Title

Deforestation in Paraguay and the Impacts of the Zero Deforestation Law

Author

Fenton, Marieke

Publication Date

2024

Peer reviewed|Thesis/dissertation

Deforestation in Paraguay and the Impacts of the Zero Deforestation Law

By

MARIEKE CHRISTINA FENTON

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Agricultural and Resource Economics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Rachael Goodhue, Chair

Kevin Novan

Brittney Goodrich

Committee in Charge

2024

Abstract

On December 13, 2004, the Zero Deforestation Law made deforestation illegal in Eastern Paraguay. Deforestation did not stop, however the law may have decreased deforestation. Most deforestation in this region is conducted to clear land for agriculture. In this dissertation, I conduct three analyses to investigate patterns of deforestation in Eastern Paraguay and the impact of this law on aggregate deforestation, predictors of deforestation, and post-deforestation agricultural land uses.

The first analysis investigates drivers of deforestation in Eastern Paraguay before the Zero Deforestation Law. I use satellite-derived deforestation data in a linear probability model to investigate how physical and anthropological land characteristics are correlated with forest loss between 2001 and 2004. I find that physical land characteristics including slope, elevation, soil group, and ecoregion were useful for predicting deforestation. Only some anthropological land characteristics predicted deforestation. Distances to towns or roads were not correlated with deforestation. Proximity to patchy forest cover, such as in areas of existing agricultural clearings, was correlated with more deforestation.

The second analysis investigates whether deforestation fell after the Zero Deforestation Law was implemented, and whether the relationship between land characteristics and deforestation changed. I use a linear probability model on a panel dataset of deforestation and land characteristics. My results show that deforestation fell by around a quarter after implementation. In addition, the predictive power of some drivers of deforestation weakened. One relationship reversed. Pre-policy, deforestation was more likely with less tree cover nearby, while post-policy, deforestation was more common near more densely forested areas, possibly indicating a desire to hide. My results cannot be explained by other events that took place during this time, and they provide strong evidence that the Zero Deforestation Law successfully lowered deforestation in Eastern Paraguay.

The third analysis investigates how these decreases in deforestation post-policy were distributed across small-scale subsistence-oriented agriculture, large-scale commodity-oriented agriculture, and large-scale cattle ranches. Data does not exist on post-deforestation agricultural land use, so I generate my own data using a three-step process. First, I manually identify the post-deforestation land use for randomly sampled deforested locations in Eastern Paraguay. Second, I use these observations to train random forest models that predict post-deforestation land use from physical and anthropological land characteristics. Third, I use these models to generate land use predictions for all pixels deforested in Eastern Paraguay within four years before or after the Zero Deforestation Law came into effect. I find that the decrease in deforestation after implementation can mostly be attributed to a decrease in deforestation for large-scale agriculture. Clearing for large-scale agriculture fell by over 60%, from 34 thousand hectares annually on average before the policy to 12 thousand hectares annually on average after the policy. This was due to a composition effect, under which different locations were deforested after the policy was implemented than before, and a land use effect, under which the same deforested locations were used for different agricultural classes in the pre- and post-policy periods. Levels of deforestation for rangeland and small-scale agriculture remained relatively steady between the pre- and post-policy periods.

Acknowledgements

I am extremely grateful to my dissertation committee for encouragement and guidance. I learned so much from each of you about how to do research, and also how to prioritize people and maintain balance. Rachael, your considerate mentoring style is aspirational. Kevin, I admire your ability to stay grounded and curious through all steps of the research process. Brittney, thank you for being supportive through the dissertation process, and for teaching me about bees and bringing me on California specialty crop field work for our additional research.

I am grateful for the supportive and welcoming environment of the ARE department. The people made the program so much more than an academic experience, and made my time here immensely enjoyable. While there are far too many to name, a few follow. Charlotte, Dasha, Julian, Sarah, and TJ, thank you for being great study buddies early on and great friends throughout. Olivia, thank you for being a dependable office mate from day one. Laura and Isy, thank you for being there for me in the latter years. I am grateful to Peace Corps Paraguay for hosting me for two years prior to this work. I am grateful to Joey for supporting me through all the ups and downs that come with pursuing a PhD, and to my parents and siblings in North Carolina for supporting me from far away and making it so enjoyable to go back. I could not have done this without each of your support.

Table of Contents

List of Figures

List of Tables

I. Introduction

Deforestation and agricultural expansion have proceeded hand in hand in Eastern Paraguay. National policy encouraged farmers to clear land in the interior of the country for crops and livestock since the mid-20th century. This benefited a national economy that relies heavily on agriculture. However, this agricultural expansion came with tradeoffs. Forests provide valuable local and global ecosystem services that are lost or diminished when the forest is cleared. These ecosystem services include carbon sequestration, maintenance of local and regional climates, biodiversity habitat, water filtration, and the provisioning of food and timber resources (Taye et al., 2021).

Recognizing the potential harms of excessive deforestation and that deforestation was proceeding at an unsustainable rate, Paraguay implemented the Zero Deforestation Law in December 2004. This law made deforestation in Eastern Paraguay illegal. However, the ban did not stop deforestation. This is not unexpected. Legal bans commonly fail to fully eliminate their targeted offense. For example, this is often true of drug use bans. Regardless, the passage and implementation of the Zero Deforestation Law was important because it marked the transition of deforestation from a legal to an illegal activity in Eastern Paraguay.

In this dissertation I investigate three questions after providing background information and a discussion of data in section II. First, I investigate which land characteristics were correlated with deforestation in Eastern Paraguay before the implementation of the Zero Deforestation Law, when deforestation was legal. A rich literature exists linking deforestation with characteristics that make deforestation more likely to occur in a location. Generally, land that is physically

1

better suited for agriculture, or land that is nearer to populations or infrastructure, is more likely to be deforested. I use satellite-derived data on deforestation from Global Forest Watch (Hansen et al., 2013) in a linear probability model to investigate how physical and anthropological land characteristics were correlated with forest loss in Eastern Paraguay from 2001 through 2004. I find that before the Zero Deforestation Law was implemented, physical land characteristics including slope, elevation, soil group, and ecoregion were useful for predicting where deforestation will occur. Only some measures relating to human activity were useful for predicting deforestation. Distances to towns or roads were not correlated with deforestation. Proximity to patchy forest cover, such as in areas of existing agricultural clearings, was correlated with more deforestation. My results underline the importance of physical suitability for agriculture in predicting where deforestation will occur. Accessibility was less important, contrary to findings in prior literature. This may be because Eastern Paraguay was already extensively developed and deforested by 2001. Deforestation did not occur along a remote frontier, but rather patchily and situated amongst areas of existing development. Understanding whether these commonly identified patterns held in Eastern Paraguay pre-policy sets a baseline for further study of the impacts of the Zero Deforestation Law.

In the second analysis, I investigate how the Zero Deforestation Law impacted aggregate deforestation and how it altered the drivers of deforestation. My analysis of deforestation data from Global Forest Watch shows that while the Zero Deforestation Law did not stop deforestation in Eastern Paraguay, deforestation slowed after the law was implemented. Deforestation fell from 93 thousand hectares per year on average between 2001 and 2004 to 71 thousand hectares per year on average between 2005 and 2008, a decrease of nearly 25%. Identifying this drop in the rate of deforestation is informative but leaves the story incomplete. I argue that this decrease in deforestation can be attributed to the Zero Deforestation Law by considering and dismissing alternative explanations for a drop in deforestation that coincide with the law. I then analyze whether the change in aggregate deforestation between the pre- and post-policy periods is significant, and whether the relationship between commonly identified drivers of deforestation changed following the law. To do this, I run a linear probability model with deforestation as the outcome on a panel dataset of forested locations before and after the Zero Deforestation Law was implemented. I find that the decrease in deforestation post-policy remains statistically significant after controlling for characteristics of forested land. I also find that, in addition to lowering deforestation, the ban shifted how some land characteristics examined in the previous section correlated with deforestation post-policy. Many drivers became less useful in predicting where deforestation will occur. For example, while land that was more suitable for agriculture in terms of slope and elevation remained more likely to be deforested, this trend was not as strong as it was before the ban. This change has important policy implications. Deforestation became more dispersed across physical characteristics in the post-policy period. If Paraguay wishes to pursue a further decrease in deforestation, the increased uncertainty in where deforestation is most likely to occur will make this decrease more difficult to achieve.

