.

Figure 3.4 displays comparable estimates of naïve (column 4), matching (column 6), and fixed-effects estimates (table 4, column 2) for strict, multi-use, and indigenous parks. Matching model estimates for strict and multi-use parks fall between larger naïve estimates and small fixed-effects estimates: strict matched estimates correct for approximately half of the naïve bias, and multi-use matched estimates correct for most of the naïve bias. Again, matching model estimates for indigenous parks seem to be equally biased as naïve models. Apparently, there are important unmeasured determinants of deforestation in indigenous areas, which are effectively purged by the fixed effects models used for the primary analysis.

In general, the bias of matching estimators varies greatly across park type, suggesting that fixed-effects estimators are better suited for comparing the relative effectiveness of park types.

Inferring Avoided Carbon Emissions

Returning to the preferred difference-in-difference panel technique presented in equation (3.1), the estimates above measure the effectiveness of parks in terms of reduced rates of deforestation and generate an area-based forest loss measure that is standard in existing literature. For these estimates to serve as valid metrics for program evaluation, however, imposes the faulty assumption that every hectare of forest in the Amazon provides equal benefits in terms of programs goals. Due to the vague term “forest,” most often defined as land with 30 percent tree cover or greater, an area-based metric of “deforestation” treats a 100-percent tree covered hectare the same as a 30-percent tree covered hectare. Furthermore, area-based metrics do not distinguish across forest quality: a hectare of 150-year old-growth forest is equal to a hectare of 5-year-old trees in a regenerated forest, despite vast differences in biomass, carbon storage, and other ecosystem services.

To address limitations of deforestation rates, I use an alternative metric of carbon loss per square kilometer. Carbon loss is equal to the share of forest loss multiplied by the carbon density of the original forest. The carbon density map is created at the 1 km^2 unit by combining maps of percentage tree cover with above-ground carbon density maps (Hansen et al., 2012; Sacchi 2001). To examine the effect of parks on carbon loss, I begin with the specification of deforestation rates from column 3 of table 4 (containing simultaneous estimates of government and strictness treatment effects) and introduce in table 3.4 comparable estimates using carbon loss as the dependent variable.

The second column of table 3.5 shows the effect of park designations on annual levels of carbon loss, measured in units of megagrams (Mg) of carbon per square kilometer per year (Mg C/km^2/year), where 1 Mg is equal to 1 metric ton (1.1 US tons). The federal park coefficient of -18.46 in column 2 indicates that after removing mean park carbon loss levels, common state-by-year effects, and park-level carbon loss time trends, carbon loss slowed by approximately 18.5 Mg/km^2 per year more in parks with a federal park designation compared with non-designated candidate parks. The state park point estimate is significantly positive, and suggests that carbon emissions grew by 10.7 Mg/km^2 per year in state parks compared with non-designated candidate parks.

Comparing columns 1 and 2 of table 3.4, carbon loss point estimates for federal and state parks match the signs of deforestation rate estimates, but the relative magnitude of federal effects grew considerably more compared to state effects when estimated using carbon loss as the dependent variable. Simply dividing carbon loss coefficients by deforestation rate coefficients suggests that forests in federal parks prevented deforestation in carbon-dense forests containing approximately 200 Mg of carbon per hectare—or a carbon density 38 percent greater than forests affected by state parks, which contain approximately 145 Mg per hectare. All remaining coefficients are insignificant in the carbon loss model. An F-test of the hypothesis that the six designation coefficients are jointly zero is rejected at a 1-percent significance level with a p-value of 0.00002.

Without additional context, it is unclear whether federal parks are sited in areas with high carbon density, or whether federal parks effectively protect those areas with the highest carbon storage. The relative carbon density of state and federal parks at the time of site selection, but prior to park designation, can help identify the relative impact of selection bias on carbon-based treatment effects. In 2000, prior to park designation, the average carbon density of federal parks was 167 Mg of carbon per hectare, surprisingly similar to that of state parks with 161 Mg of carbon per hectare. Since federal and state parks had similar carbon densities prior to site selection, it appears that federal parks effectively protect the high-carbon forests, while state parks permit deforestation primarily in low-carbon forests.

Relative to forest loss treatment effects, carbon estimates suggest that park networks are more effective, in aggregate, because avoided forest loss occurs in high-carbon areas and increased forest loss occurs in low-carbon areas. Table 3.4 shows that federal parks prevented deforestation in high-quality forests, a mean carbon density more than 30 Mg higher than average, perhaps by restricting access to primary forests in the park interior. In contrast, state parks facilitated deforestation in low-quality forests with a mean carbon density of about 15 Mg lower than average, perhaps due to forest clearing activities in secondary forests. These results are consistent with a recent study that finds deforestation of unprotected areas in Rodonia occurred in in forests with lower than average carbon density (Song et al., 2014).

Estimates of avoided carbon emissions have direct policy implications.[24] For example, in the context of REDD+, one could use estimates in column 2 of table 3.4 to estimate the avoided carbon emissions from all federal parks in the Amazon designated between years 2001 and 2009. During these 8 years, 38 federal parks were established over 310,733 km^2 of the Amazon forest. Using the estimate of 18.5 Mg/km^2 implies avoided carbon emissions of approximately 5.7 million metric tons per year, which at a carbon price of $20 per metric ton has an annual value of $115 million from post-2000 federal parks. A comparable estimate shows that the adverse selection of state park designations actually detracted from the net value of carbon sequestration by $53 million per year from post-2000 state parks. This calculation assumes state parks generate 2.7 million metric tons of carbon emissions per year, using estimates of 10.5 Mg/km^2 additional carbon emissions sourced from each of the 250,146 km^2 contained in the 52 state parks established after 2001.

