Do Energy Savings Grow on (Shade) Trees?[1]

Joe Maher[2]

JOB MARKET PAPER

Abstract

In warm climates, trees could provide natural air-conditioning by shading homes. Yet, there is little rigorous empirical evidence on the shade benefits of green infrastructure. This paper uses data on tree removal permits from a tree protection policy in Gainesville, Florida, between 2000 and 2016, and applies a difference-in-difference estimation model to examine whether tree shade reduces electricity demand. I find that a typical tree removal leads to 4-6% more electricity consumption annually, and 8-10% more during summer months of peak air-conditioning demand. Effects increase with the amount of shade loss: from zero effect (no shade) to 8-12% (large shade loss) and 10-20% (very large shade loss). I then use these estimates to calculate the private benefits of shade trees ($450 to $1,900) and find that they are similar to other energy efficiency investments (retrofits, building codes) assessed using comparable evaluation methods. I also calculate the social benefits (shade benefits to neighbors, avoided emissions, and avoided generation costs) of the tree protection policy and find that they are equal to the city expenditures on this policy. Overall, this study presents a road map for utilities and policy makers to assess returns from investments in green infrastructure.

---

[1] This version: November 28, 2016. For the latest version, please see www.josephamaher.com. This paper draws upon the second chapter of my dissertation. I thank my advisors Richard Just, Maureen Cropper, Sebastien Houde, Rob Williams, and Ted McConnell for their guidance and support. I am also grateful to Magda Tsaneva, Xiaopeng Song, Binesh Maharjan, Ryan Bolt, Stephanie McGuire, Philippe Marchand, Leah Bevis, and participants in various seminars at the University of Maryland, AERE summer meetings, AEA summer meetings, Camp Resources, and the Socio-Environmental Synthesis Center. In addition, many individuals in the City of Gainesville, Gainesville Regional Utilities, and University of Florida have been very generous with their time. I gratefully acknowledge funding for this work from: a University of Maryland graduate fellowship, the USDA Forest Service Urban Forestry South, the Resources for the Future Postdoctoral Fellows Program, and the National Science Foundation.

[2] Postdoctoral Fellow, National Socio-Environmental Synthesis Center (SESYNC), jmaher@sesync.org.

Buildings account for 42 percent of energy use and 38 percent of CO_2 emissions in the United States, making buildings' energy efficiency a major 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 in demand-side management programs that subsidize energy-saving improvements for homes and businesses. This surge in public investment has spurred a renewed interest in building codes, demand-side management policies, and behavioral interventions, but studies typically show that these approaches have limited effects and fail cost-benefit tests (e.g. Fowlie et al. 2016; Jacobsen and Kotchen 2012; Davis et al. 2014; Alcott 2011). Despite the well-established limitations of conventional approaches, empirical research continues to overlook the role of environmental factors known to affect energy consumption. For example, sun exposure significantly affects building energy performance and drives numerous architectural and engineering decisions, such as building orientation, roof reflectance, heating, ventilation, and shading.

Considering the large role that environmental factors and green infrastructure play in energy efficiency, surprisingly little empirical research considers how they help conserve energy. Trees shade buildings during summer months, reducing air-conditioning usage, and can also serve as windbreaks during winter months, reducing heating requirements. Today, more than 3,400 U.S. cities have ordinances that protect urban tree canopies. Such ordinances may be effective demand-side management policies if shade trees and windbreak trees substantially improve buildings' energy efficiency (Tree City USA 2015).

Current empirical research suggests a correlation between tree cover and slightly lower summer energy consumption but fails to find large enough effects to attract policy attention. Donovan and Butry (2009) use monthly billing data from 460 homes in Sacramento, California, and find that homes with trees shading their south and west walls 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. Yet, these studies lack causal inference; both identify effects based on cross-sectional comparisons of energy use between houses with varying levels of tree shade. Moreover, both papers include only basic controls for home characteristics (e.g., house square footage and presence of a pool) and rely almost entirely on average energy use in early spring to serve as a proxy for baseline energy use without shade. Cross-sectional estimators will be biased if tree size is correlated with other variables that affect energy efficiency, such as the age of a structure. Pandit and Laband (2010) also require participants' permission to have their electricity meters read, introducing selection bias into the estimates.

This paper uses quasi-experimental variation in tree shade associated with residential tree removals to identify the energy savings from tree shade in Gainesville, Florida. Since 1999, the Gainesville tree ordinance has been regulating the removal of high-value shade trees within the city. A permit is required for removal of any mature tree within city limits. This analysis uses 2000 to 2016 data on household energy consumption, date of tree removals, and canopy cover from multi-temporal aerial imagery to quantify how shade-loss (tree removal) affects energy-consumption.

I apply a difference-in-difference model that estimates the treatment effects of residences with tree removals between 2001 and 2011 using households which had a tree removal later (between 2012 and 2016) as controls. Results suggest that tree removals cause a 5 percent increase in annual electricity consumption, yet have no effect on natural-gas consumption. Moreover, the difference in energy effects by month-of-year, are consistent with increased consumption of electricity for air-conditioning (summer electricity use), which is the main end-use affected by cooling shade, and no effect on winter energy demand. I further examine the heterogeneous treatment effects using variation in loss of shade. I find that monthly energy effects increase with amount of shade loss: from 8-12% (large canopy area shade loss), to 15-20% (very large canopy area shade loss). Yet, canopy size has no effect on energy during months when a tree does not directly shade a building.

After estimating the energy-savings from tree shade, I calculate the private benefits of a shade tree based on reduced energy expenditures from extending the life of a tree. Depending on tree size and location, the private benefits of shade trees range from $450 to $1,900 and are comparable to benefits from common energy-efficiency investments, such as pool pumps, air conditioners, and refrigerators, estimated using comparable difference-in-difference methods (Maher 2016).

Finally, I conduct a cost benefit analysis of city tree protection policies based on city expenditures and social benefits related to energy savings. In contrast to the private calculation, the social one accounts for spillover benefits of shade to neighboring residences, the benefits of reduced greenhouse gas and local pollutant emissions, and the avoided costs of investments in energy generation infrastructure. When the monetized social benefits are added to private benefits, the annual rate of return that would justify tree protection investments ranges between 0.7% and 8.9%, depending on the assumed policy scenario. This implies that the costs of enforcing the Gainesville tree ordinance could be justified based upon tree shade benefits alone. Finally, I also calculate the average cost per ton of avoided CO_2 and average cost per kilowatt-hour of avoided electricity demand under a range of assumptions. The plausible cost per ton of avoided CO_2 ranges from -\$57 and \$319, and cost per avoided kWh ranges from \$0.01 to \$0.23. These estimates suggest that tree protection has potential to be a cost-effective climate and demand side management policy, albeit an ambiguous potential. These results have immediate policy relevance for Gainesville, a city that is receiving long-standing tree ordinances and has a municipal owned utility undergoing major infrastructure investments to meet growing demand.

Overall, the analysis makes several contributions by providing robust causal evidence that tree shade affects air-conditioning demand, by showing that trees offer comparable savings as energy efficiency investments, and by performing a cost-benefit analysis to derive the private and social returns from tree protection policies.

This paper proceeds as follows: Section 2 describes the data sources and provides summary statistics. Section 3 discusses my estimation strategy and econometric model. Section 4 reports results. Section 5 develops measures of the returns to tree shade based on previously estimated results, and section 6 concludes.

2 DATA

2.1 Energy Data

This study uses residential utility data from January 2000 to May 2016 for residential customers serviced by Gainesville Regional Utilities (GRU). The utility data track monthly kilowatt-hours (kWh) of electricity and therms of natural gas for each residence. Utility data from 2000 to 2013 are downloaded from gainesvillegreen.com, a website designed to encourage energy conservation through the provision of publicly available information on household energy consumption.[3] To extend the study period, I merge utility data from 2013 to 2016; these utility data are updated each month by the City of Gainesville at gainesvilleopendata.com and include detailed address information.

2.2 Tree Removal Permit Data

This study uses residential tree removal permit data from March 2001 to March 2016. Permit data include the date (issue date, expiration date), location (property address), and number of trees (requested for removal, approved for removal). Therefore, the data reveal which residences removed trees and when the trees were removed. Information on the location and timing of tree removals, combined with billing data, allows a comparison of energy consumption before and after a change in tree cover. Permit data were acquired from the Gainesville city arborist, who visits the residence of each applicant to assess whether a tree is eligible for removal. The data set includes the street addresses of residents who were issued permits, the number of trees approved for removal by each permit, and the issue date for each permit. The permits are issued prior

[3] Unfortunately, the utility deleted records from May 2007 to February 2008 for approximately 80% of households.

to tree removal, of course, and not on the actual date a tree is cut down. In Gainesville, a tree removal permit is valid for six months after the issue date. Although the complete set of permits covers 1999 through 2016, the analysis includes only permits issued after 2001 so that tree-removing residences have a full 12 months of billing data both pre- and post-treatment.

To quantify the size and location of removed trees, an additional data set 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 is generated over 7,417 properties, including all tree removal sites and adjacent residences.[4] 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 (ACPA), the owners of each tree canopy area are then identified.

2.3 Structural Characteristics Data

In this study, structural characteristics are used to further confirm whether houses treated at different times have similar heating and cooling requirements. Data come from the ACPA property sale database and include 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 an indicator for central air-conditioning) and heating fuel (electricity, gas, or oil) relate to seasonal variation in

[4] See Appendix 1 for additional details about input data and methods used to classify tree canopy loss.

electricity and natural gas consumption, and type of roof (shingle, tile, wood, or metal) may affect heat transfer from solar radiation.

2.4 Mapping Canopy Loss

Given that the amount of shade a tree provides depends on 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 are used to construct a single measure of tree cover change for the period from 2001 to 2011. The resulting data set traces, at a spatial resolution of 1 meter by 1 meter, the patterns of tree cover loss. Appendix 1 describes how the tree cover change data set is constructed from raw aerial images.