Finally, I investigate how changes in deforestation after implementation of the Zero Deforestation Law were distributed across the diverse types of agriculture in Eastern Paraguay. The agricultural groups that I consider are small-scale subsistence-oriented farms, large-scale commodity-oriented farms, and large-scale cattle ranches. Understanding the distribution of impacts is important because policies to slow deforestation have additional and sometimes unanticipated environmental, social, and economic consequences. For example, a law that shifts the use of newly deforested land from small-scale to large-scale agriculture could have negative distributional impacts since access to newly cleared land decreased for the more economically vulnerable group, as well as potentially more adverse environmental impacts since small-scale systems are typically more biodiverse.

The distribution of policy impacts across agricultural groups is difficult to analyze because there is no dataset summarizing what each deforested plot is subsequently used for. I present a method to do so. Specifically, I manually identify the post-deforestation agricultural use for a sample of deforested locations using satellite imagery and knowledge of agricultural patterns in the regions. These data points are used to train random forest models to predict the post-deforestation land use of all land deforested in Eastern Paraguay in the four years directly before and after the Zero Deforestation Law was implemented. These predictions are based on physical land characteristics, such as slope and soil type, and anthropological land characteristics, such as proximity to population centers and roads. Data on these physical and anthropological characteristics are available globally, meaning that this method can be replicated for other data-scarce regions, offering a contribution to the distribution literature.

When I disaggregate deforested areas by their predicted post-deforestation agricultural use a clear pattern emerges. Deforestation for large-scale agriculture was most strongly impacted by the Zero Deforestation Law. Before the Zero Deforestation Law was implemented, an average of 34 thousand hectares of forest were cleared per year for large-scale agriculture. Afterwards, this fell to an average of 12 thousand hectares per year. In contrast, deforestation for small-scale agriculture or rangeland continued at similar levels for the entire period. The shift away from clearing for large-scale agricultural use is because different locations are cleared, referred to as the composition effect, and because deforested locations are being put to a different use after the ban, referred to as the land use effect.

The remainder of the dissertation is organized as follows. Section II provides background on deforestation in Eastern Paraguay and describes the data used in this dissertation. Section III investigates the drivers of deforestation in Eastern Paraguay, and how they compare to drivers commonly identified in the literature. Section IV looks at the impact of the Zero Deforestation Law on aggregate deforestation and how relationships between land characteristics identified as drivers of deforestation and deforestation change after the law is implemented. Section V analyzes how post-deforestation land use changed after the Zero Deforestation Law was implemented by describing the manual generation of land use data, the random forest models used to predict agricultural land use, and revealed land use patterns both before and after the policy is enacted. Section V concludes.

Deforestation in Eastern Paraguay

Paraguay is a landlocked country at the heart of South America. Eastern Paraguay, the region of interest in this analysis, is separated from Western Paraguay by the Paraguay River and bordered by Argentina and Brazil. Most deforestation is driven by agriculture, therefore an understanding of the agricultural landscape of Eastern Paraguay is useful.

Land in Paraguay's interior was primarily under state control until the War of the Triple Alliance in the late 1800s, after which some areas passed to large-scale private ownership (Nickson, 1981). Widespread development did not take off until the 1950s when government policy began encouraging settlement and clearing of idle and underutilized land for agriculture (Richards, 2011; Blanc, 2015). The development of road networks and the 1957 establishment of the city Ciudad del Este on Paraguay's eastern border encouraged expansion. A bridge soon connected Ciudad del Este to the adjacent Brazilian port of Paranagua, facilitating the export of agricultural commodities (Richards 2011). Legislation also drew people to the interior. The Agrarian Statute of 1963 granted land titles for settlements. This attracted small-scale farmers eastward from the areas surrounding Paraguay's capital city Asuncion to the agricultural frontier in the interior of the region. The statute also led large landowners to preemptively deforest to avoid land grabs (Nagel 1999). An amendment to the Agrarian Statue in 1967 allowed foreigners to purchase property within 150 kilometers of the border, leading to an influx of Brazilian farmers along the eastern border. The Itaipu hydroelectric project upriver from Ciudad del Este, begun in 1975 and completed in 1982, brought thousands of laborers to the region, many of whom established agricultural settlements.

Satellite images, available since 1972 (Finer et al., 2018), document the later years of this transformation from forest to agriculture. Huang et al. (2007) found that in the early 1970s three quarters of the Atlantic Forest region, which extends halfway across Eastern Paraguay from the border with Brazil, was forested. By 2000, only one quarter of the Atlantic Forest region remained under forest cover. This represents a loss of over four million hectares of forest, an area approximately the size of Denmark, in just over 25 years.

The agricultural groups that settled and deforested the region over the last century still dominate the landscape today. The main crops are strongly correlated with farm size (Weisskoff, 1992). Eastern Paraguay's most expansive crop is soy, which is produced primarily on large-scale systems. In 2010, soy was cultivated on over 2.5 million hectares in Eastern Paraguay and covered more than three times the area of the second highest area crop, maize. Mechanized soy took off after being introduced by Brazilian settlers and expanded further with the illegal introduction of genetically modified soy in 1996 (Correia, 2019). A lack of land, capital and/or credit have prevented small-scale farmers from adopting mechanized soy production (Peters, 2015). Ownership patterns in soy production persist, with a substantial portion of large-scale soy operations under Brazilian ownership (Galeano, 2012). Large-scale cattle ranching also plays an important role in the Paraguayan economy. Over half of Paraguay's 12 million head of cattle were produced in Eastern Paraguay in 2010 (Ministerio de Agricultura y Ganaderia, 2010). On the other side of the spectrum, small-scale farmers use labor-intensive methods to produce crops largely for personal consumption. Small-scale crops with the largest area include mandioca, beans, corn, and peanuts, as well as small quantities of livestock for family and local consumption. Cotton was an important market crop for small-scale farmers until the mid-1990s after which it declined due to a dismal global market (Richards, 2011).

II. Data

The three analyses in this dissertation use the same data, which I have compiled from a variety of sources to create a single coordinated spatial dataset. I use the 30-meter pixel from the Global Forest Watch data as my level of observation. All additional data sources are resampled to match this resolution. Different subsets of the data are used in the different analyses, as noted in their respective methodology sections.

Data on Forest Cover

I use a 30 percent forest cover cutoff to define forested area. The cutoff matches the official definition of forest used by Paraguay (UN FAO, 2020), and is supported by measures of percent tree cover in forested areas of Eastern Paraguay (Huang et al., 2009). In 2000, the baseline year for this analysis, Eastern Paraguay had around 6.5 million hectares of forest.

Deforestation is provided as annual pixel-level indictors of loss, beginning in 2001 (Hansen et al., 2013). Figure 1 plots cumulative deforestation in Eastern Paraguay between 2001 and 2008. In the four years before the Zero Deforestation Law came into effect, 93 thousand hectares of forest were lost on average annually in Eastern Paraguay. In the four years after implementation, between 2005 and 2008, an average of 71 thousand hectares of forest cover were lost per year. The locations of deforestation in the four years before and after the Zero Deforestation Law was implemented is mapped in Figure 1.