In sum, results suggest that parks designated between 2001 and 2009 collectively reduced annual carbon emissions approximately 3 million metric tons per year through 2012. By comparison, these estimates are lower than the avoided carbon emissions reported in another recent study. Ferraro et al. (2015) estimate that Brazil's post-2000 protected areas reduced carbon emissions by 749 million metric tons between 2000 and 2008. Using results in this study, the cumulative effects over a 9-year period would

---

[24] It should be noted that the estimated forest carbon loss only represents "committed carbon emissions" rather than the total carbon emissions from deforestation. A full accounting of carbon emissions associated with deforestation must include estimates of other fundamental carbon pools, such as soil carbon, dead organic matter, and carbon emissions from the succeeding land uses after deforestation (e.g. ranching), etc.

imply a reduction in carbon emissions of 27 million metric tons—less than 4 percent of total reductions estimated by Ferraro et al. (2015). Although this study uses a different source of carbon data than Ferraro et al. (2015), disagreement in carbon-based treatment effects is most likely due to differences in methodology. Compared to cross-sectional methods, superior panel estimators show a loss of significance for indigenous parks and a reversal in sign for state park effects. If one simply applied the carbon-effects estimated for federal parks to all 1,166,106 km^2 of total park area established between 2001 and 2009, and multiplied by 9-years, the estimates in this study would increase to 194 million metric tons per year—a magnitude that is reasonably similar to estimates of Ferraro et al. (2015).

Inferring Causality via the Timing of Park Designation

Panel data techniques inherently identify effects based on changes in treatment status over time. In the case of multiple experiments, where parks are designated at different points in time, panel analyses can also provide insights on whether park effectiveness has evolved over time.

The difference-in-difference estimates in equation (3.1) provide no sense of the dynamics of park designation and forest loss: how quickly deforestation declines after a park is designated and whether this effect accelerates, stabilizes, or reverts to a mean. If deforestation declines lead to the designation of parks, rather than vice versa, the previous estimates would obscure this reverse causality. On the other hand, if a temporary surge in deforestation leads to the designation of parks, then previous estimates would obscure this reversion to the mean. To explore these dynamics, table 3.5 provides estimates of a subset of the models in table 3.4, augmented with leads and lags of park designation. Specifically, I add indicator variables for years 1, 2, and 3 before designation, years 0-3 after designation, and year 4 forward.[25] Of these eight indicator variables, the first seven are equal to one only for a single year, while the final variable is equal to one in each year starting with the fourth year of designation.

The first column of table 3.5 presents the base specification augmented with the leads and lags. The coefficients on the two and three year designation lags are close to zero, showing little evidence of an anticipatory response within parks about to receive federal protection status[26]. In the year prior to adoption, however, deforestation decreases substantially by -0.2 percent and remains at this level during the year of designation. Since this initial decline occurs within one year of designation, the negative coefficient on the one year designation lag may simply reflect measurement error in the forest loss dataset. This follows because I am unable to identify the date of forest loss within a time interval shorter than a year, so forest loss occurring in the time period surrounding designation cannot be accurately assigned to a pre-designation or post-designation regime. After the year of designation, however, annual deforestation rates drop an additional -0.3 percentage points in the subsequent two years, after which annual effects stabilize to average -0.5 percentage points in year 4 forward.

---

[25] Additional models were estimated with additional leads and lags, including eleven indicator variables for 1, 2, 3, and 4 years before designation, years 0-5 after designation, and year 6 forward. Alternative estimates are similar in sign, significance, and magnitude to the estimates presented in the model with eight indicator variables representing leads and lags and so are not included here.

[26] The three-year lead has a substantially smaller standard error than the two-year lead, although both are statistically insignificant, and than other leads and lags. In alterative specifications, leads further from the year of designation consistently have relatively small standard errors. This may be due to limited variation in forest loss in the period before effects. In addition, earlier periods are less likely to have any variation caused by anticipation effects or measurement error in forest loss date assignment, and thus much less change in variation.

Subsequent columns repeat these estimates using a sample of parks established after 2001, and using carbon loss as the dependent variable for the sample of all parks established since 1985 and a sample restricted to parks established after 2001. The pattern of coefficients is comparable for each case, providing robust evidence that designation of federal parks has led to a reduction in forest loss rather than vice versa. Figure 3.5 depicts this pattern, with 95-percent confidence bars, for carbon loss estimates in table 3.5 column 4, the preferred specification using all parks established since 1985. Anticipation effects are not indicated for 3 and 2 year leads. An initial effect of 30 Mg C/km^2/year begins with a 1-year lead and in the year of protection, followed by a second drop through the 2 year lag, followed by stabilization at levels between 70 and 80 Mg C/km^2/year where estimates become statistically significant at the 5-percent level.

Estimates of dynamic effects offer compelling evidence that park establishment has a near-immediate causal effect on avoided carbon emissions. The results on the timing of effects make two distinct contributions to the literature. First, the discrete onset of effects around the timing of designation strengthens causal claims that the park designation does, in fact, drive the results presented in this essay. Second, the quick onset of effects provides important policy insights that are missing from previous literature. In particular, it appears that "paper parks" work. Legal statues appear to become effective immediately at the time of designation, perhaps by eliminating access to titles for would-be land grabbers. Meanwhile, parks seem to be effective 2 years after designation—well before new infrastructure, staffing, and enforcement protocols take effect. These are planning steps that typically require 5 years or more to implement.

CONCLUSIONS

This essay makes three central contributions to existing literature. First, the fixed effects model finds that government-ownership is a key determinant of effectiveness. The surprising finding is that federal and state governments have dramatically different effects that are of opposite qualitative directions. Specifically, federal parks are immediately effective in reducing deforestation, while state parks cause a considerable increase in deforestation over time. Contrary to conventional wisdom, panel results suggest that the establishment of indigenous parks causes no change in deforestation.

Second, this essay demonstrates the limitations of cross-sectional matching estimators that dominate the deforestation literature. For example, cross-sectional matching estimators applied to the same data find that indigenous parks are very effective at reducing deforestation, which illustrates the inherent bias of studies that rely on cross-sectional matching. The findings in this essay suggest that indigenous people monitor forests regardless of legal park status, which will bias matching estimators that fail to control for the presence of indigenous populations.

Third, results highlight the importance of spatial variation in forest carbon storage. These results show that studies that measure only avoided deforestation are misleading as a reflection of avoided carbon emissions. This is particularly true given evidence that deforestation effects target either high-carbon or low-carbon forest areas across different governance regimes.

To date, the effectiveness of protected areas has been assessed using cross-section matching estimators. In this essay, new evidence is derived using quasi-experimental models on two panel datasets of deforestation. The results reveal dramatic new insights into the importance of government oversight of protected areas and the jurisdiction from which government oversight is provided—findings that counter economists' prior notions of the avoided deforestation from park designation. I extend the analysis to estimate avoided carbon emissions. I show that this key metric, which is of much more policy importance than deforestation, varies considerably from deforestation trends. For the both science and policy communities, I underscore uncertainty in policy evaluation based on imperfect satellite-derived deforestation data products compared to data products that measure carbon.