2.4.1 Tree Location Variables

The identification strategy employed in this study requires details about the date of tree removal and the tree location relative to nearby buildings. 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 is assigned to each area of canopy loss based on GIS property boundary data provided by ACPA. In addition, a GIS map of building footprints was downloaded from the ACPA website. Figure 1 illustrates the three main GIS data sets: tree canopy loss, property boundaries, and building footprints. Maps of tree loss and building locations characterize tree position using two metrics: distance from house, and angle to house.

One way to quantify the size of tree removal is to calculate the total amount of canopy cover within a certain distance of the house. I calculate total area as a first step to test whether the effect of tree removal increases with the amount of canopy loss. One limitation of a total canopy area metric is that it fails to differentiate between trees located to shade homes and trees otherwise located. Failing to account for canopy cover location will bias effects downwards due to measurement error. This motivates alternative ways of quantifying the effect of tree size that accurately account for the shade provided by the tree prior to removal.

One alternate way to quantify the effect of tree removal is to geographically partition the area around the removed tree. According to this method, each tree removal is classified by one of 16 regions around each house (Figure 2). For example, the northern quadrant canopy loss variables would 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 would then be identical for remaining quadrants to the east, south, and west of a structure. If a tree canopy spans 2 or more regions, then the entire canopy area would be assigned to the region nearest to the house. For example, in Figure 2, the dark green tree is assigned to region N-1. These canopy loss variables permit flexible tests of how a residence's energy savings vary with differences in the canopy area, proximity, and direction.

In fact, in this paper, I do not treat tree removals as described above. One disadvantage of direction-by-distance dummies is the loss of precision from discretizing continuous measures of angle, distance, and area into arbitrary bins. Estimating the effect of shade loss using 16 categorical variables complicates inference and reduces precision of estimates. Moreover, location dummies average effects for all months and ignore the fact that trees only provide shade during certain months. Measurement error in the duration of treatment will cause a downward bias in shade effects.

Instead, I choose to parametrically model shade for each tree removal. This preferred approach reduces the direction-distance relationship into a single dimension and adds flexibility that allows me to identify which months a tree shades a house. The following section describes how I model shade using data on residence location, tree location, and sun position.

2.4.2 Shade Intensity Variables

Additional data are needed to measure the loss of shade associated with a tree removal. Tree shade can be determined from four metrics: tree location, tree size, tree height, and sun position. Two factors—tree size and location—are quantified by mapping canopy loss from aerial imagery. This section provides an overview of the movement of the sun. Then, I describe the data and methods used to simulate tree shade.

Sun position is defined by two angles: Azimuth angle (north, east, south, and west) determines shadow direction, and altitude angle (0 to 90 degrees) determines shadow length.[5] Sun movement describes temporal changes in sun position, which varies with time of day and by season. Short shadows are cast when the sun is at a high altitude angle, which reaches a daily maximum at solar noon and an annual maximum on the summer solstice. Conversely, the longest shadows are cast at sunrise and sunset and during winter.

Figure 3 diagrams seasonal sun paths over Gainesville and the temporal changes in sun position relative to a house.[6] During the spring and fall equinoxes, the sun rises

[5] From the object, azimuth angle measures sun direction (north, east, and south, west) around the circular plane; altitude angle is measured as the vertical degree angle from the object to sun.
[6] 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.

due east, sets due west, and arcs across the south, reaching a maximum altitude of 60 degrees (Figure 3, center arc). 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 maximum angle of 35 degrees (Figure 3, right arc). The summer solstice is a 14-hour day, with sunrise in the northeast, sunset in the northwest, and a midday sun passing overhead at a maximum altitude of 84 degrees (Figure 3, left arc).

2.4.3 Modeling Shade Loss from Tree Removals

The following section describes how I model shade based on tree canopy area and the movement of the sun. As discussed previously, tree removal permits are combined with aerial data on tree cover change and building location to determine a tree's position relative to a building. To isolate the effect of shade, I simulate the shadow cast by each tree using information on the movement of the sun in a given month and the tree crown area. I incorporate resulting data on sun movement to create a metric characterizing whether a tree provided shade for the residence before it was removed.

Shade is a function of sun position, tree height, and distance and angle between tree and home. Data on sun position are publicly available from sunearthtools.com for Gainesville, Florida in 2010. Sun path data include information on sun azimuth and sun elevation for each hour of the year or 8,700 observations of sun position (8,700 hours = 24 x 365). For each hour of daylight, I calculate the average azimuth and elevation angles in each month, creating approximately 288 sun position parameters (288 angles = 2 x 12 x 12).

Monthly shade zones around each house are calculated using a formula for shadow direction and shadow length. Shadow direction equals 360 minus shadow azimuth angle, reflecting the fact that shadows are cast in the opposite direction of the sun. Shadow length equals tree height divided by the tangent of sun elevation. I approximate tree height using tree growth function of common species combined with information on minimum trunk diameter of trees eligible for regulation. For estimating shadow length, I apply a tree height of 50-feet for all tree removals. The simulations create a shade zone around each house, whereby any tree within a distance of the defined shadow length for any hour-of-day is assigned as providing shade in that month.

Figure 4 illustrates the shadows of a 50-foot tree at various times of the day and during various seasons. In the summer, a tree on the west side of a house and set back as far as 75 feet provides cooling shade in late afternoon, leading to an expectation that such a tree may reduce summer electricity use from air-conditioning (Panel A). Interestingly, despite providing cooling midafternoon shade, a south-side tree may reduce electricity consumption only if it is less than 25 feet from the house. This follows because south-side trees cast very short shadows during the Gainesville summer, when the midafternoon sun is positioned overhead at a high angle. In the winter, on the other hand, south-side trees provide almost all shading (Panel B). In fact, because of the low altitude of the sun, a south-side tree set back as far as 100 feet can provide cooling shade that may increase a residence's heating demand. Since most Gainesville homes use natural gas during the winter, one could expect that south-side trees may increase natural gas used for heating.

Finally, I define a variable characterizing the size of canopy loss within a shade zone for each residence-month. This variable equals the sum of all canopy area positioned to shade a home at any hour-of-day for each month. A measure of shade canopy area provides a useful metric for comparing the effect of tree removal for shade trees of different sizes that provide shade to residences during different months-of-year.

In sum, shade simulation provides two distinct advantages for identifying shade-based treatment effects from tree removals. First, shade simulation provides a single metric containing all relevant information on the distance and angle between a house and home. Angle and distance are completely contained in shadow length and shadow direction calculations. Second, shade variables provide additional information about the timing of shade loss treatments, including the months-of-year and hours-of-day that effects from shade are expected to occur. Finally, in addition to reducing measurement error on shade effects, shade modeling provides a natural falsification test to demonstrate the absence of shade effects during months when trees never provided shade prior to removal.

2.5 Sample Construction

Previous sections describe the data and variables that will be used to estimate the effects of tree removal. This section describes the sample of residences used in this analysis.

Data on tree removals from the city of Gainesville include about 4,000 permits from commercial, government, and residential applicants. I drop observations with incomplete address information (19%), locations with multiple permits during the sample period (4%), and permits dates prior to March 2001 (<1%). For the remaining 3,090 permits, I use ArcGIS to geocode addresses.

Matching on geocoded addresses, I merge tree permit data with information on structural building characteristics. I use information on building style and year of construction to define the treatment sample. First, I restrict the sample to residences built before 2000, prior to the period of energy consumption. By limiting the sample to buildings that exist at the beginning of the study periods, I exclude trees removed as part of new developments.

I further restrict the sample to single-family residences with tree removals. Single-family homes account for approximately 57% of all tree permits in Gainesville. I only consider single-family homes for two reasons. First, tree ordinance regulations remain the same for single-family homes during the study time period. Single-family residences face a common selection process, regarding the size, species, and health status of trees eligible for regulation, and mitigation requirements. Second, single-family residences are likely to have similar energy demand from heating and cooling loads, and similar responses to tree removal treatments. Data include a total of 1,663 tree removal permits issued to single-family homes built before 2000.

Energy consumption records are merged with residences in the tree removal sample using address information with a match rate of 95%. The high rate of matching is possible using address geocoding algorithms available in ArcGIS that link standardized addresses to residence-specific identifier codes. The resulting sample includes 1,545 tree removals linked to a panel of monthly electricity and natural gas bills from 2000 to 2016. Finally, I trim the sample to reduce the effects of outliers. Specifically, I remove residences with a combined energy consumption above 15,000 kWh (99^{th} percentile) in any given month, and those with consumption below the 1,500 kWh (1^{st} percentile) in any given month.

The final sample is comprised of 1,383 residences with trees eligible for regulation and that received tree removal permits after March of 2001, when aerial imagery begins. I divide the sample into two groups: an early group that receives tree removal permits during 2001 to 2011 (early permits), and a late group that receives permits during 2012 to 2016 (late permits). The early group includes 834 residences, while the late group includes 574 residences. For the early permit group, I have information on tree canopy loss because these tree removals occurred between 2001 and 2011, a period during which corresponding aerial images are available. For the late group, by contrast, I do not have canopy loss information.[7] Since I am unable to distinguish between shade and non-shade trees, I drop posttreatment observations from the late permit group (2012 to 2016 permits). Nonetheless, I include pretreatment observations from the late treatment group to increase the sample of comparison residences.

2.6 Descriptive Statistics

The empirical strategy in this study is based on a comparison of energy bills before and after a change in tree cover (more on this below). Thus, the best comparison involves residences with adequate billing information prior to and after tree removal. As a result, for this analysis, treated residences were excluded if they had fewer than 12 monthly billing observations in either the pretreatment or the posttreatment period.[8]

[7] Unfortunately, it is impossible to characterize canopy loss information for residences with removals occurring after 2011.

[8] Results are robust to an alternative treatment definition that requires 24 monthly billing observations both before and following the permit issue date.

Even without this restriction, most treated residences have multiple years of pre- and posttreatment observations for each season. For example, the median treated household includes 8 years (90 months) of pre-treatment and post-treatment billing data. All treated residences have at least 4 years of post-treatment billing data, which guarantees at least one observation in each month of the year for both periods. The median comparison residence has 14 years (168 months) of pre-treatment billing data.