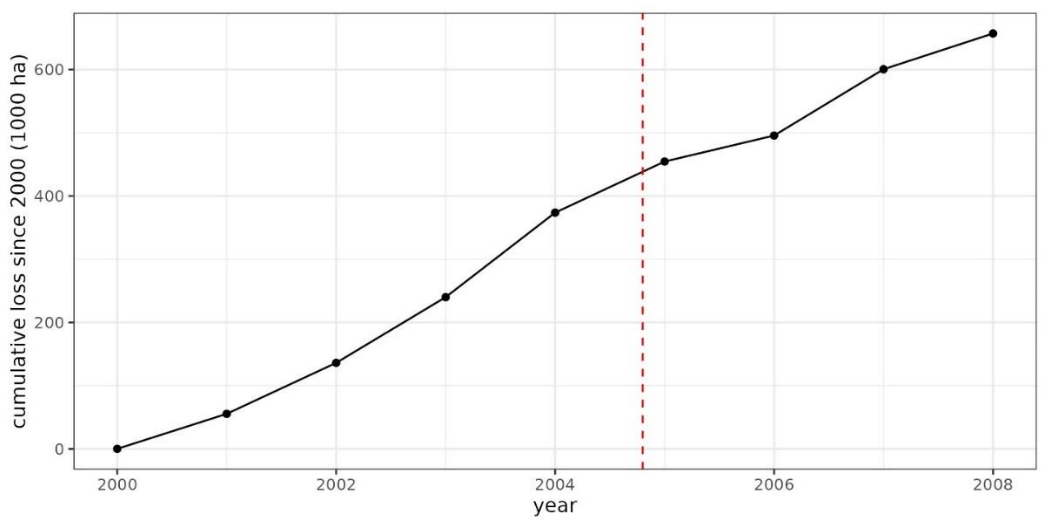

Figure 1: Cumulative deforestation in Eastern Paraguay from 2001 through 2008, based on Global Forest Watch tree cover and deforestation data (Hansen et al., 2013). The dotted line indicates the implementation of the Zero Deforestation Law at the end of 2004.

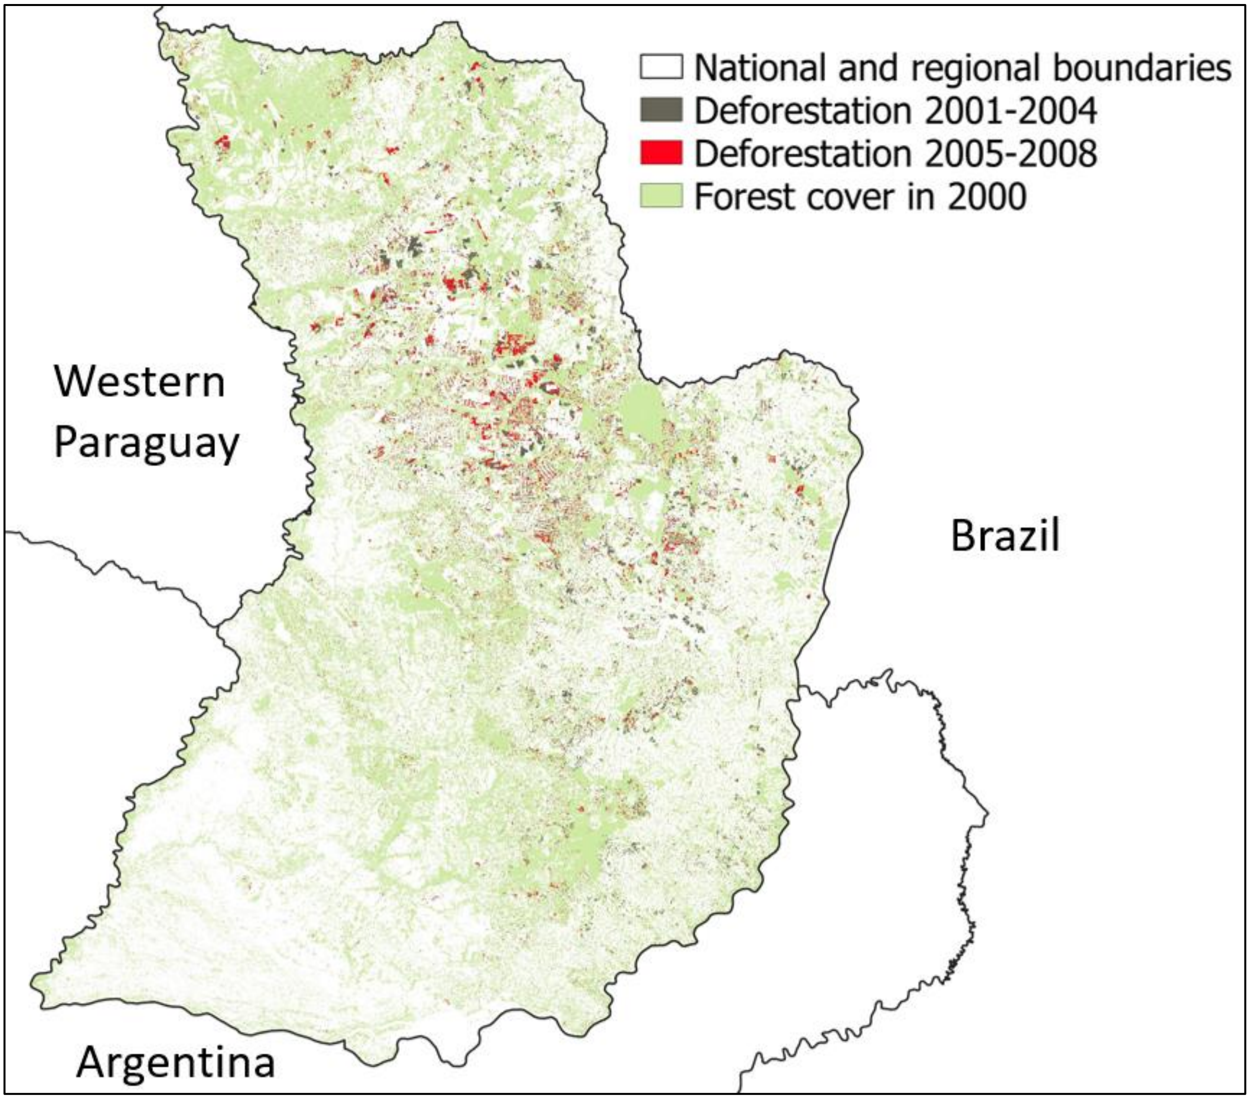

Figure 2: A map of baseline forest cover and deforestation in Eastern Paraguay in the four years preceding and following the implementation of the Zero Deforestation Law.

Two features of the Global Forest Watch data are particularly important. First, tree cover is defined as all vegetation over five meters in height. The data does not differentiate between natural forest and tree plantations or agroforestry. This is unlikely to affect the analysis in a substantive way because tree plantations and agroforestry play a small role in the agricultural economy of Paraguay. Eucalyptus, the most common tree crop, was cultivated on only around 55 thousand hectares in 2008. Yerba mate, a traditional tree product of the region, was cultivated on less than half this area. In comparison, soy was cultivated on nearly 2.5 million hectares in 2008 (Ministerio de Agricultura y Ganaderia, 2009). An alternative deforestation dataset, Tropical Moist Forests (Vancutsem et al., 2021), omits agricultural tree cover. Deforestation in the two datasets is similar, as can be seen in a comparison of the datasets in appendix A1. Second, Global Forest Watch does not allow the same pixel to be lost multiple times. If tree cover were lost multiple times, the first loss since 2000 is recorded. This is again unlikely to substantively impact results as cumulative reforestation data from Global Forest Watch shows minimal planting or regrowth in Paraguay.

Data on Land Characteristics

I utilize data on physical and anthropological land characteristics. These characteristics have been identified in the literature as drivers of deforestation, or factors that make deforestation more likely to occur in a location (Busch and Ferretti-Gallon, 2017). In my analyses, these land characteristics are used two ways. First, I use them to explore the relationships between land characteristics and deforestation in the context of Eastern Paraguay. Second, I use them to predict the post-deforestation agricultural land use of deforested pixels in Eastern Paraguay.

Physical characteristics used in my analysis include slope, elevation, soil group, and ecoregion. Slope and elevation come from GLAD (Potapov et al., 2020). These measure the topological characteristics of the land. Elevation is measured in meters, and slope is measured in degrees. Maps of elevation and slope can be seen in Figure 3 and Figure 4.