REFERENCES FOR ESSAY 3

Andam, K., P. Ferraro, A. Pfaff, G. Sanchez-Azofeifa, and J. Robalino. (2008). Measuring the Effectiveness of Protected Area Networks in Reducing Deforestation. *Proceedings of the National Academy of Sciences,* 105(42), 16089-16094.

Blackman, A. (2013). Evaluating Forest Conservation Policies in Developing Countries Using Remote Sensing Data: An Introduction and Practical Guide.*Forest Policy and Economics*, 34, 1-16.

Busch, J., K. Ferretti-Gallon, J. Engelmann, M. Wright, K.F. Austin, F. Stolle, and A. Baccini. (2015). Reductions in Emissions from Deforestation from Indonesia's Moratorium on New Oil Palm, Timber, and Logging Concessions. *Proceedings of the National Academy of Sciences,* 12(5), 1328-1333.

Deininger, K., and B. Minten. (2002). Determinants of Deforestation and the Economics of Protection: An Application to Mexico. *American Journal of Agricultural Economics,* 84(4), 943-60.

Farr, T.G., et al. (2007). The Shuttle Radar Topography Mission. *Reviews of Geophysics,* 45. Available at http://www2.jpl.nasa.gov/srtm.

Ferraro, P., and M. Hanauer. (2011). Protecting Ecosystems and Alleviating Poverty with Parks and Reserves: Win-win' or Tradeoffs? *Environmental and Resource Economics,* 48, 269-28.

Ferraro, P. J., M.M. Hanauer, D.A. Miteva, G.J. Canavire-Bacarreza, S.K. Pattanayak, and K.R. Sims. (2013). More Strictly Protected Areas Are Not Necessarily More

Protective: Evidence from Bolivia, Costa Rica, Indonesia, and Thailand. *Environmental Research Letters*, 8(2), 025011

Ferraro, P.J., M. Hanauer, D. Miteva, J. Nelson, S. Pattanayak, C. Nolte, and K. Sims. (2015, In Press). Estimating the Impacts of Conservation on Ecosystem Services by Integrating Modeling and Evaluation. Accepted. *Proceedings of the National Academy of Sciences*.

Fischer, G., H. van Velthuizen, F. Nachtergaele, and S. Medow. (2002). Global Agro-Ecological Zones (Global- AEZ). Food and Agricultural Organization/International Institute for Applied Systems Analysis (FAO/IIASA). See (http://www.iiasa.ac.at)

Nepstad, D., McGrath, D., Stickler, Alencar,C., Azevedo, A., Swette, B., Bezerra, T., Maria DiGiano, M., João Shimada, J., Ronaldo Seroa da Motta, R., Armijo, E., Castello, L., Brando, P., C. Hansen, M., McGrath-Horn, M., Carvalho, O., and Hess, L. (2014). Slowing Amazon Deforestation through Public Policy and Interventions in Beef and Soy Supply Chains. *Science*, 344(6188), 1118-1123.

Nolte, C., A. Agrawal, and P. Barreto. (2013a). Setting Priorities to Avoid Deforestation in Amazon Protected Areas: Are We Choosing the Right Indicators?.*Environmental Research Letters*, 8(1), 15039-15045.

Nolte, C., A. Agrawal, K., Silvius, and B. Soares-Filho. (2013b). Governance Regime and Location Influence Avoided Deforestation Success of Protected Areas in the Brazilian Amazon. *Proceedings of the National Academy of Sciences*, 110(13), 4956-4961.

Pfaff, A., J. Robalino, E. Lima, C.Sandoval, and L. Herrera. (2014). Governance, Location and Avoided Deforestation from Protected Areas: Greater Restrictions

Can Have Lower Impact Due to Differences in Location. *World Development*, 55, 7-20.

Saatchi, S.S., N.L. Harris, , S. Brown, M. Lefsky, E.T. Mitchard, W. Salas, B.R. Zutta, W. Buermann, S.L. Lewis, S. Hagen, and others. (2011). Benchmark Map of Forest Carbon Stocks in Tropical Regions Across Three Continents. *Proceedings of the National Academy of Sciences*, 108, 9899-9904.

Sims, K. (2010). Conservation and Development: Evidence from Thai Protected Areas. *Journal of Environmental Economics and Management,* 60, 94-114.

Soares-Filho, B., Moutinho, P., Nepstad, D., Anderson, A., Rodrigues, H., Garcia, R., Dietzsch, L., Merry, F., Bowman, M., Hissa, L., Silvestrini, R., and Maretti, C. (2010). *Role of Brazilian Amazon protected areas in climate change mitigation. Proceedings of the National Academy of Sciences*, 107, 10821-10826.

Stern, N., *The Economics of Climate Change: The Stern Review*. (2006). Cambridge, UK: Cambridge University Press.

Townshend, J. R. G., M. Carroll, C. Dimiceli, R. Sohlberg, M. Hansen, and DeFries, R. (2011). *MODIS Vegetation Continuous Fields product MOD44B, Collection*, 5.

Uchida, H., and A. Nelson. (2009). Agglomeration Index: Towards a New Measure of Urban Concentration. Background paper for the World Bank's World Development Report. See (http://bioval.jrc.ec.europa.eu/products/gam/index.htm)

United Nations Environment Program, World Conservation Monitoring Center (UNEP-WCMC). (2007). World Database on Protected Areas. Cambridge, UK: World Conservation Union (IUCN) and UNEP-World Conservation Monitoring Center.

Van der Werf, G., Morton, D., DeFries, R., Olivier, J., Kasibhatla, P., Jackson, R., Collatz, J., and Randerson, J. (2009). CO_2 Emissions from Forest Loss. *Nature Geoscience*, 2, 737-738.