Table 1 reports basic summary statistics. Mean electricity consumption is 1,036 kilowatt-hours (kWh) per month. Mean natural gas consumption is approximately 28 therms per month. The average residence has 2,326 square feet of livable space, 3.1 bedrooms and 2 bathrooms. Most residences are single-story buildings on a 1/10th acre lot (1,409 square meters), and nearly all residences have central air-conditioning, forced-air heating systems, and shingled roofs. Most residences use natural gas for heating.

The average tree permit approves the removal of 1.5 trees. Early permits are associated with an average loss of 180 square meters of canopy, and an average loss of 66 square meters of shade canopy per month. On average, tree removals cause a loss of some shade for 6 months of the year, and a loss of heavy shade (over 100 square meters) for 2 to 3 months of the year.

3 IDENTIFICATION STRATEGY

This paper uses a difference-in-difference identification strategy for the sample of households that applied for tree removals after March 2001 and have customer accounts with GRU. I compare patterns in energy consumption among early tree removal residences and residences that applied for tree removal permits but had not had a tree removal by mid-2011, when remote sensing imagery ends.

I estimate the following main model:

$$ln(y_{imt}) = \beta TR_{imt} + \alpha_{im} + \alpha_{mt} + \varepsilon_{imt}, \quad (1)$$

where $ln(y_{imt})$ measures the natural log of energy consumption (natural gas, electricity, or a combined measure) at household i in month m and year t. The TR indicator variable switches from zero to one in the month after a household's tree removal is complete. The equation includes household-by-month-of-year fixed effects α_{im} to account for permanent differences in a household's energy consumption across months. It is possible to include such a rich set of fixed effects that account for household-specific seasonal variation in energy consumption because I observe households across multiple years. The model also includes month-by-year fixed effects α_{mt} to adjust for the average effects of time-varying factors (e.g., summer temperature) that generate variation in average consumption across all households. Standard errors are clustered at the residence level.

The parameter of interest is β, which measures the mean difference in energy consumption subsequent to the completion of tree removals, after adjustment for the fixed effects. It is a difference-in-differences estimator that compares the change in energy consumption after tree removal to before tree removal, relative to changes in consumption among households that have not yet removed a tree.

One potential concern in this empirical methodology is the presence of 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. Yet, in this study, the tree ordinance minimizes self-selection concerns because, for most residents, removing trees at will is not an option: the city arborist approves permits only for trees that are dying, diseased, or damaging property and denies permits for healthy trees that pose no safety hazards. Thus, residents do not control 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, Gainesville residents obey the tree removal ordinance with near-perfect compliance, since the regulations are well known and 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 (average cost of $250).[9] To further protect healthy trees, the ordinance also prohibits actions that degrade tree health, such as over-pruning or treetop removal.

Another potential threat to identification is that while the difference-in-difference estimator accounts for any time-invariant characteristics that affect energy consumption, there may be time-varying factors that affect household demand for energy and also influence the timing of tree removal. For example, households might choose to remove a

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

tree in declining health due to a drought or freeze, in which case weather factors may affect the trajectory of energy consumption and the timing of tree removal. Alternatively, if households are more likely to remove trees following a period of *low* energy consumption, then the estimated tree removal effect could reflect a reversion to mean energy consumption. I test the identifying assumption of common trends in Figure 5. Figure 5 compares the seasonal patterns in pre-treatment energy consumption for these two groups among residences with permits issued after 2004. The top panel shows that electricity consumption occurs primarily during hotter summer months (May-October), when demand for air-conditioning is highest. The bottom panel shows that natural gas consumption occurs primarily during colder winter months (November-April), when demand for heating is highest. The two groups of residences have similar energy consumption trends during pre-treatment years of 2001 to 2004. In Figure 8, I augment specification (1) by adding leads and lags. The results indicate that during the one, two, and three years prior to tree removal, households have no change in energy consumption.

4 RESULTS

4.1 The Effect of Tree Removal on Energy Consumption

If tree removal has the expected effect, summertime electricity consumption should increase and wintertime natural gas consumption should decrease because of the additional exposure to the sun. Table 2 reports the estimates of specification (1) for electricity, natural gas, and a combined measure of energy. Columns 1, 3 and 5 report estimates in natural logs; columns 2, 4, and 6 report estimates in levels.

Overall, the results suggest that tree removals have large, economically important effects on electricity consumption. Column 1 shows that a tree removal increases electricity consumption by 5.5 percentage points. A substantial effect remains when electricity is measured in levels, rather than logs. Column 2 shows an increase of 36 kilowatt-hours per month, or 4 percentage points for the median residence. [10] The large tree removal effect for electricity likely reflects the fact that shade reduces summer air-conditioning demand, which means that, prior to tree removal, residences had lower electricity demand.

Columns 3 and 4 show that, by contrast, tree removals do not increase natural gas consumption. In fact, point estimates are negative for natural gas, though not statistically significant. A negative tree removal effect for natural gas is consistent with the hypothesis that shade increases winter heating demand. Column 3 shows a decrease of only 0.007 therms per month, a change equivalent to less 0.2 kilowatt hours per month. The size of the natural gas point estimate is small compared to that of the electricity estimate (0.2 compared to 34). Note however that the level estimates are less precise than log estimates, and natural gas estimates are less precise than electricity estimates. Nonetheless, the results in Table 2 are consistent with the hypothesis that tree shade has a greater effect on summer electricity demand than on winter natural gas demand.

Columns 5 and 6 show effects reported in terms of combined energy consumption quantified as kilowatt-hours per month. The size of tree removal effects for combined energy consumption is about the same as effect size for electricity, with regard to both percentage effects (4.7 compared to 5.5) and level effects (36 compared to 34). The

[10] Percentage conversion based on the median consumption of 856 kWh per month.

implication is that electricity effects account for the lion's share of overall energy effects from tree removals.[11]

Because electricity consumption is strongly (positively) correlated with summer air-conditioning demand, the large effect on electricity is consistent with the hypothesis that shade loss from tree removals increases air-conditioning demand. This motivates estimating tree removal effects separately during summer months when air-conditioning demand is greatest.

4.2 Monthly Effects of Tree Removal on Energy Consumption

Next, I estimate the increased electricity consumption from tree removals over month-of-year. In this section I augment specification (1) by interacting the after-treatment variable with month-of-year dummies. I model electricity and natural gas separately to establish whether decomposing estimates by fuel type does, in fact, imply differences in seasonal effects. Figure 6 illustrates the results graphically for electricity (top panel) and natural gas (bottom panel).

Figure 6 shows that tree removal effects are greatest during warmer summer months, when air-conditioning demand is highest. Panel A shows that for electricity, the effects are greater during summer months than winter months. In colder winter months, point estimates for electricity are indistinguishable from zero. Similarly, Panel B shows that natural gas has no effect during any months, indicating that tree removals have no effect on heating demand. In hotter summer months (May to October), by contrast, point estimates for electricity range from 45 to 60 kilowatt-hours, with high levels of statistical

[11] This is in contrast to many traditional energy efficiency measures, such as building codes and attic insulation, which have larger effects for natural gas than electricity.

significance. On average, point electricity estimates during May to October are larger than the estimate in Table 2 (60 compared to 34 kilowatt-hours). The larger summer effect for electricity likely reflects the fact that, prior to tree removal, shade reduces air-conditioning demand. Table 2 estimates of tree removal effects on air-conditioning suffer from downward measurement error bias since effects are averaged across all months of the year. The estimates in Figure 6 are therefore consistent with the hypothesis that tree removals increase electricity savings most during hotter summer months when demand for air-conditioning is greatest.

4.3 The Effects of Shade and Tree Size on Energy Consumption

Although the models presented in Table 2 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 because tree removal is an imperfect proxy for change in tree shade, which contains measurement error. For example, the removal of trees that did not shade houses would not affect energy use. Such measurement error in the treatment biases estimates toward zero and underestimates the true effect of shade on energy use. The magnitude of this bias will depend on whether the average tree that is removed was positioned to shade a home. To address these concerns and achieve more precise estimates of the energy savings from tree shade, I next account for heterogeneity in shade changes across tree removals.

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 it was removed. To isolate the effect of shade, I simulate the shadow cast by each tree based on the movement of the sun in a given month and the tree crown area. I estimate the main regression model, additionally controlling for an interaction term between treatment and shade.

I further estimate a heterogeneity in the treatment effect by the size of tree removal, discretizing shade into shade area bins, where $A \in \{0, 25, \ldots, 275, 300\}$ and each area bin represents one 25 square-meter interval of canopy loss affecting shade.[12]

Binning coefficients provides a flexible treatment of the effect of canopy area loss in shade zones on energy consumption. The top panel of Figure 7 shows results that account for variation across the size of tree canopy loss, and variation over the month-of-year that a tree provides direct shade to buildings.

Figure 7 shows that a larger amount of shade loss appears to increase treatment effects.[13] Although the sign and trend of shade intensity coefficients are predictable, the magnitude of effects is unexpected. Surprisingly, energy increases to as much as 20 percentage points for months with large shade loss (above 200 square meters), compared to months that tree removals do not affect shade. The point estimates for large shade loss represents the potential role of specific shade trees in regulating residential energy consumption, rather than the average role of trees. By contrast, for months affected by moderate shade loss (100 to 200 square meters) energy increases by about 10 percentage points. The top panel shows that shade effects are highly statistically significant for all shade loss categories above 100 square meters.

12 For example, the first bin A_0 is equal to 1 if $\{0m^2 \leq ShadeArea_{im} < m^2\}$ and equals zero otherwise, where $ShadeArea_{im}$ corresponds to canopy area shading house i in month-of-year m prior to removal.
[13] I use the combined energy consumption measure here because it yields an overall effect that accounts for potential substitution between electricity and natural gas.