Soil group provides information on potential agricultural productivity, and comes from the International Soil Reference and Information Centre (ISRIC) Soil and Terrain Database for Latin America and the Caribbean (ISRIC, 2005). Soil groups vary across Eastern Paraguay's geography and elevation, including tropical yet productive nitisols in the southeast, wetland gleysols in the west, and weakly developed regosols in the north, among others. A map of Paraguay's soil groups can be seen in Figure 5. Descriptions of soil groups can be found in appendix A2 (Driessen et al, 2001).

Ecoregions data is from terrestrial ecoregions of the world (TEOW), which classifies areas based on the dispersion of communities and species prior to major land use change and are intended for use in conservation (Olson et al. 2001). Most land area in Eastern Paraguay falls into one of three ecoregions: the Humid Chaco, the Atlantic Forest, and the Cerrado. The Humid Chaco, the westernmost ecoregion, is a mosaic of woodland and savannah where frequent flooding turns low elevation areas into bogs. The Atlantic Forest in the east is naturally forested from river plains to mid-level plateaus. Pockets of Cerrado, an ecoregion whose natural vegetation ranges from open fields to dense forest, are found on the northern plateaus. A map of the ecoregions can be seen in Figure 6.

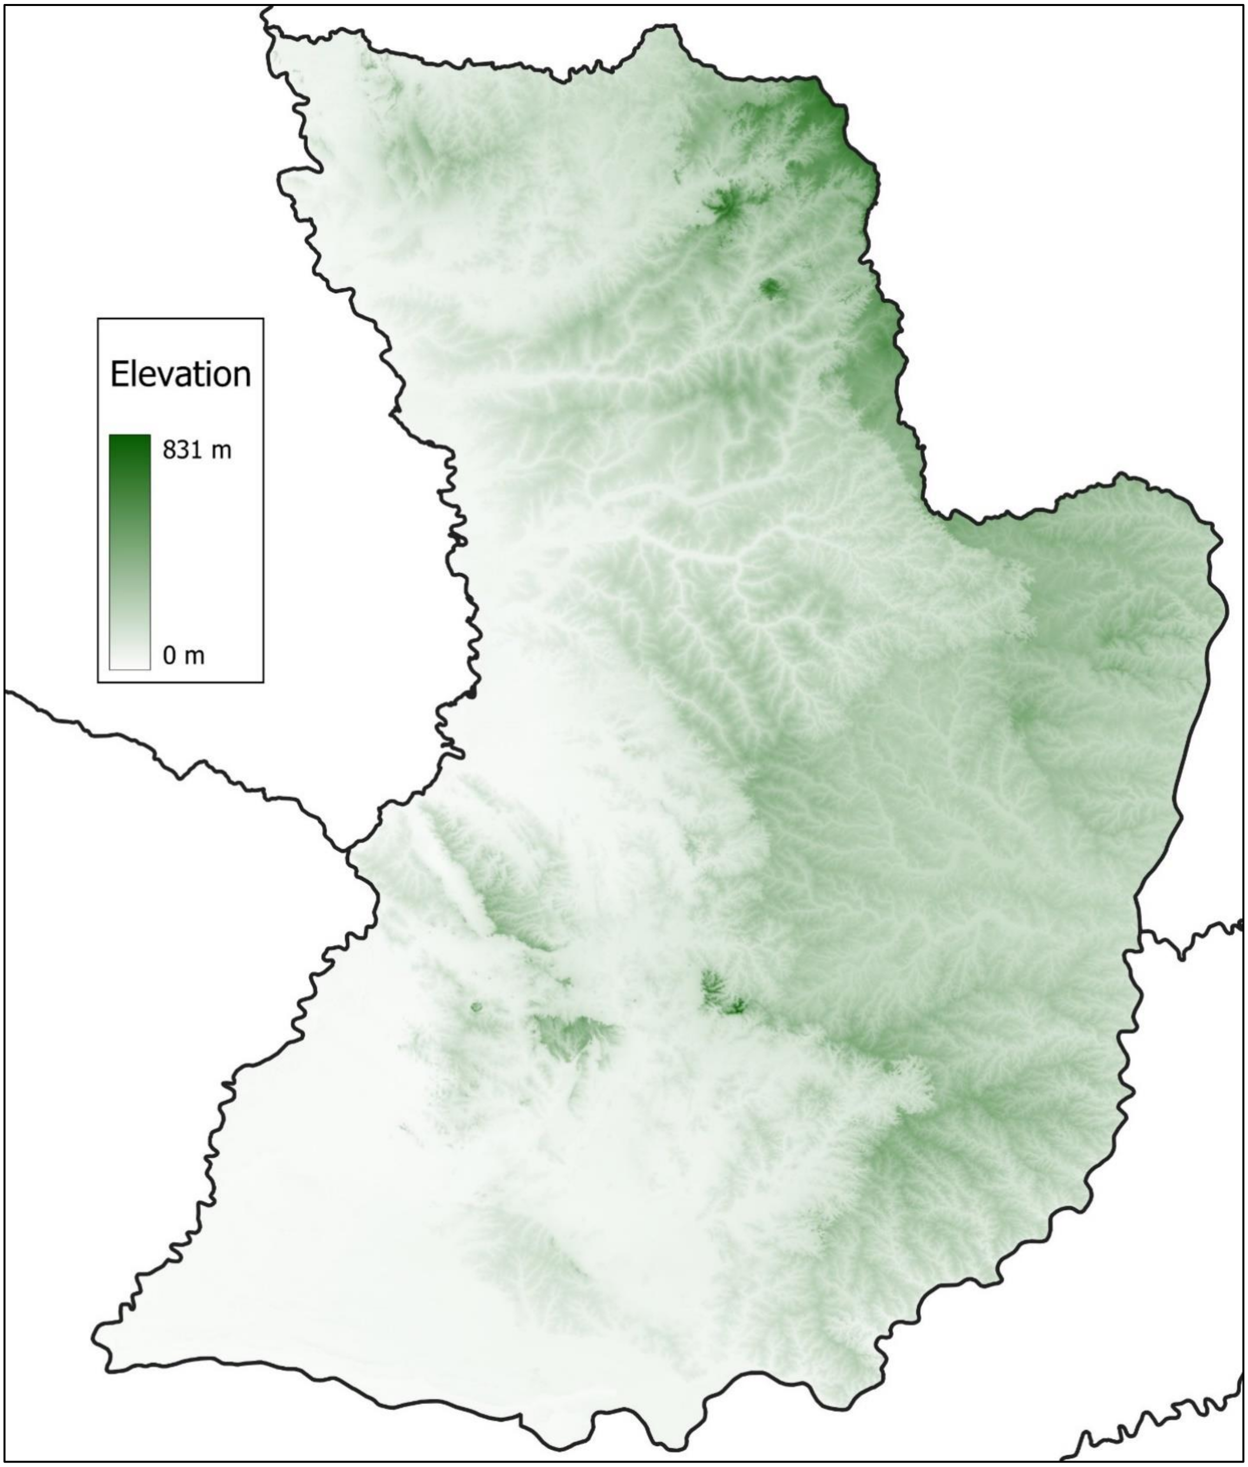

Figure 3: A map of elevation in Eastern Paraguay, measured in meters (Potapov et al., 2020). In general, lower elevation areas are found in the east and higher elevation areas are found in the west.