FIGURES FOR ESSAY 3

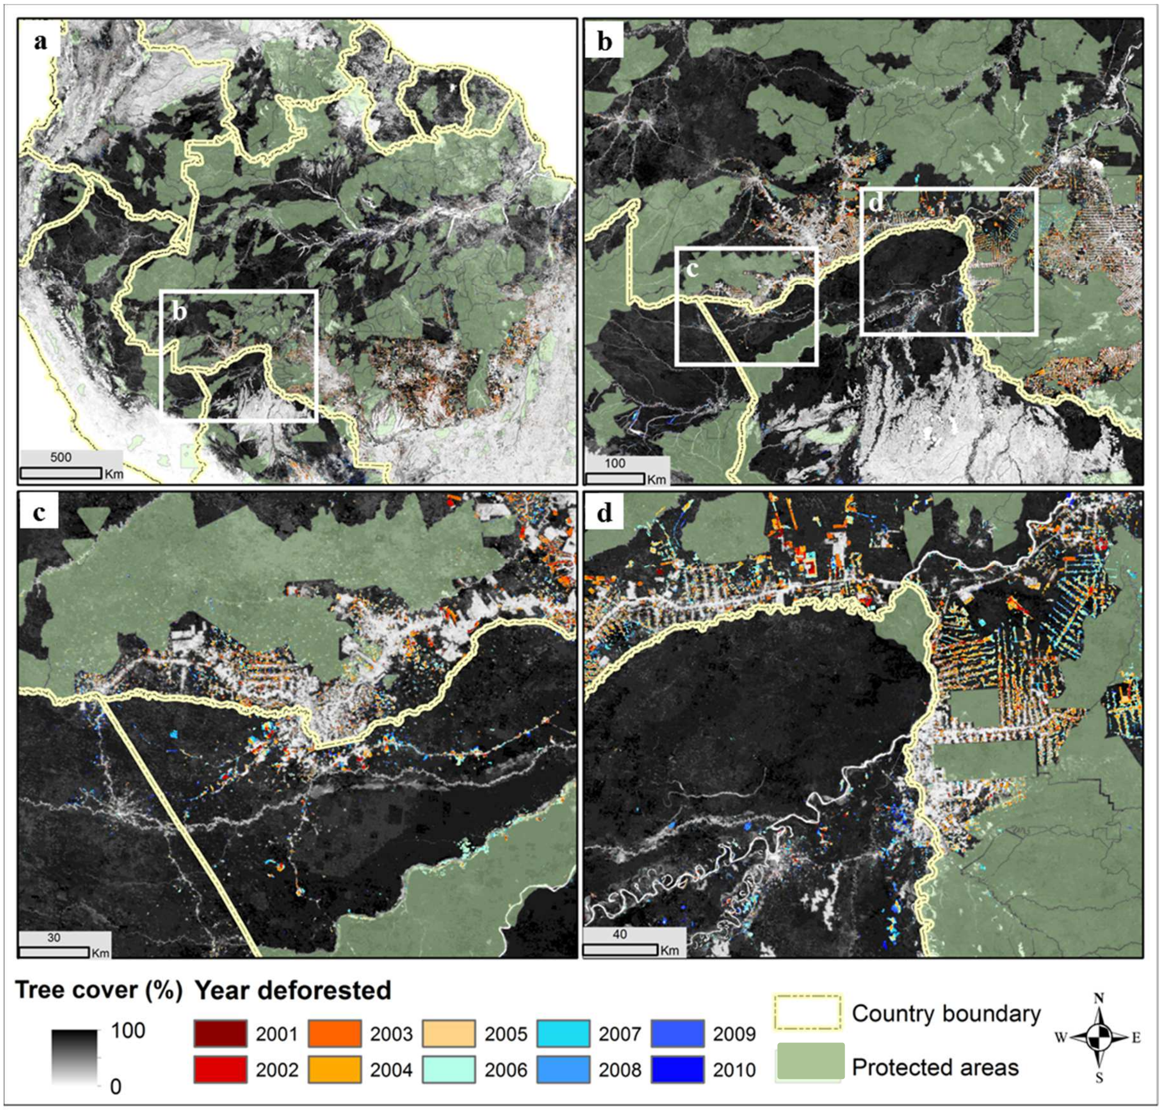

Figure 3.1 Forest, deforestation and protected areas in the Amazon basin, where over 50% of forests have been designated as protected areas. The four map layers are overlaid in the order from top to bottom of country boundaries, protected areas, deforestation year, and tree cover percentage in the year 2000. Panel (a) is a full-view of protected areas in the entire study area. Panel (b) is a close-up in the Brazil/Bolivia/Peru tri-national border region where forests on the Brazil side are either cleared or protected. Panel (c) zooms-in over the city of Cobija, the capital of the Bolivian Pando Department. The inter-oceanic highway begins in this region. Panel (d) zoom-in over the city of Guayaramerin, where more deforestation is observed on the Bolivia side after 2006.