The estimates for treatments without shade loss provide a falsification test, which can address concerns about potential selection in the timing of tree-removing residents. For months with little or no shade loss (0 to 100 square meters) residences have energy effects that are statistically indistinguishable from zero. These coefficients suggest that small trees (less than 100 square meters) provide shade of little economic value for energy conservation. In fact, the estimate that includes observations without any shade change (0 to 25 meters) is negative, though not significant. This provides a falsification test, since prior to tree removal, household energy demand was not being reduced by additional shade. This is important because most trees only provide direct building shade for part of the year, which implies that treatment effects vary with time based on sun movement. In sum, results in Figure 7 (top panel) provide strong evidence that shade is the causal mechanism through which tree removals affect energy consumption.

The results presented above suggest that tree size matters when located in shade zones, where treatment intensity varies by month-of-year and only considers canopy area within shade zones. One may still be concerned that residences selecting large tree removals may have different energy consumption trajectories than residences with small tree removals. To address this concern, I test whether the size of tree removal matters, regardless of tree location. Specifically, I augment specification (1) interacting the treatment indicator with the total canopy loss. If shade does indeed drive results, then including all tree removals introduces measurement error and causes a downward bias of estimates for tree size effects.

Figure 7 (bottom panel) shows that tree size (overall canopy) has little effect on energy consumption effects of tree removal. Compared to shade area results (top panel) the total canopy area results (bottom panel) have smaller coefficients and less precision. In fact, among the 13 area bins, 11 coefficients are statistically indistinguishable from zero. The two larger bins that are statistically significant are likely correlated with increasing area of tree canopy in shade zones. Figure 7 results are consistent with the hypothesis that estimates including non-shade related tree removals bias treatment effects downwards.

Figure 7 provide causal evidence that residences increase energy consumption after tree removals associated with shade loss. Furthermore, shade effects increase with the area of tree canopy loss positioned to shade buildings. The implication is that the energy benefits from trees are a function of the timing and intensity of shade to buildings.

4.4 Inferring Causality from the Timing of Tree Removal

Finally, in Figure 8, I examine how the effect of tree and shade removal may vary over time. The model augments the main specification by adding leads and lags. Results show an absence of preexisting trends, and also provide evidence that tree removal effects begin at the time of tree removal and persist over time. Estimates in the top panel indicate that in the year following tree removal, residences use 4.25 percent more energy overall, an effect that persists in subsequent years. The treatment effect for year four forward increases to about 5.25 percent, an effect that is statistically significant at the 10 percent level.

The bottom panel shows a similar pattern in the timing of tree removal effects for large shade treatments; however, estimates have a larger magnitude and higher precision. In one, two, and three years following tree removal, residences with a large amount of shade loss (over 100 square meters) increase energy consumption by 10 to 12 percent, with estimates significant at the 1 percent level. Treatment effects continue to persist beyond three years, likely because replacing shade loss requires growth times longer than the study period.

5 VALUATION OF SHADE BENEFITS

This section evaluates the private and social returns to investments in residential tree shade. To conduct this part of the analysis, I use the quasi-experimental estimates of the effect of tree removals on energy consumption to estimate the average energy savings from extending the life of existing shade trees.

5.1 Private Returns of Residential Tree Shade Protection

To express the estimates of monthly energy savings in dollar terms, I use the measure of electricity savings in percentage terms. Estimates of average monthly electricity savings are multiplied by the product of average monthly electricity in the control group and the residential retail price of electricity ($0.15/kWh). The estimate of average energy savings from a tree preservation is approximately $103 per year. The estimate of average energy savings increases to $110 per year for a large shade tree, and $166 per year for very large shade trees.[14]

Next, I evaluate the net present value of energy savings from extending the useful life of a tree. I invoke several assumptions to compute the discounted value of energy savings. First, I rely on assumptions about tree lifespans. Based on inspections by a city arborist, the average tree in this study has an expected remaining lifespan of 15 years or

[14] Large shade tree estimates use monthly percentage estimates for trees with 100-200 m^2 shading canopy multiplied by the average number of months that at least 100 m^2 of canopy are in the range of shading a tree (7 months). Similarly, very large shade tree estimates use percentage estimates for 200-300 m^2 shading canopy loss based on 7 months of shade.

less. In my calculations, I consider discounted benefits assuming extended lifespans of 10 and 20 years. I also assume that the effect of retrofits on energy consumption – and real energy prices – do not vary over the increased lifespan.[15] Discounted benefits are calculated using annual discount rates of 3%, 6%, and 10%.

Table 3 shows the net present value of energy savings from extending the useful life of tree. The discounted benefits vary by tree size and location. Estimates of present value of the savings for average trees range from $388 to $938 depending on the time horizon and discount rates. The present value of shade trees ranges from $682 to $1,651 for large shade trees, and $1,023 to $2,476 for very large shade trees.

Compared with advantages of other energy efficiency investments, the value of benefits from shade trees is substantial. Three recent studies apply GRU billing data in order to study returns from home retrofits and building codes using difference-in-difference methods. One study reports that the net present value of energy savings for 9 retrofit technologies ranges from $13 to $6,500, assuming a 6% discount rate (Maher 2016). [16] By comparison, the private benefits of trees are comparable to the average building retrofit. Specifically, the average tree, large shade tree, and very large shade tree, outperform 2, 5 and 6 out of the 9 retrofit measures, respectively.[17] Two studies consider the impact of Florida energy codes in Gainesville and estimate annual energy savings of

[15] Based on GRU's block-rate pricing in 2011 and the distribution of monthly consumption levels observed in the data, the average marginal price that GRU charges consumers for electricity is 15.3 cents per kilowatt-hour. This consists of a 10.2-cent average energy charge plus a 5.1-cent fuel adjustment charge. The average marginal price that GRU charges consumers for natural gas is 76.4 cents per mmbtu (1 therm = 10 mmbtu). This consists of a 68.9-cent average energy charge plus a 3.5-cent fuel adjustment charge.

[16] The nine retrofits include air-conditioner maintenance ($13), low-income grants ($128), refrigerator removal ($1,043), high-efficiency pool pump ($1,264), room air-conditioner ($1,309), duct leakage repair ($1,474), attic insulation ($3,428), high-efficiency central air conditioners ($4,934), and super-seer air conditioners.

[17] Comparisons assume a 20-year time horizon and 6% discount rate present value estimates for average tree ($723), large tree ($1,273), and very large tree ($1,909).

$106 per year – strikingly similar to the annual savings of $103 for the average tree (Jacobsen & Kotchen 2013, Kotchen 2016).

The energy savings from shade capture only one dimension of the welfare benefits of trees. In addition to reducing air-conditioning demand, shade may also increase indoor and outdoor thermal comfort. Moreover, homeowners may choose to invest in trees for other benefits, or choose to remove them due to costs imposed by trees. Assessing the socially optimal level of tree cover is beyond the scope of this study. Nonetheless, the results of this study quantify one value of residential tree cover – the value of shade.

5.2 Social Returns of Residential Tree Shade Protection

Tree investment decisions also generate avoided energy expenditures for neighboring households when shade crosses property boundaries. Other widely recognized externalities of energy-efficiency investments include health benefits from cleaner air and avoided costs of capital investment in energy infrastructure. This section considers these energy-related externalities from private trees.

Panel A of Table 4 evaluates the internal rate of return (IRR) on investment from the social perspective. More precisely, I report the discount rate at which the discounted value of the average avoided energy expenditures exactly equals the average upfront costs to the city. I impose several additional assumptions. First, I assume the benefits of tree ordinance come from delaying tree removals for some households, effectively extending the useful life of mature trees by either 10 or 20 years.

Column 1 of Table 4 computes the internal rate of return using average costs to the city[18] and the average reduction in annual energy savings from encouraging 100 households to defer mature tree removals. The measure of energy savings is derived from annual avoided energy expenditures derived from the quasi-experiment, plus an additional 30% to adjust for shade spillovers to neighboring properties. When an adjustment for externalities is added to the energy savings, the annual benefits are roughly $130 per year. Using this measure of annual benefits, the IRR is -17.8% for the 10 year horizon and -4.2% for 20 years. This finding of highly negative returns suggests that reduced energy expenditures alone may not justify the public costs of tree ordinance.

Column 2 assumes the tree ordinance encourages 200 households to defer mature tree removals each year. Over the 15-year sample period, this implies about 3,000 residences (5% of single-family homes) delay tree removals to comply with the tree ordinance. This implies that for each tree permit issued, another tree removal is delayed by at least 10 years. Using this policy scenario, the rate of return is -7.8% for the 10 year horizon and 2.2% for 20 years. The finding of low, or negative, returns remains.

Panel B repeats calculations adding the value of avoided emissions to benefits. Carbon emissions are valued at $35 per ton. Nitrogen oxide, sulfur dioxide, and particulate matter ($PM_{2.5}$) emissions from electricity consumption are valued using regional emissions and ground-level damage factors (Holland et al., 2016).[19] Accounting

[18] I apply a value of $265,421 as the annual management costs associated with the Gainesville tree ordinance in a 2015 budget provided by the Gainesville City Arborist. This includes salaries of four full time employees and operational costs. Employee salaries include a City Arborist ($67,195), an Urban Forestry Inspector ($76,818), an Operation Supervisor ($58,224), and a Staff Assistant ($47,202). Operational costs include overhead ($10,065), professional training ($2,867), gas and oil ($1,150), and vehicle fleet mileage ($1,900).

[19] I calculate total marginal damages of $0.04978 per kWh from all sources of emissions based on a marginal damage estimate of $0.03010 per kWh for local pollutants ($SO_2 + NO_X + PM_{25}$) and CO_2 damages of $0.01960 per kWh. Holland et al. (2016) estimate local marginal damages per kWh ($SO_2 + NO_X + PM_{25}$)

for benefits of avoided emissions, IRRs increase modestly for the 10 and 20 year time horizons to -3.4% and 5.3%, respectively.