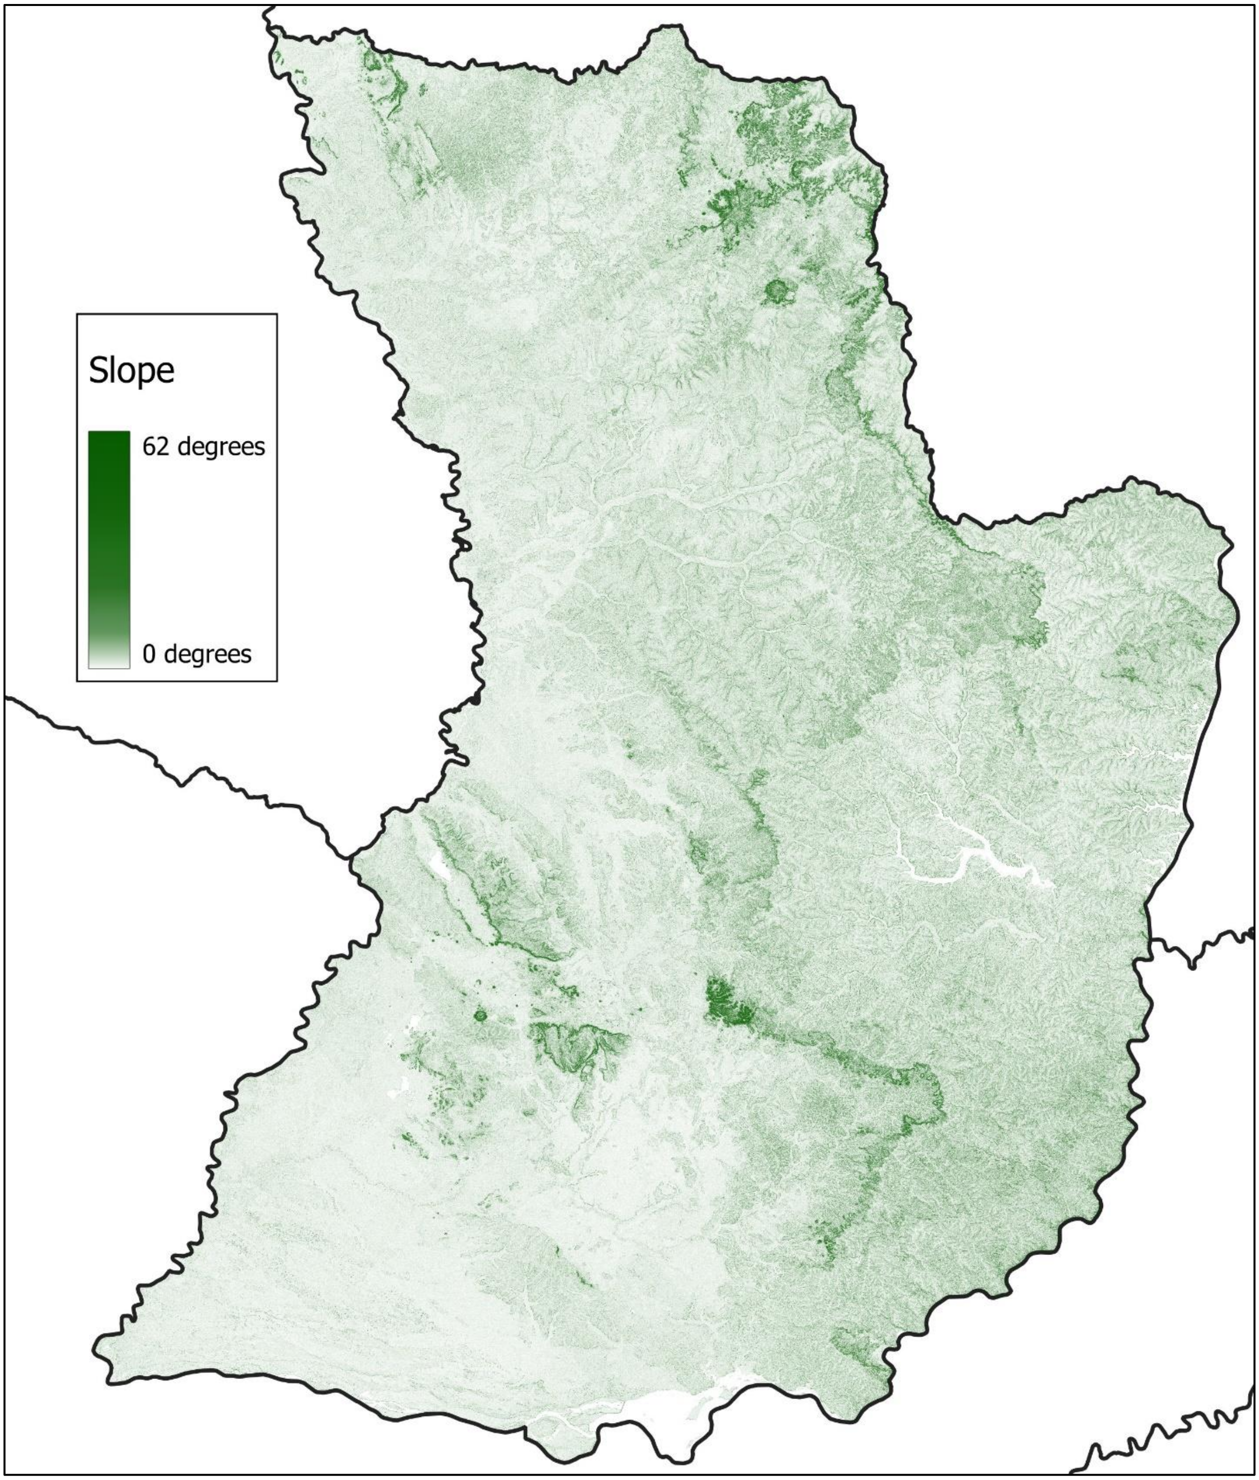

Figure 4: A map of slope in Eastern Paraguay, measured in degrees (Potapov et al., 2020).