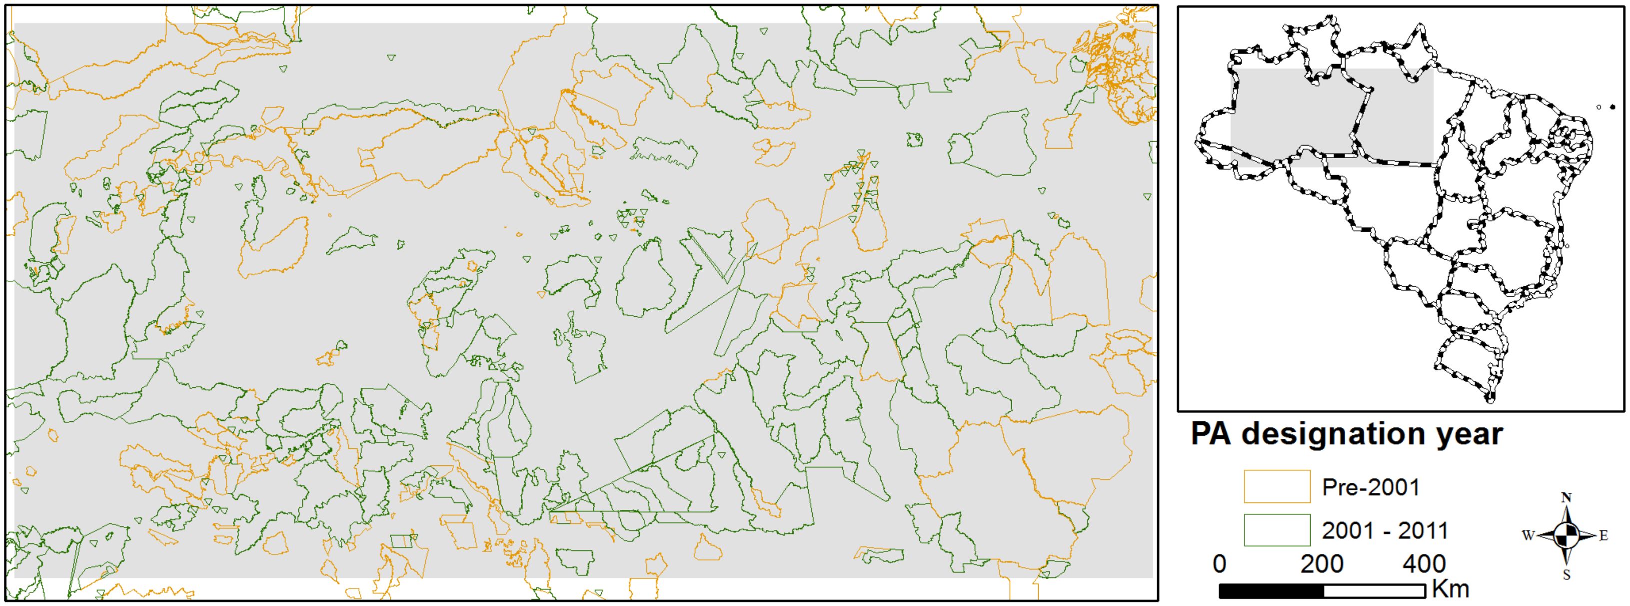

Figure 3.2 A sub-sample that includes approximately 50% of the Amazon study area, centered over the two largest Brazilian states of Amazonas and Pará. Parks designated as pre-2001 (the comparison group) are shown as orange polygons, while parks designated post-2001 (the treatment group) are shown as green polygons.

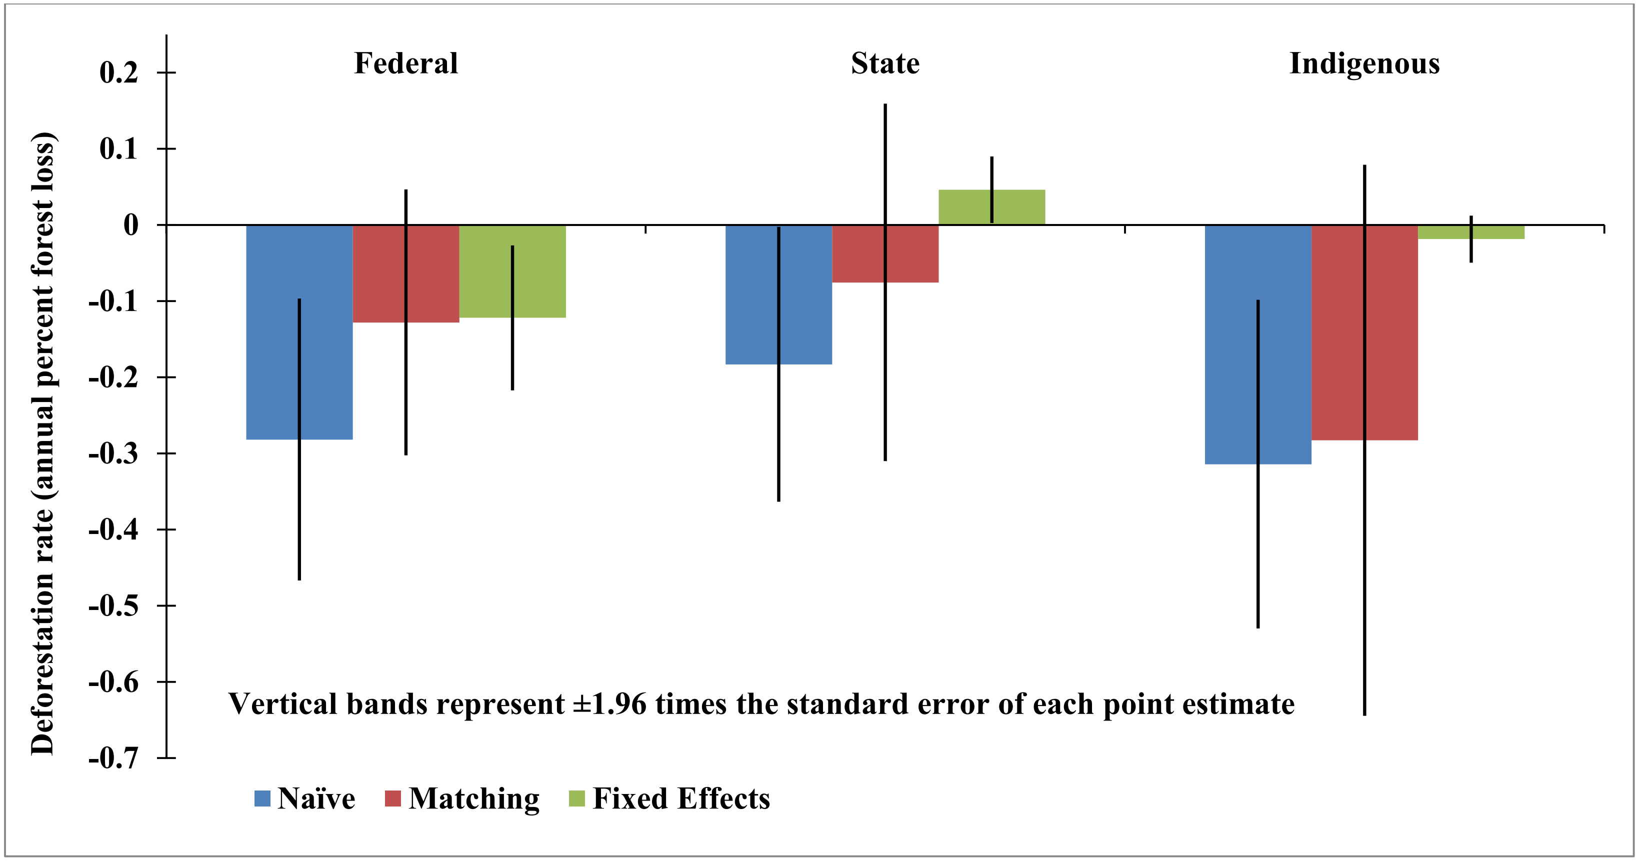

Figure 3.3 Effects of government park designations on percentage forest loss estimated from naïve, matching, and fixed-effects models.