The social IRR additionally includes the value of avoided costs of generation and capacity.[20] The demand-side management perspective is an especially important one from which to judge investments given Gainesville's priority of delaying infrastructure upgrades to meet a growing energy demand. In fact, since 2005, GRU has invested in a wide set of policies to encourage residential energy efficiency investments. I extend the value of avoided infrastructure costs to monetize the benefits of demand reduction from tree ordinance policies.[21]

Panel C adds the value of deferred or avoided infrastructure investments to benefits from avoided emissions and household expenditures. This adjustment leads to a meaningful increase in the IRR relative to the IRRs in Panels A and B. In particular, social IRR for the 10- and 20-year time horizons are 1.5% and 8.9%, respectively. Column 1 shows social IRRs at 10 and 20 years of -9.9% and 0.7%. Overall, city tree ordinance investments can have positive rates of returns based on energy-saving benefits alone for reasonable policy scenarios and time horizons.

Two assumptions suggest that IRRs from Table 4 represent conservative estimates of the social rate of returns from tree ordinance policies. First, cost calculations consider the full costs of program administration, including staff, overhead, and capital costs.

for 24 hours for each of nine district regions set by the North American Electric Reliability Corporation (NERC). I use the average hourly marginal damage estimate of $0.03010 per kWh for the Florida Reliability Coordinating Council (FRCC) region that includes Alachua County, Florida.

[20] Avoided costs are the expenses GRU would have incurred had it generated or purchased electricity.

[21] I use the value of $0.0643 kWh as the marginal avoided costs of generation and capacity, discounted at a rate of 6%. This represents the average energy benefits during off-peak demand. Using off-peak benefits generates conservative estimates of avoided infrastructure benefits since energy savings from trees likely correspond to summer peak electricity demand. Since 2005, this figure has justified $27.8 million in energy conservation initiatives (GRU 2005).

Since residential tree removals only account for about 50% of all tree permits related to these resources, I may overestimate the costs associated with tree ordinance enforcement of residential permits. Second, I assume conservative scenarios of the impact of tree ordinances on tree cover. For example, I ignore future shade benefits associated with tree planting requirements. Together, these assumptions provide an upper-bound of program costs and lower-bound of ordinance-induced energy savings, suggesting estimates represent an upper-bound of the cost per kilowatt-hour of electricity saved.

5.3 Cost Effectiveness of Tree Ordinance for Climate and Energy Policy

Panel D considers the returns from the tree ordinance policy using an estimate of the cost per ton of CO_2 avoided. This is calculated as the ratio of the net cost of tree ordinance enforcement (the annual rental cost of upfront expenditures less the social value of annual energy savings) and the tons of CO_2 reduced per year.[22] Using a 6% discount rate, Column 1 shows a cost per ton of CO_2 of $725 and $319, respectively – roughly an order of magnitude greater than the social cost of carbon used by the United States federal government. By contrast, Column 2 shows a cost per ton of CO_2 of $149 and -$57, respectively. Although estimates range widely, any cost below $35 per ton of CO_2 reduced could be considered an efficient approach to mitigating climate change. Indeed, a cost of -$57 would imply that tree ordinance investments simultaneously reduce local social costs and carbon emissions.

Panel E conducts a similar exercise to create an estimate of the cost-effectiveness of the tree ordinance from a demand-side management perspective, after accounting for

[22] Subtracted social value includes all avoided energy expenditures, avoided local pollutants, and avoided infrastructure benefits. CO_2 includes avoided emissions as described previously, plus 200 pounds per year of CO_2 sequestered by a mature tree.

other social benefits of energy savings. These are calculations of costs per kilowatt-hour saved based on all administration costs and estimated energy savings based on extended tree lifespans discounted at a rate of 6%. This is calculated as the ratio of the net cost of tree ordinance enforcement discounted at a rate of 6% and the kilowatt-hours of electricity demand reduced per year.[23] The 20-year horizon calculations of cost per kWh avoided of $0.01 and $0.23 suggest that tree ordinances may be cost-effective relative to other demand-side management programs run by GRU (Maher 2016).

6 CONCLUSION

As public agencies search for effective policies to reduce energy consumption, there is hope for a solution arising from nature itself. So far policymakers have been able only to speculate about the value of trees. Florida's hot summers and high-air conditioning demand would suggest that it may be reasonable to expect high returns from tree shade. Taking a quasi-experimental approach to estimate accurate causal impacts, this study establishes that shade trees are, in fact, effective in reducing energy consumption. Moreover, the magnitude of effects is substantial.

Results suggest that a typical tree removal leads to 4-6% more electricity consumption annually. Treatment effects are largest during summer months (8-10%), and months with large amounts of shade loss (8-20%). Falsification tests show zero effects for natural gas consumption, winter energy consumption, and no-shade loss tree removals. Taken together, these findings provide robust causal evidence that tree shade reduces residential air-conditioning demand.

[23] Subtracted social value includes all avoided energy expenditures and avoided local and CO_2 emissions.

Private benefits from shade trees are economically significant. Depending on tree size and location, the private benefits of shade trees range from $450 to $1,900, and are comparable to energy-efficiency investments such as pool pumps, air conditioners, and refrigerators. Further, this study compiles significant evidence that can help homeowners locate trees strategically to lower summer electricity bills. Such recommendations could become an important element of consumer education to reduce future energy use. This paper underscores that it is critical to develop a body of credible evidence on the energy benefits of green infrastructure that can be directly compared with other energy efficiency investments.

From a policy perspective, the Gainesville tree ordinance appears to pass a cost-benefit test based on energy savings alone. From a utilities perspective, the paper finds that tree ordinances may prove to be cost-effective demand-side management policies. In terms of climate policy, it is unclear whether the Gainesville tree ordinance can provide the least expensive carbon reductions. However, the dual role of residential trees in terms of adaptation to climate change may make tree protection policies more attractive investments for cities. More broadly, these results inform utilities and local governments about the social value of trees, allowing them to design policies to target energy saving benefits.

APPENDIX 1: CHANGE CLASSIFICATION METHODS[24]

Overview of Canopy Change Detection

Constructing aerial images of tree loss poses two challenges. First, urban regions like Gainesville are composed of tightly interwoven land uses with low-density tree cover. This makes it difficult to use satellite sensors, like Landsat, which are better for measuring annual tree loss in landscapes with higher-density tree cover. 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 periodically commissioned to photograph cities at an ultrafine spatial resolution as detailed as 1 foot by 1 foot. This aerial imagery dramatically increases the spatial resolution and accuracy of classifications. However, low-flying aircraft also identify tree loss over a 10-year interval instead of the annual time interval possible with Landsat.

To generate the data for this study, two sets of aerial imagery are used as inputs. Images from 2001 characterize initial, pretreatment tree cover at a spatial resolution of 1 foot by 1 foot and include spectral signals for three bands, collectively known as the visible bands because they can be seen with the human eye: blue (459–479 nm), green (545–565 nm), and red (620–670 nm). Images from 2011 characterize final, posttreatment tree cover; they also have a spatial resolution of 1 foot by 1 foot and include the same three visible bands plus a fourth band of near-infrared signals (841–876

[24] 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.

nm). Images were acquired separately, with the 2001 data purchased from ACPA, and the 2011 data provided by Florida State University. Both data sets were collected during summer, with trees in leaf. (Attempts to enhance 2001 images with an available band of 1999 near-infrared imagery were 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.

Unlike satellites, which collect data over time with the same sensor, aircraft may collect data with different sensors, and these instrument changes introduce spectral differences in images across time. Spectral differences can also occur when images are collected from different altitudes, at different lens angles, or at different times of day. 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 data sets are therefore aggregated at 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 detection (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 are 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 thus 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, shrubs, and gardens—which remain green year-round in Florida. Several methods for change detection can exploit multidimensional data, including classification and regression tree models; however, such objective methods are most effective for comparing 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

To take full 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 change detection model determines 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 (Figure 1).

Technical Details of Canopy Change Detection Methods

Land use change classification occurs in two stages: change detection and land use 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 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 categories of bare soil 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 difficult because of spectral differences between the 1999 and 2011 NIR sensor and differences in the season and time of day of imagery. A fusion of 1999 NIR and 2001 red imagery created similar problems. A more objective classification and regression tree (CART) model also produced unreliable results because of spectral differences between 2001 and 2011 imagery. Although CART methods are typically preferred, they are most effective when two sets of data come 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, unsupervised (ISODATA) classification, and feature extraction methods permit greater flexibility in defining the criteria for land-use changes.

REFERENCES

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

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): 207–38.

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: 662–68.

ERDAS, Inc. 1999. Erdas IMAGINE Field Guide, fifth ed. Atlanta, Georgia. Available at http://www.gis.usu.edu/manuals/labbook/erdas/manuals/FieldGuide.pdf. Accessed May 4, 2015.

Fowlie, M., M. Greenstone, C. Wolfram. 2015. Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program? Revision requested by *Quarterly Journal of Economics*.

Gainesville Regional Utilities (GRU). 2013. Planning and Building a Strong Gainesville: Gainesville Regional Utilities Annual Report 2012-2013. Available at: https://www.gru.com/Portals/0/Legacy/Pdf/AnnualReports/2012-2013GRUAnnualReport.pdf

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

Holland, S., E. Mansur, N. Muller, and A. Yates. 2016. Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors, *American Economic Review*. Forthcoming.

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.

Kotchen, M. 2016. Longer-Run Evidence on Whether Building Energy Codes Reduce Residential Energy Consumption, *Journal of the Association of Environmental and Resource Economists*, forthcoming.

Maher, J. 2016. Measuring the Accuracy of Engineering Models in Predicting Energy Savings from Home Retrofits: Evidence from Monthly Billing Data. Draft paper, University of Maryland, College Park, August.

Pandit, R., and D. Laband. 2010. Energy Savings from Tree Shade. *Ecological Economics* 69: 1324–29.

Tou, J.T., and R.C. Gonzalez. 1974. Pattern Recognition Principles. Reading, MA: Addison-Wesley Publishing Co.

Treiman, T., and J. Gartner. 2004. Community Forestry in Missouri, U.S.: Attitudes and Knowledge of Local Officials. *Journal of Arboriculture* 30(4): 243–50.

U.S. 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.

U.S. Green Building Council. 2011. Green building research. Accessed on March 2, 2011 at http://www.usgbc.org/DisplayPage.aspx?CMSPageID=1718.