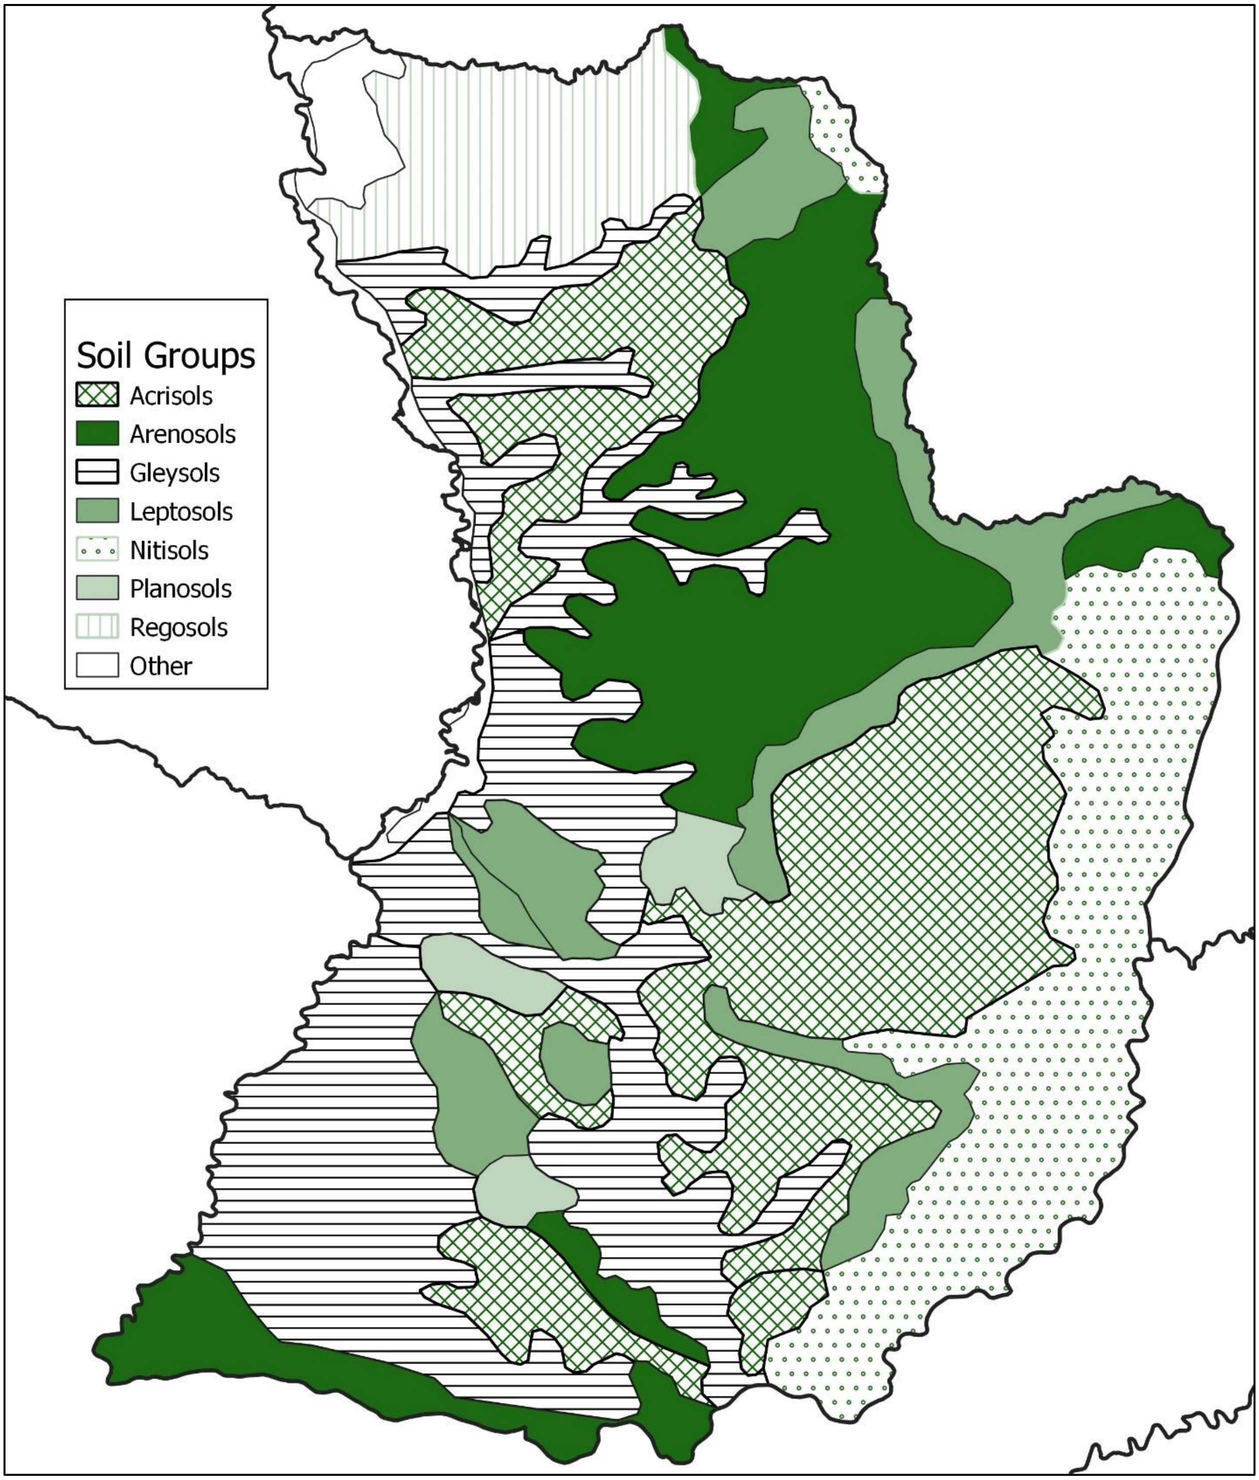

Figure 5: A map of soil groups in Eastern Paraguay (ISRIC, 2005). See appendix A2 for information on soil groups.