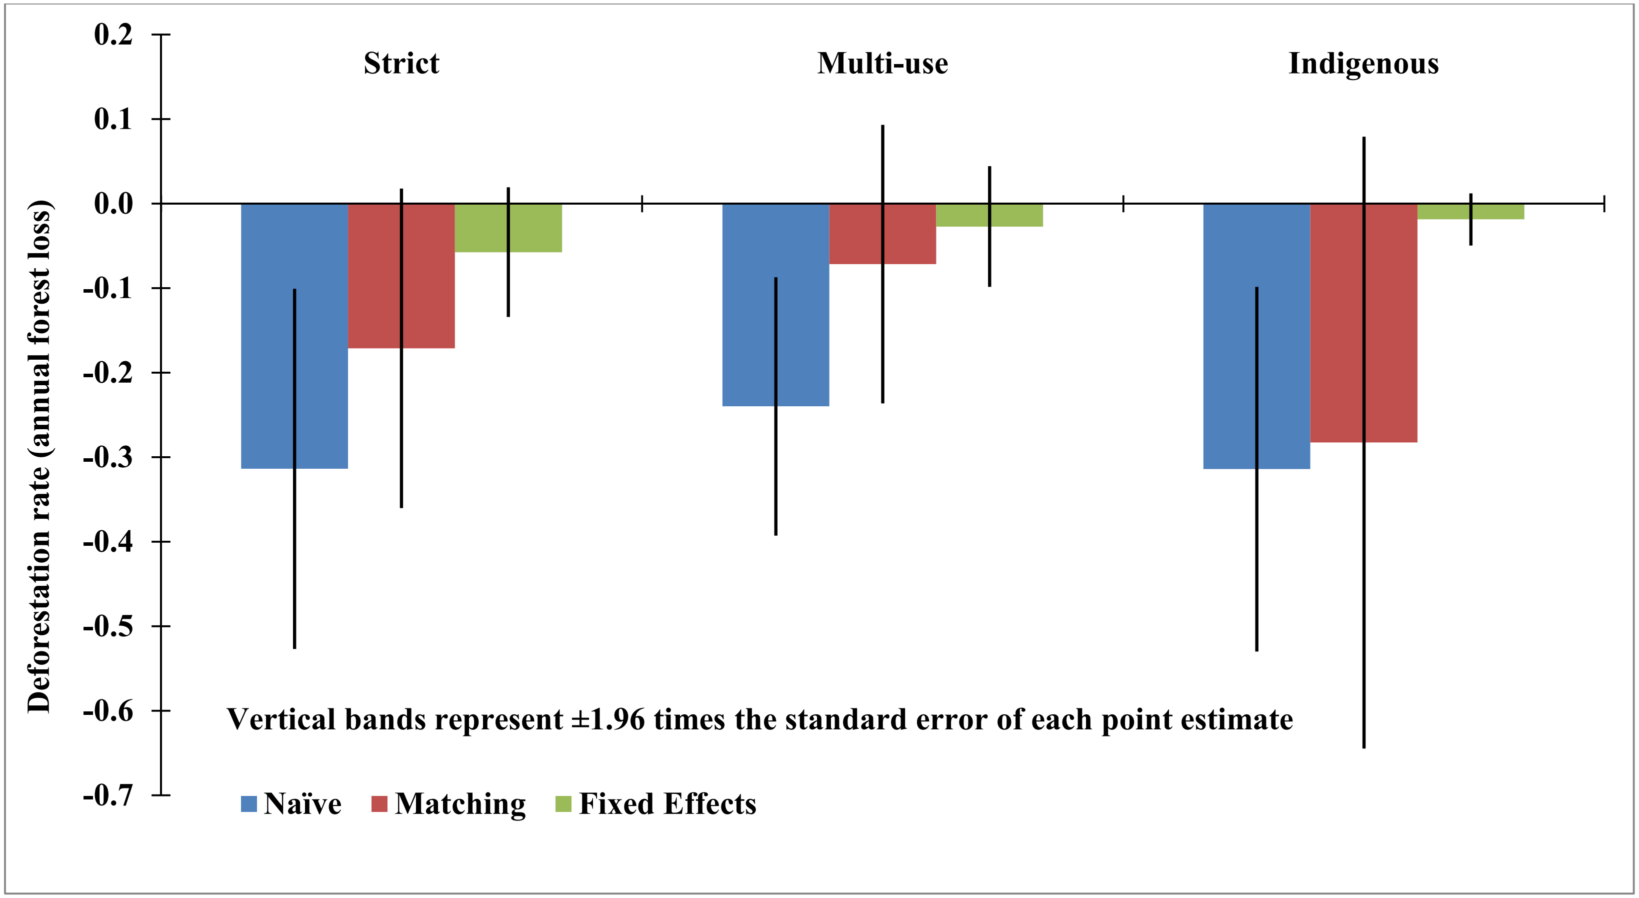

Figure 3.4 Effects of park strictness designations on percentage forest loss estimated from naïve, matching, and fixed-effects models.