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): 103–14.

FIGURES

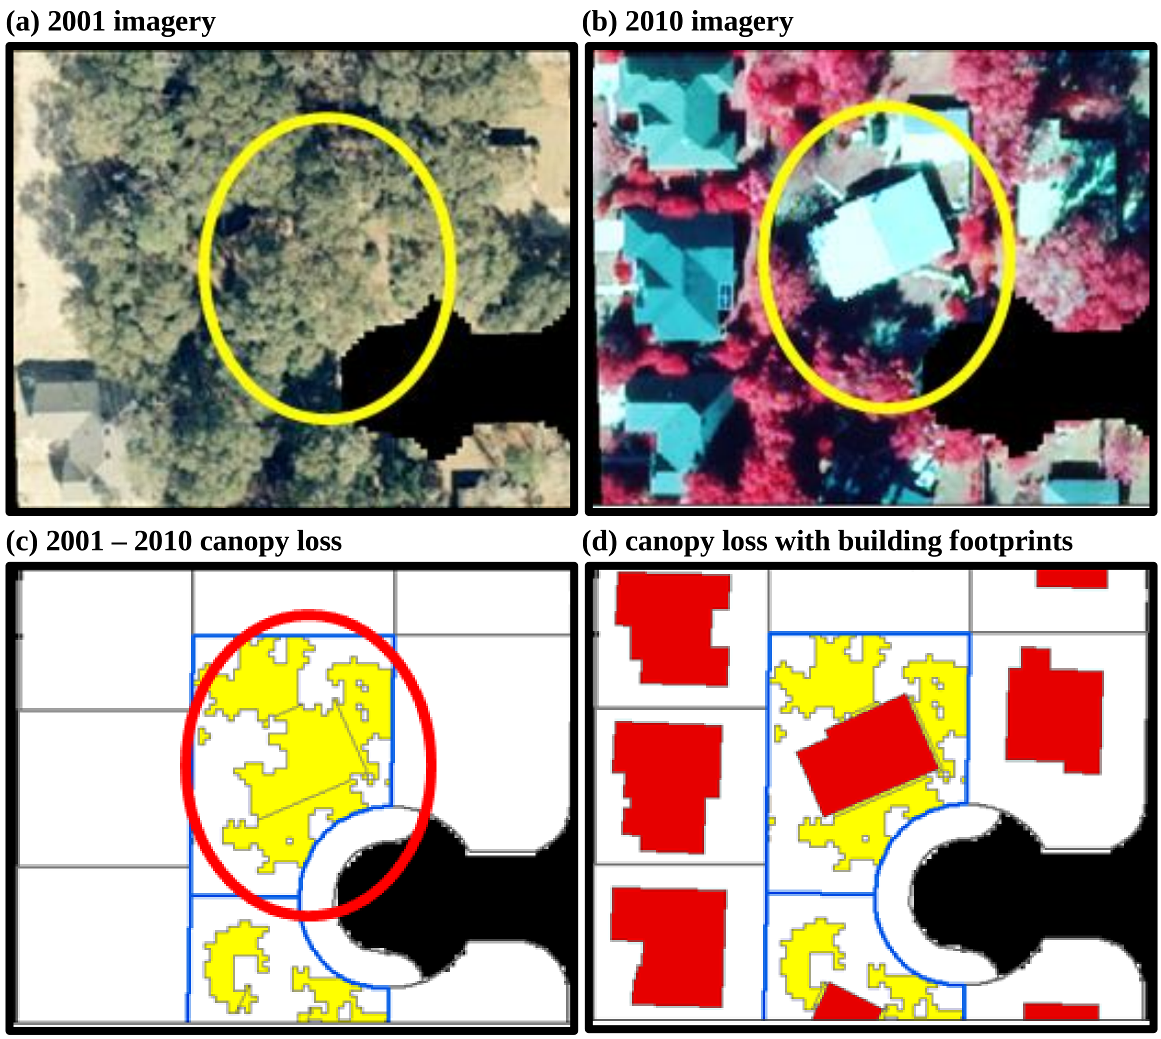

Figure 1. Inputs and outputs used to create canopy loss map, with auxiliary cadastral data used to define shade-intensity variables. Panels (a)–(d) cover identical spatial extent in a Gainesville neighborhood to illustrate the remote sensing data inputs. Panel (a) shows 2001 color imagery input (1-foot resolution) and panel (b) shows 2010 color infrared imagery input (1-foot resolution). Panel (c) shows the tree canopy loss map (1-meter resolution) for two properties with tree removal permits (blue borders). Yellow polygons represent areas of canopy loss during 2001–2010, blue borders represent property boundaries for residents with tree removals, and gray borders 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.

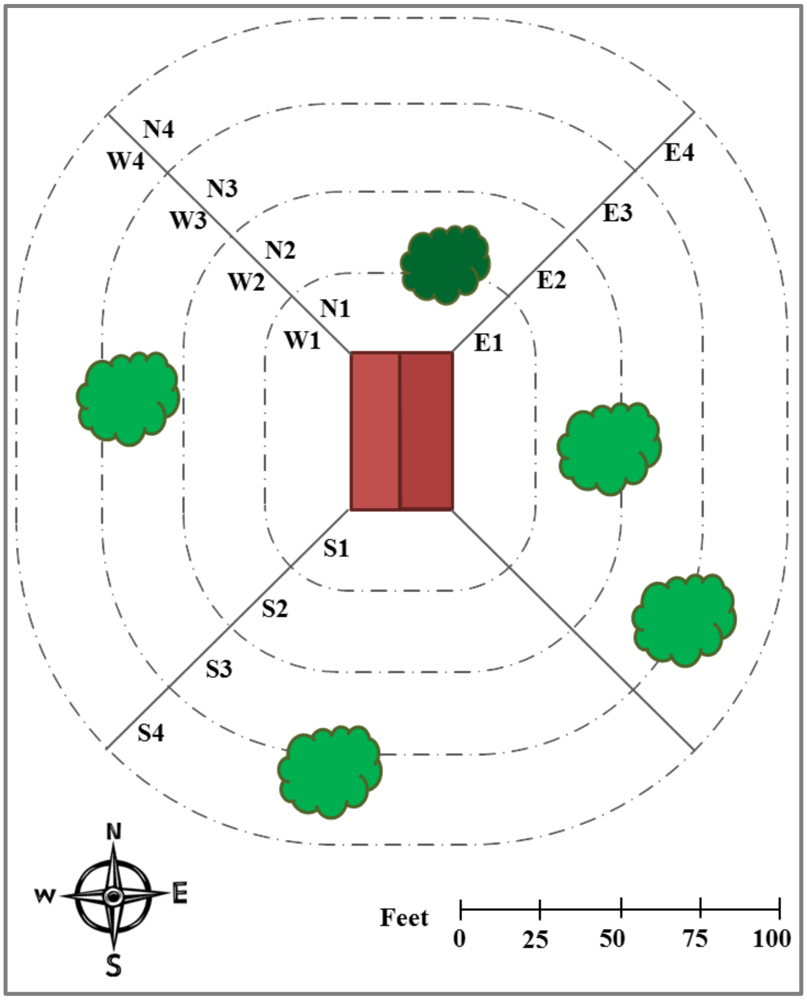

Figure 2. Illustration of 16 zones of tree position for a house defined by four directional quadrants (north, east, south, west) and four distance buffers (25, 50, 75, and 100 feet).

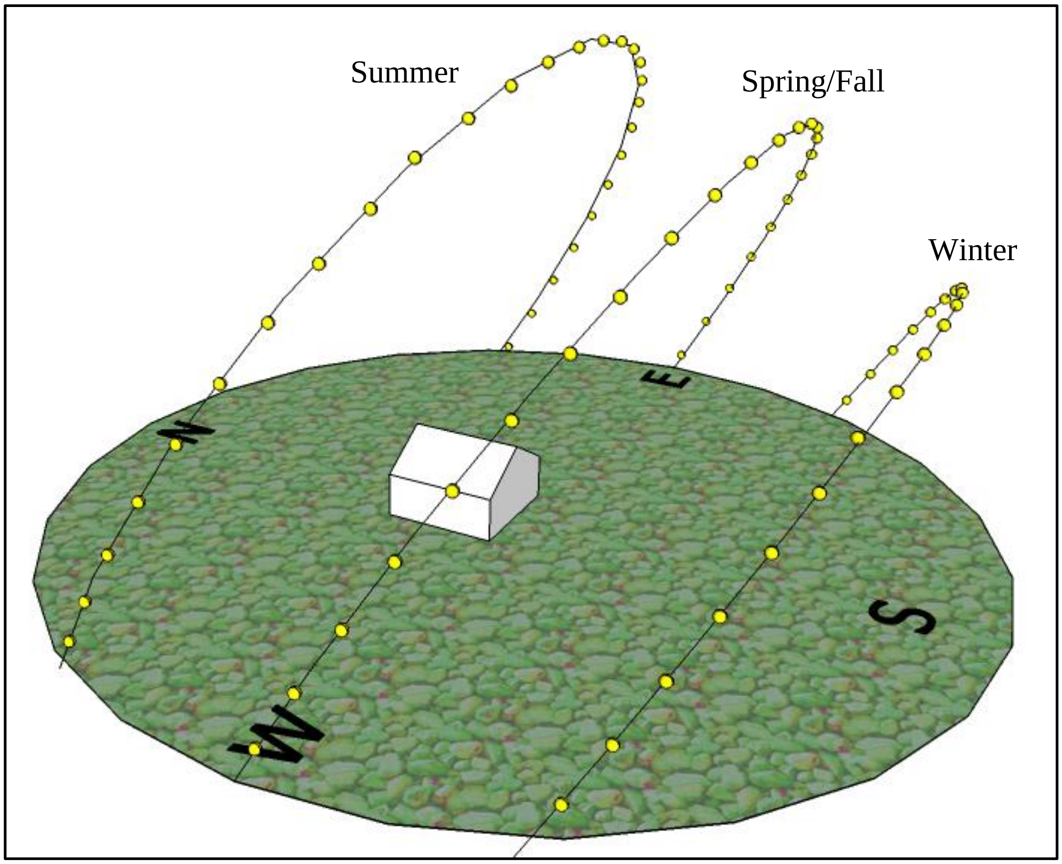

Figure 3. Sun path on the summer solstice (June 21), spring and fall equinoxes (March 21 and September 21), and the winter solstice (December 21). Each yellow sphere marks the position of the sun in 30-minute increments between sunrise and sunset. From the house, azimuth angle measures sun direction (north, east, and south, west) around the circular plane; altitude angle is measured as the degree angle from house foundation to sun.