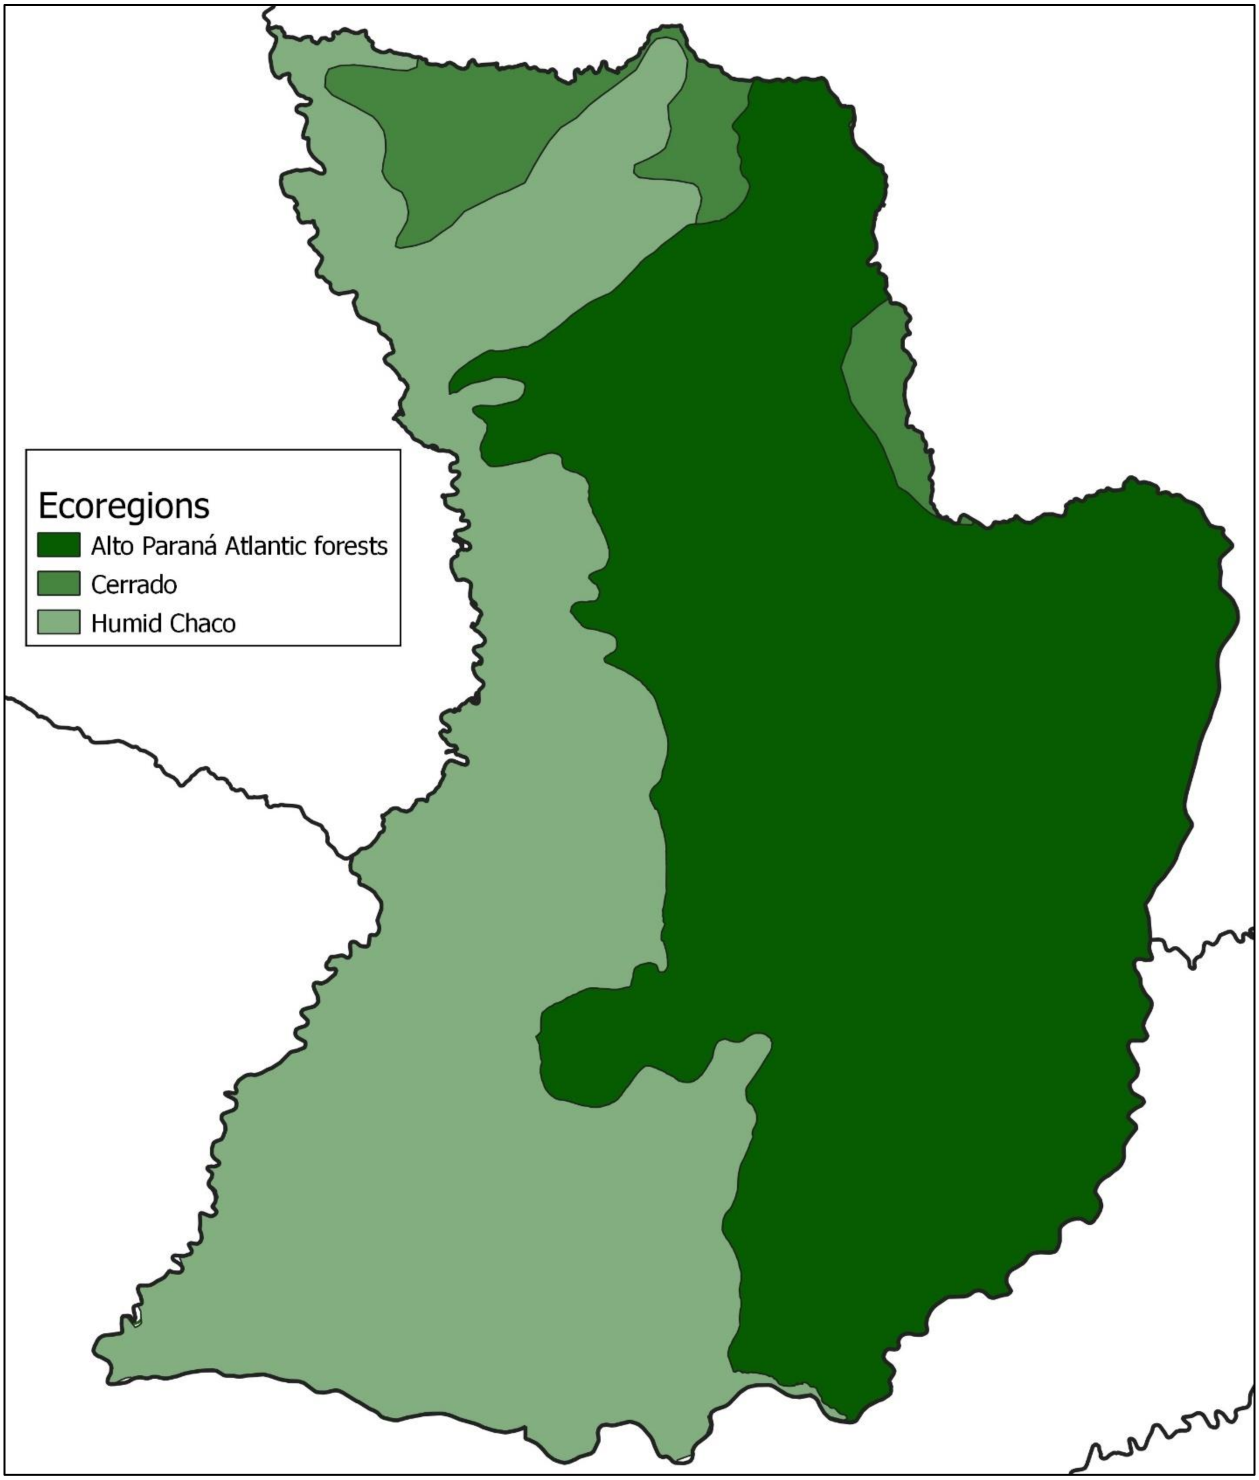

Figure 6: A map of ecoregions in Eastern Paraguay (Olson et al. 2001). Ecoregions classify areas based on the dispersion of communities and species prior to major land use change.

I consider anthropological variables including distance from main roads and towns, brightness of nearby nighttime lights, and measures of nearby tree cover. Distance from a main road is calculated using spatial road data from Open Street Maps (Open Street Maps, 2021). The focus is on roads classified with codes 5112 and 5113. These include major highways connecting urban areas, as well as important rural transit routes. The latter includes relatively low traffic routes that may contain large unpaved segments but are important for the transport of agricultural goods. A map of this road network can be seen in Figure 7.

Nighttime lights (NTL) is a measure of nocturnal visible and near-infrared lights visible from satellite (Roman et al., 2021). The data are cleaned to correct for cloud cover, moonlight, atmospheric radiance, and other distortions. This represents human activity patterns. Annual nighttime lights values from 2012, the earliest year available, are used. The values are relative, with higher values indicating brighter lights and lower values indicating weaker or no light. In Eastern Paraguay, a value of 1000 corresponds to cities, a value of 200 corresponds to a large town, and single digit values correspond to settlement towns with population of a few hundred or less, or other rural infrastructure such as agricultural buildings with lights. Values between these indicate areas of human development between these scales. NTL data is used to generate distance from a town, where a town is defined as areas with NTL values of 200 or more. This data is also used to generate the NTL value at a pixel location, the brightest NTL value within one kilometer and the brightest NTL within ten kilometers. The latter two indicate whether there is a settlement, town, or city within these radii of a pixel. The location of towns and larger population centers can be seen in Figure 8.

Some locations are preserved or otherwise protected by public or private entities. I take data on protected areas from the World Database on Protected Areas and other effective area-based conservation measures (UNEP-WCMC, 2019). The map in Figure 9 displays area that were protected before 2000, the baseline year of my study, and additional locations that were protected between 2001 and 2008.

Mean tree cover in the year 2000 is a measure of the initial forest cover. This is given for the pixel location, and calculated for a one-hundred-meter and a one-kilometer radius. A mean tree cover value of 50 could indicate that all pixels have fifty percent tree cover, or could indicate a densely forested area that is partially cleared. The standard deviation of tree cover in 2000, calculated for the same radii, helps complete the forest cover picture by providing a measure of patchiness. In areas with uniform tree cover standard deviation is low. Where tree cover is partially cleared the standard deviation is higher. A map of tree cover in the year 2000 can be seen in Figure 10.