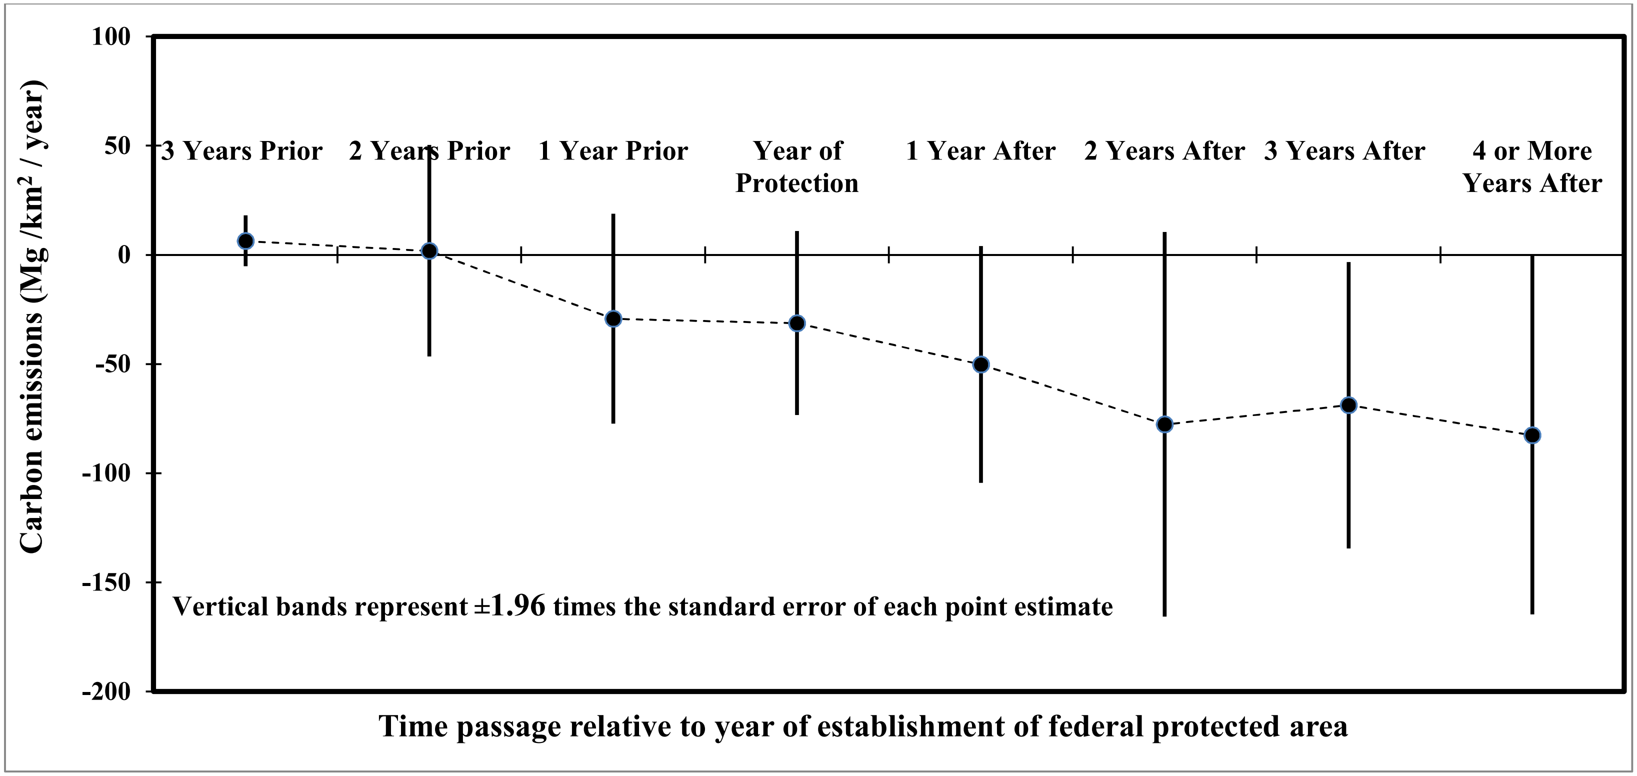

Figure 3.5 Estimated effects of federal park designation on carbon loss for years before, during, and after park designation. (See table 3.5, model 4).

TABLES FOR ESSAY 3

Table 3.1 The Estimated Effect of Park Designations on Annual Deforestation Rates in the Amazon, 2000-2012

	(1)		(2)		(3)		(4)		(5)		(6)	
Park designation:												
Any type	-0.039	***	-0.027	**	-0.025	*						
	(0.014)		(0.011)		(0.013)							
Federal							-0.071	***	-0.066	***	-0.122	**
							(0.021)		(0.022)		(0.049)	
State							-0.059		-0.035		0.046	**
							(0.038)		(0.035)		(0.022)	
Indigenous							-0.041	**	-0.022		-0.019	
							(0.017)		(0.019)		(0.016)	
Non-government							0.044		0.025		0.000	
							(0.031)		(0.017)		(0.028)	
Park and year dummies	Yes		Yes		Yes		Yes		Yes		Yes	
State x year dummies	No		Yes		Yes		No		Yes		Yes	
Park x time trends and park x time^2 trends	No		No		Yes		No		No		Yes	
R^2	0.015		0.070		0.378		0.02		0.072		0.377	

Source: The dependent variable is constructed from Hansen et al. 2012, and is defined as average percent forest loss per year in each of 610 parks in each of 12 years (7,320 observations). Ordinary least squares estimates are weighted by park area. The sample includes parks established since 1985. Standard errors in parentheses are clustered by park to allow for arbitrary correlation of residuals within each park. Asterisks denote significance at levels of 1, 5, and 10 percent (***, **, *).

Table 3.2 The Estimated Effect of Park Designations on Annual Deforestation Rates in the Amazon, 2000-2012: Testing the Effect of Government Ownership versus Protection Strictness.

	(1)		(2)	(3)	
Park designation:					
Federal	-0.122	**		-0.093	*
	(0.049)			(0.052)	
State	0.046	**		0.074	**
	(0.022)			(0.033)	
Non-government	0.000				
	(0.028)				
Indigenous	-0.019		-0.019	-0.019	
	(0.016)		(0.016)	(0.016)	
Strict			-0.057	-0.033	
			(0.039)	(0.033)	
Multi-use			-0.027	-0.025	
			(0.036)	(0.030)	
Not reported			0.009	0.009	
			(0.033)	(0.033)	
R^2	0.377		0.609	0.636	

Source: For the dependent variable, see Hansen et al. 2012
Dependent variable: percent forest loss per year; n=7,320. Sample includes parks established between 1985 and 2009. Ordinary least squares estimates include weights for park area. Standard errors in parentheses are clustered by park to allow for arbitrary correlation of residuals within each park. Asterisks denote significance at levels of 1, 5, and 10 percent (***, **, *). All models include park main effects, state-by-year effects, and park specific linear and quadratic time trends. Omitted treatment group in model 3 is non-government park designations.

Table 3.3 The Estimated Effect of Park Designations on Annual Deforestation Rates in the Amazon: Cross-section Estimators