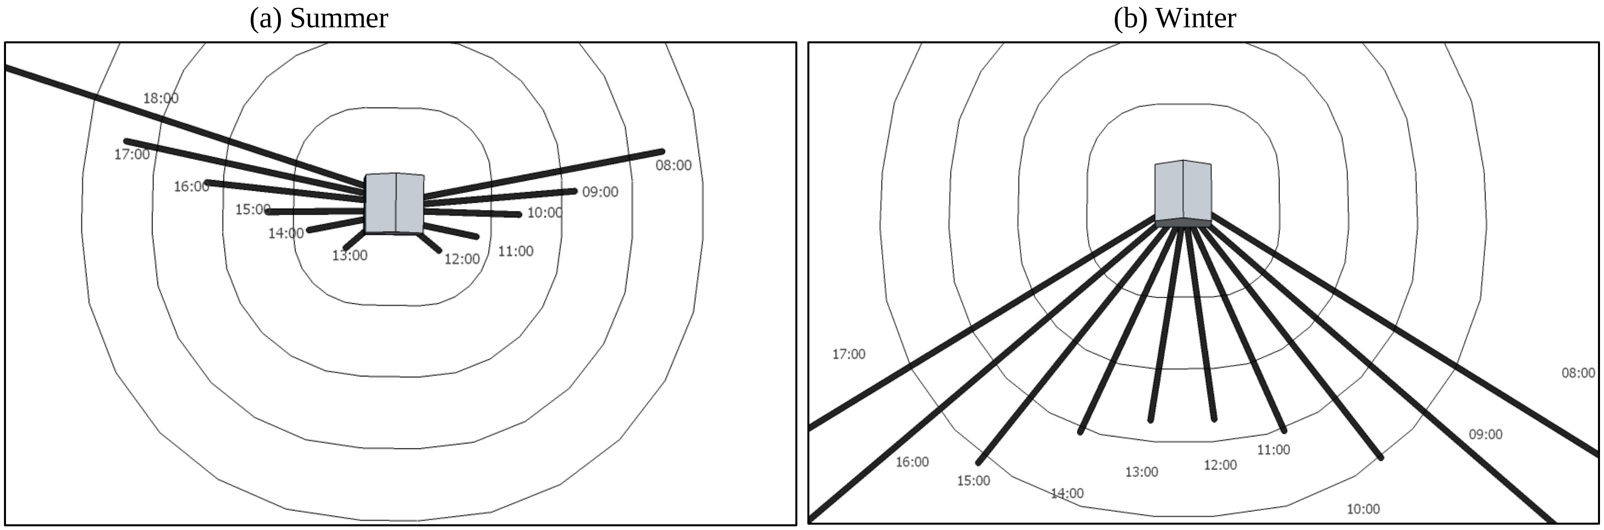

Figure 4. Hourly shadows cast by a 50-foot tree located at the maximum distance to shade the 15-foot roof of a two-story house in Gainesville, Florida. Each black line represents the length and direction of a shadow on the hour from 8 a.m. to 6 p.m. Concentric circles represent areas within 25-, 50-, 75-, and 100-foot buffers surrounding a square house. Panel (a) shows the shadows on June 21 (summer solstice). Panel (b) shows the shadows on December 21 (winter solstice).

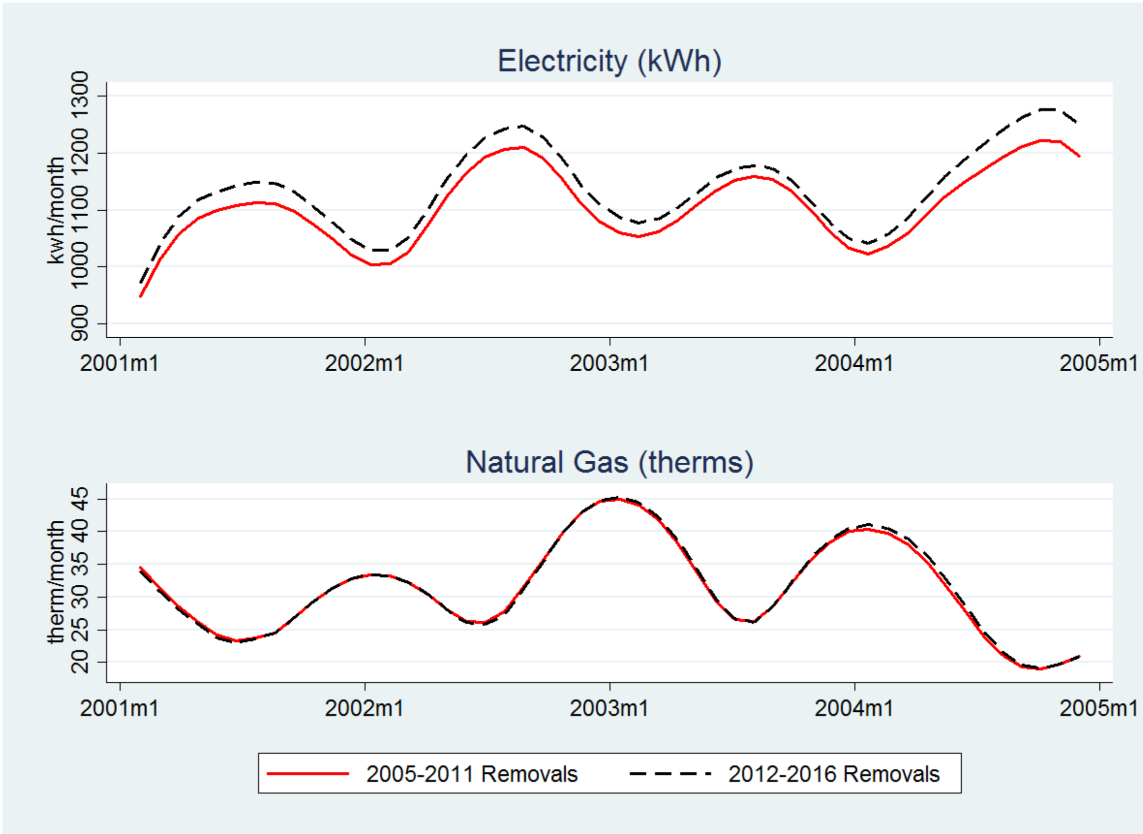

Figure 5. Comparison of pretreatment energy trends for treated residences (2005 – 2011 permits) and control residences (2012 – 2016 permits). Each line represents a separate kernel-weighted local polynomial regression on pretreatment electricity and natural gas from January 2001 (2001m1) to January 2005 (2005m1). Data include the sample of all residences with both electricity and natural gas consumption, including 934 residences (444 treated residences and 490 control residences) and 41,322 monthly energy observations.

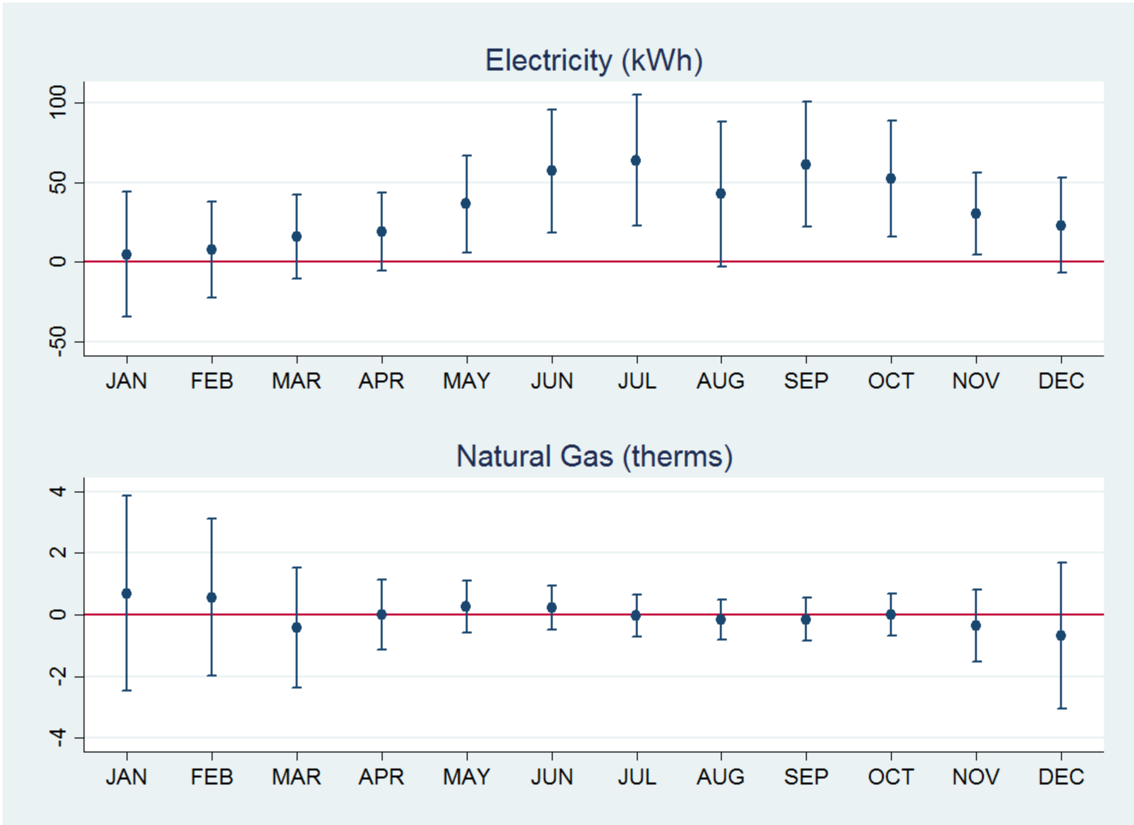

Figure 6. Monthly effects of tree removal treatment. Dependent variable is monthly energy consumption in either kilowatt-hours (electricity) or therms (natural gas). The treatment-by-month interaction terms switch to one after tree removal for one month-of-the-year. Each regression includes residence-by-month-of-year dummies and month-by-year dummies. Data include 1,383 residences with tree removal treatments and 220,459 monthly energy observations. Standard errors are clustered by residence. Bars represent 95-percent confidence intervals.

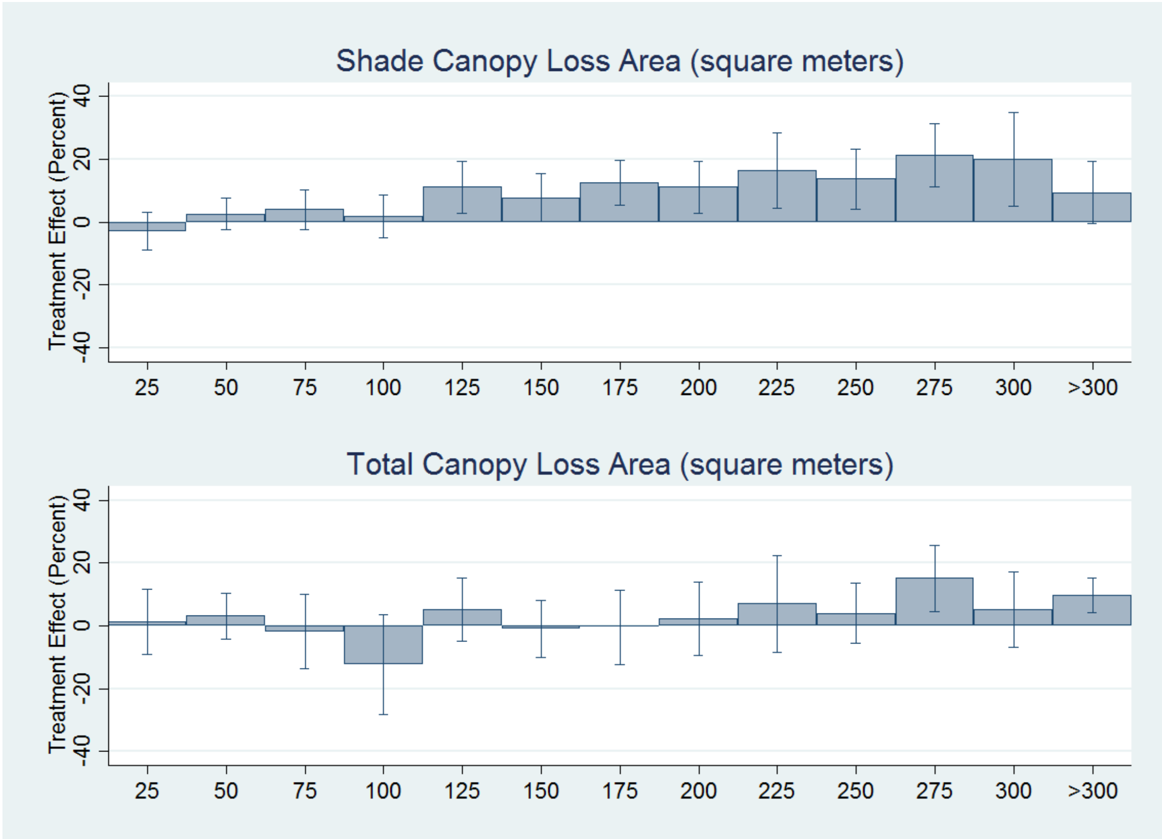

Figure 7. Estimated effects of tree removal by shade intensity (top panel) and tree canopy area (bottom panel). Treatment effects are estimated as a percent change in overall energy consumption. Top panel "shade intensity" bins represent increments of $25m^2$ of canopy area positioned to shade a residence in month m. Bottom panel "total canopy area" bins represent $25m^2$ increments for total canopy loss area associated with removal. Each regression includes residence-by-month-of-year dummies and month-by-year dummies. Data include 1,383 residences with tree removal treatments (834 residents with canopy area classified) and 220,459 monthly energy observations. Standard errors are clustered by residence. Bars represent 95-percent confidence intervals.

Figure 8. Estimated effects of tree removal for years before, during, and after removal event. Treatment effects are estimated as a percent change in overall energy consumption. Top panel treatment indicator switches to one in the month after tree removal. Bottom panel "shade loss" dummy equals one in months with over 100 square meters of shade canopy loss after tree removal (heavy shade treatment). Each regression includes residence-by-month-of-year dummies and month-by-year dummies. Data include 1,383 residences with tree removal treatments and 220,459 monthly energy observations. Standard errors are clustered by residence. Bars represent 95-percent confidence intervals.

TABLES

Table 1. Summary Statistics

Variables	Mean	Standard Deviation	Minimum	Maximum
Electricity (kWh)	1,036	680	1	10,311
Natural Gas (therms)	28.4	35.5	0	473
Effective Year Built	1979.7	7.7	1937	1999
Living Space (ft^2)	2,326	912	696	6,548
Heated Space (ft^2)	1,809	688	604	5,313
Bathrooms	2.0	0.6	1	5
Bedrooms	3.1	0.6	1	5
Lot Size (m^2)	1,409	1,079	186	19,843
Central Air Conditioner	0.96	0.20	0	1
Forced Air Heating	0.97	0.17	0	1
Shingled Roof	0.82	0.38	0	1
Trees Removed	1.5	2.4	0	45
Year of Tree Permit	2009.8	4.2	2001	2016
Month of Tree Permit	6.0	3.2	1	12
Total Canopy Loss (m^2)	179	306	0	3,575
Shade Canopy Loss (m^2)	62	138	0	1,810
Shade Loss (> 0 m^2)	0.48	0.50	0	1
Shade Loss (> 100 m^2)	0.20	0.40	0	1

Notes: Summary statistics are based on 220,459 observations.

Table 2. Difference-in-differences estimates of the effect of tree removal on monthly energy consumption, 2000 - 2016

Variables	(1)	(2)	(3)	(4)	(5)	(6)
	Electricity		Natural gas		Combined energy	
	log(kWh)	kWh	log(therms)	therms	log(kWh)	kWh
Tree removal effect	0.0546***	34.35**	-0.034	-0.0074	0.0468***	35.76*
	(0.018)	(14.05)	(0.022)	(0.53)	(0.018)	(20.30)
R-squared	0.562	0.673	0.777	0.777	0.599	0.686
Observations	219,118	219,118	158,515	168,909	220,459	220,459
Residences	1,383	1,383	1,110	1,110	1,383	1,383

Notes: Dependent variable is monthly consumption of electricity (1, 2), natural gas (3, 4), or a combined measure (5, 6). The “tree removal” dummy equals switches to one in the month after tree removal. Each regression includes residence-by-month-of-year dummies and month-by-year dummies. Data include 1,383 residences with tree removal treatments and 220,459 monthly energy observations. Column 3 has fewer observations than Column 4 because observations are dropped when preforming a log transformation of zero values. Standard errors are clustered by residence. Standard errors are reported in parentheses.

Table 3. Present value of (discounted) savings

Time Horizon		Discount Rate	
	3 percent	6 percent	10 percent
Panel A: Average tree in sample			
10 years	$538	$464	$388
20 years	$938	$723	$537
Panel B: Large shade tree			
10 years	$946	$817	$682
20 years	$1,651	$1,273	$945
Panel C: Very large shade tree			
10 years	$1,420	$1,225	$1,023
20 years	$2,476	$1,909	$1,417

Notes: Panel A reports the net present value of electricity savings implied by the preferred estimate reported in Column 2 of Table 2. Panel B reports the net present value of energy savings implied by the estimates for 100-200 square meter canopy area shade trees reported in Figure 2 based on the sample average of 7 months of shade per year. Similarly, Panel C reports the net present value of energy savings implied by the estimates for 200-300 square meter canopy area shade trees reported in Figure 2 based on the sample average of 7 months of shade per year. Reductions in electricity bills associated with the estimates are assumed to accrue over the remaining life of the tree using a range of discount rates and assumed time horizons.

Table 4. Estimated social returns on city investments in tree ordinance

Time Horizon	Deferred Tree Removals Per Year	
	100 trees	200 trees
Panel A : Internal rate of return		
10 years	-17.1%	-7.8%
20 years	-4.2%	2.2%
Panel B: Internal rate of return, adding avoided emissions damages		
10 years	-13.6%	-3.4%
20 years	-1.8%	5.3%
Panel C: Social internal rate of return		
10 years	-9.9%	1.5%
20 years	0.7%	8.9%
Panel D: CO_2 abatement cost - 6 percent discount ($/ton CO_2)		
10 years	$725	$149
20 years	$319	-$57
Panel E: Cost per kWh avoided energy consumption – 6 percent discount		
10 years	$047	$0.13
20 years	$0.23	$0.01

Note: All calculations use a cost of tree ordinance enforcement of $265,421. This implies a cost $2,654 or $1,327 per tree removal deferred by enforcement of the tree ordinance, depending upon policy scenario. Column 1 assumes that ever year the tree ordinance encourages 100 residents to defer a tree removal. Column 2 assumes 200 residents defer tree removals each year. All Panels reflect the annual electricity savings estimated in Column 2 of Table 2 plus an additional 30% savings to account for savings due to shade on neighboring buildings (536 kWh per year). Panels B and C also incorporate the value of estimated emissions reductions for CO_2, SO_2, NOx, and $PM_{2.5}$. Panel C additionally incorporates GRU's value of avoided or deferred infrastructure investments. Reductions in electricity consumption associated with the estimates are assumed to accrue over the extended life of the tree assuming time horizons of 10 and 20 years.