Park designation:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Federal	-0.282 **	-0.332 *	-0.128				-0.235 *	-0.218	-0.046
	(0.094)	(0.154)	(0.089)				(0.124)	(0.139)	(0.052)
State	-0.183 *	-0.239 **	-0.075				-0.140	-0.104	0.009
	(0.092)	(0.108)	(0.120)				(0.116)	(0.127)	(0.093)
Non-government	-0.301 **	0.003	-0.036				-0.271 *	0.086	0.013
	(0.122)	(0.068)	(0.046)				(0.139)	(0.102)	(0.028)
Indigenous	-0.314 **	-0.551 **	-0.289	-0.314 **	-0.548 **	-0.283	-0.314 **	-0.557 **	-0.287
	(0.110)	(0.223)	(0.188)	(0.110)	(0.225)	(0.185)	(0.110)	(0.230)	(0.189)
Strict				-0.314 **	-0.285 **	-0.171	-0.093	-0.136	-0.150 **
				(0.109)	(0.117)	(0.096)	(0.059)	(0.090)	(0.060)
Multi-use				-0.240 **	-0.267 *	-0.072	-0.028	-0.144	-0.058
				(0.078)	(0.131)	(0.084)	(0.074)	(0.144)	(0.046)
Not reported				-0.264 *	0.002	-0.001			
				(0.126)	(0.045)	(0.044)			
Geographic covariates:									
Tree cover in 2000		-0.657	-0.096		-0.534	-0.067		-0.611	-0.047
(percent)		(0.381)	(0.105)		(0.431)	(0.080)		(0.411)	(0.093)
Elevation		-0.057	-0.005		-0.042	-0.008		-0.056	-0.007
(meters)		(0.039)	(0.031)		(0.041)	(0.034)		(0.039)	(0.033)
Slope		0.040 **	0.018		0.047 **	0.027		0.042 **	0.025
(degrees)		(0.017)	(0.015)		(0.018)	(0.017)		(0.017)	(0.016)
Travel time to city		0.088 **	0.026 *		0.121 **	0.026 *		0.088 **	0.029 **
(days)		(0.043)	(0.014)		(0.058)	(0.014)		(0.043)	(0.014)
Crop suitability		-0.041	-0.044 *		-0.033	-0.042 *		-0.038	-0.040 *
(index)		(0.049)	(0.022)		(0.050)	(0.019)		(0.050)	(0.021)
Precipitation		0.030 **	0.000		0.400 *	0.000		0.300 **	0.000
(cm/month)		(0.014)	(0.005)		(0.167)	(0.064)		(0.140)	(0.054)
Mean temperature		-0.008	0.000		-0.050	-0.010		-0.080	-0.010
(C)		(0.007)	(0.006)		(0.077)	(0.062)		(0.073)	(0.061)
region x year	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
matched sample	No	No	Yes	No	No	Yes	No	No	Yes
R^2	0.078	0.814	0.335	0.077	0.81	0.34	0.078	0.815	0.342

Dependent variable: percent forest loss per year. Unit of observation is km2-year. All samples include parks established since 1985. Models 1-2, 4-5, and 7-8 also include unprotected areas (N=84,625,901); models 3, 6, and 9 include only matched unprotected areas (N=25,223,513). Standard errors in parentheses are clustered by park. All models include year effects. Significance: *1%, **5%, ***10%

Table 3.4 The Estimated Effect of Park Designations in the Amazon: Comparisons Of Percent Deforestation and Carbon Loss Estimates Across Data Sets

	(1)		(2)		(3)		(4)	
Park designation:								
Federal	-0.093	*	-18.460	**	-0.050		-10.738	**
	(0.052)		(9.251)		(0.033)		(5.254)	
State	0.074	**	10.690	**	0.043		7.465	*
	(0.033)		(5.385)		(0.027)		(4.464)	
Non-government								
Indigenous	-0.019		-3.286		-0.002		-1.722	
	(0.016)		(2.519)		(0.018)		(2.337)	
Strict	-0.033		-3.585		-0.008		-0.560	
	(0.033)		(5.681)		(0.017)		(2.802)	
Multi-use	-0.025		-4.105		-0.007		-1.321	
	(0.030)		(5.078)		(0.022)		(3.705)	
Not reported	0.009		1.098		-0.043	**	-5.657	**
	(0.033)		(4.431)		(0.018)		(2.439)	
R^2	0.636		0.377		0.609		0.636	

Source(s) - Hansen et al. forest loss 2000-2012 (models 1 and 2).
Dependent variable(s): percent forest loss per year (models 1 and 3), or metric tons of carbon loss per km^2 per year (models 2 and 4); n=7320. Samples include all parks established since 1985. Ordinary least squares estimates include weights for park area. Standard errors in parentheses are clustered by park to allow for arbitrary correlation of residuals within each park. Asterisks denote significance at levels of 1, 5, and 10 percent (***, **, *). All models include park main effects, state-by-year effects, and park specific linear and quadratic time trends. Omitted treatment group in is non-government park designations.

Table 3.5 The Estimates Effects of Federal Park Designation on Annual Deforestation Rates and Annual Carbon Loss in the Amazon, for Parks Designated After 2002 and 1985

	(1)	(2)	(3)	(4)
Federal park leads and lags:				
Status $change_{t+3}$	0.034	0.035	5.87	6.45
	(0.036)	(0.032)	(06.82)	(5.92)
Status $change_{t+2}$	-0.047	-0.012	- 2.46	1.90
	(0.153)	(0.149)	(24.78)	(24.72)
Status $change_{t+1}$	-0.202	-0.177	-34.48	-29.18
	(0.174)	(0.158)	(27.64)	(24.53)
Status $change_{t0}$	-0.220	-0.191	-37.78	-31.22
	(0.161)	(0.145)	(24.50)	(21.50)
Status $change_{t-1}$	-0.314	-0.303 *	-55.33 *	-50.20 *
	(0.201)	(0.183)	(31.04)	(27.66)
Status $change_{t-2}$	-0.500	-0.474	-84.60 *	-77.60 *
	(0.327)	(0.294)	(50.76)	(44.92)
Status $change_{t-3}$	-0.447 *	-0.413 *	-78.02 **	-68.88 **
	(0.243)	(0.215)	(39.07)	(33.49)
Federal park $status_{t-4\ forward}$	-0.535 *	-0.498 *	-92.67 *	-82.60 **
	(0.302)	(0.270)	(47.77)	(41.87)
H_0: $designation_{(t0-t4)} = 0$	0.000	0.000	0.000	0.000
R^2	0.371	0.369	0.390	0.367
No. Observations	3,684	7,320	3,684	7,332
No. PAs	307	610	307	611

Source: For the dependent variable, see Hansen et al. 2012.
Dependent variable(s): percent forest loss per year (models 1 and 2), or metric tons of carbon loss per km^2 per year (models 3 and 4). Samples include parks established since 2002 (models 1 and 3), or parks established since 1985 (models 2 and 4). Ordinary least squares estimates include weights for park area. Standard errors in parentheses are clustered by park to allow for arbitrary correlation of residuals within each park. Asterisks denote significance at levels of 1, 5, and 10 percent (***, **, *). All models include leads and lags for strict and multi-use designations, park main effects, state-by-year effects, and park specific linear and quadratic time trends.