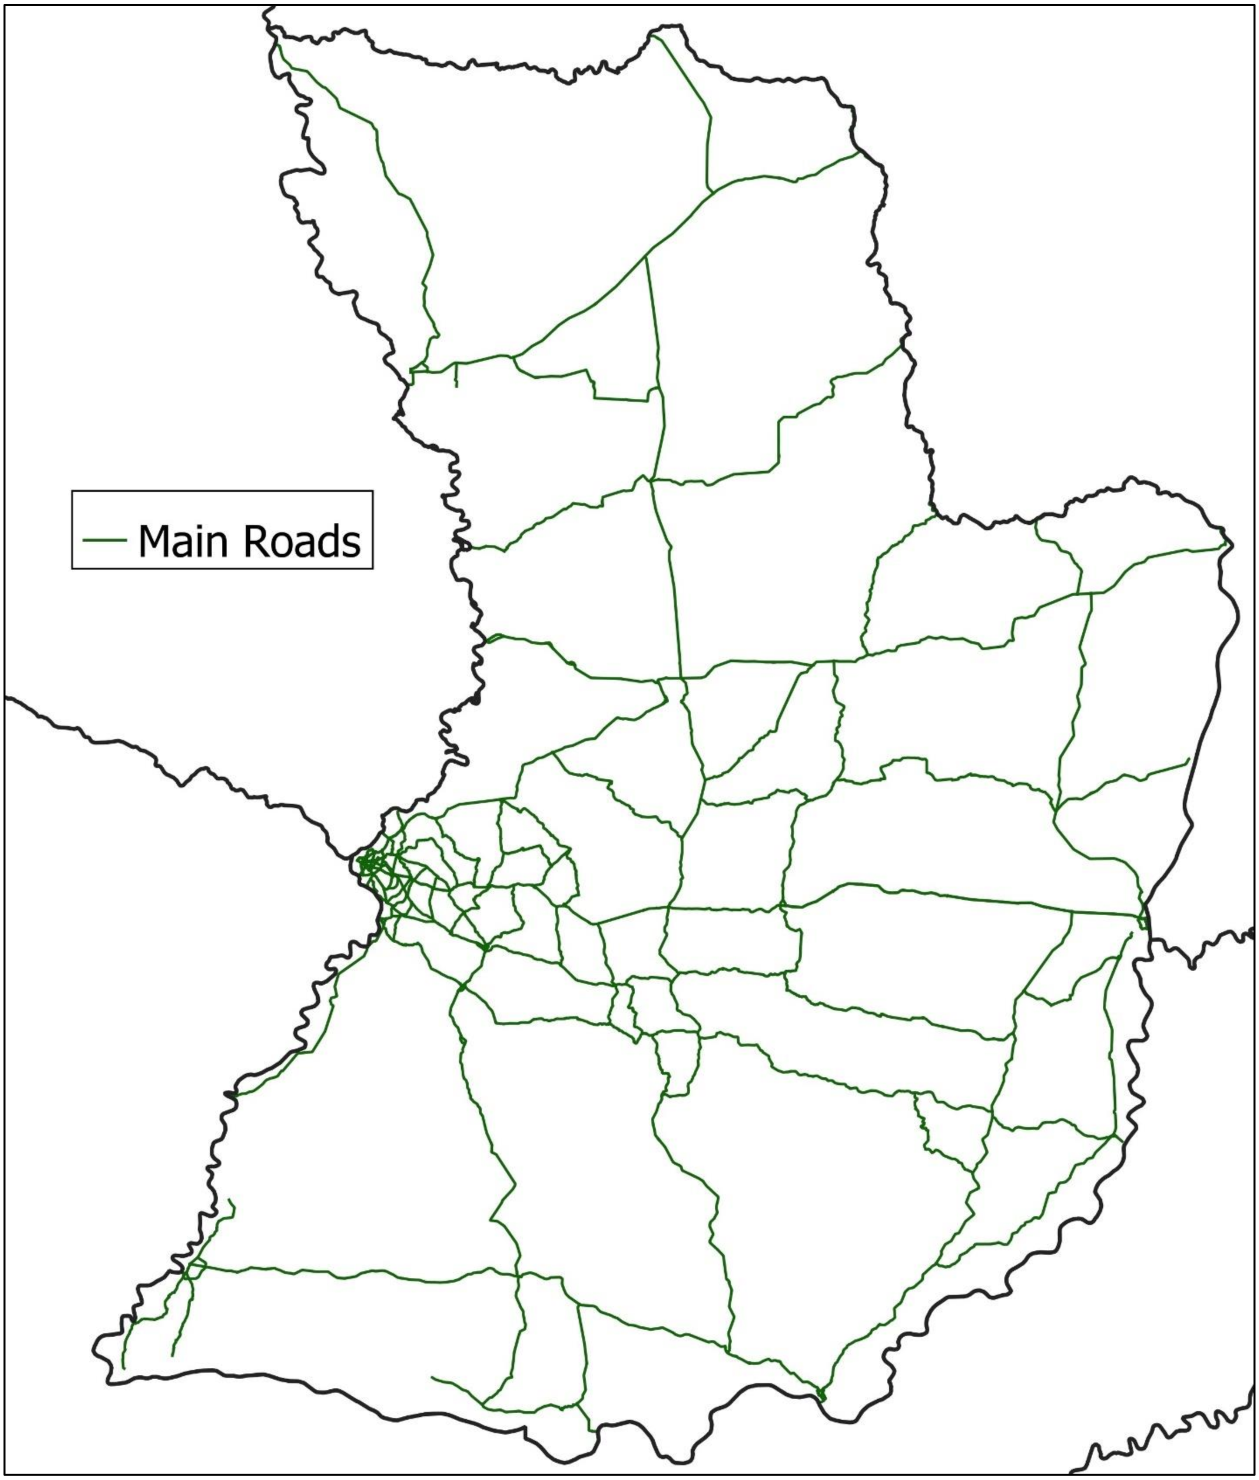

Figure 7: A map of the main roads in Eastern Paraguay, including both major highways and important rural transit routes (Open Street Maps, 2021).

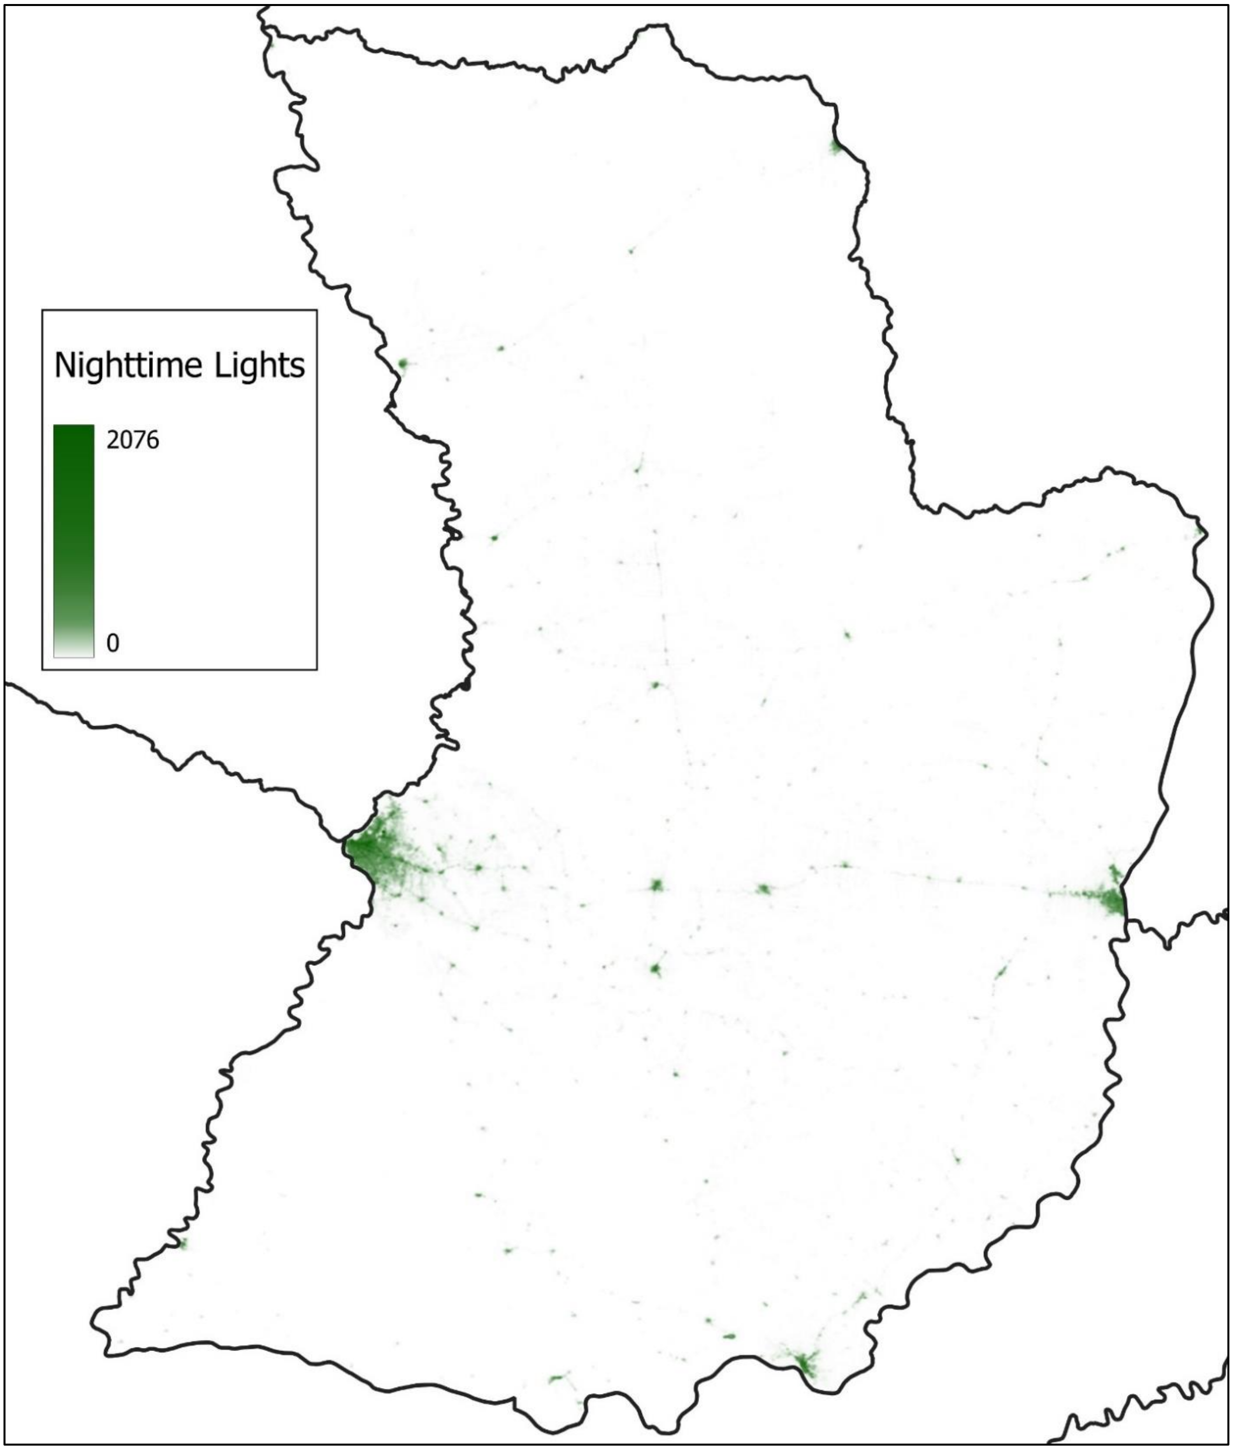

Figure 8: A map of nighttime lights values in Eastern Paraguay (Roman et al., 2021). Nighttime lights (NTL) is a measure of nocturnal visible and near-infrared lights visible from satellite, and is commonly used as a measure of economic development.

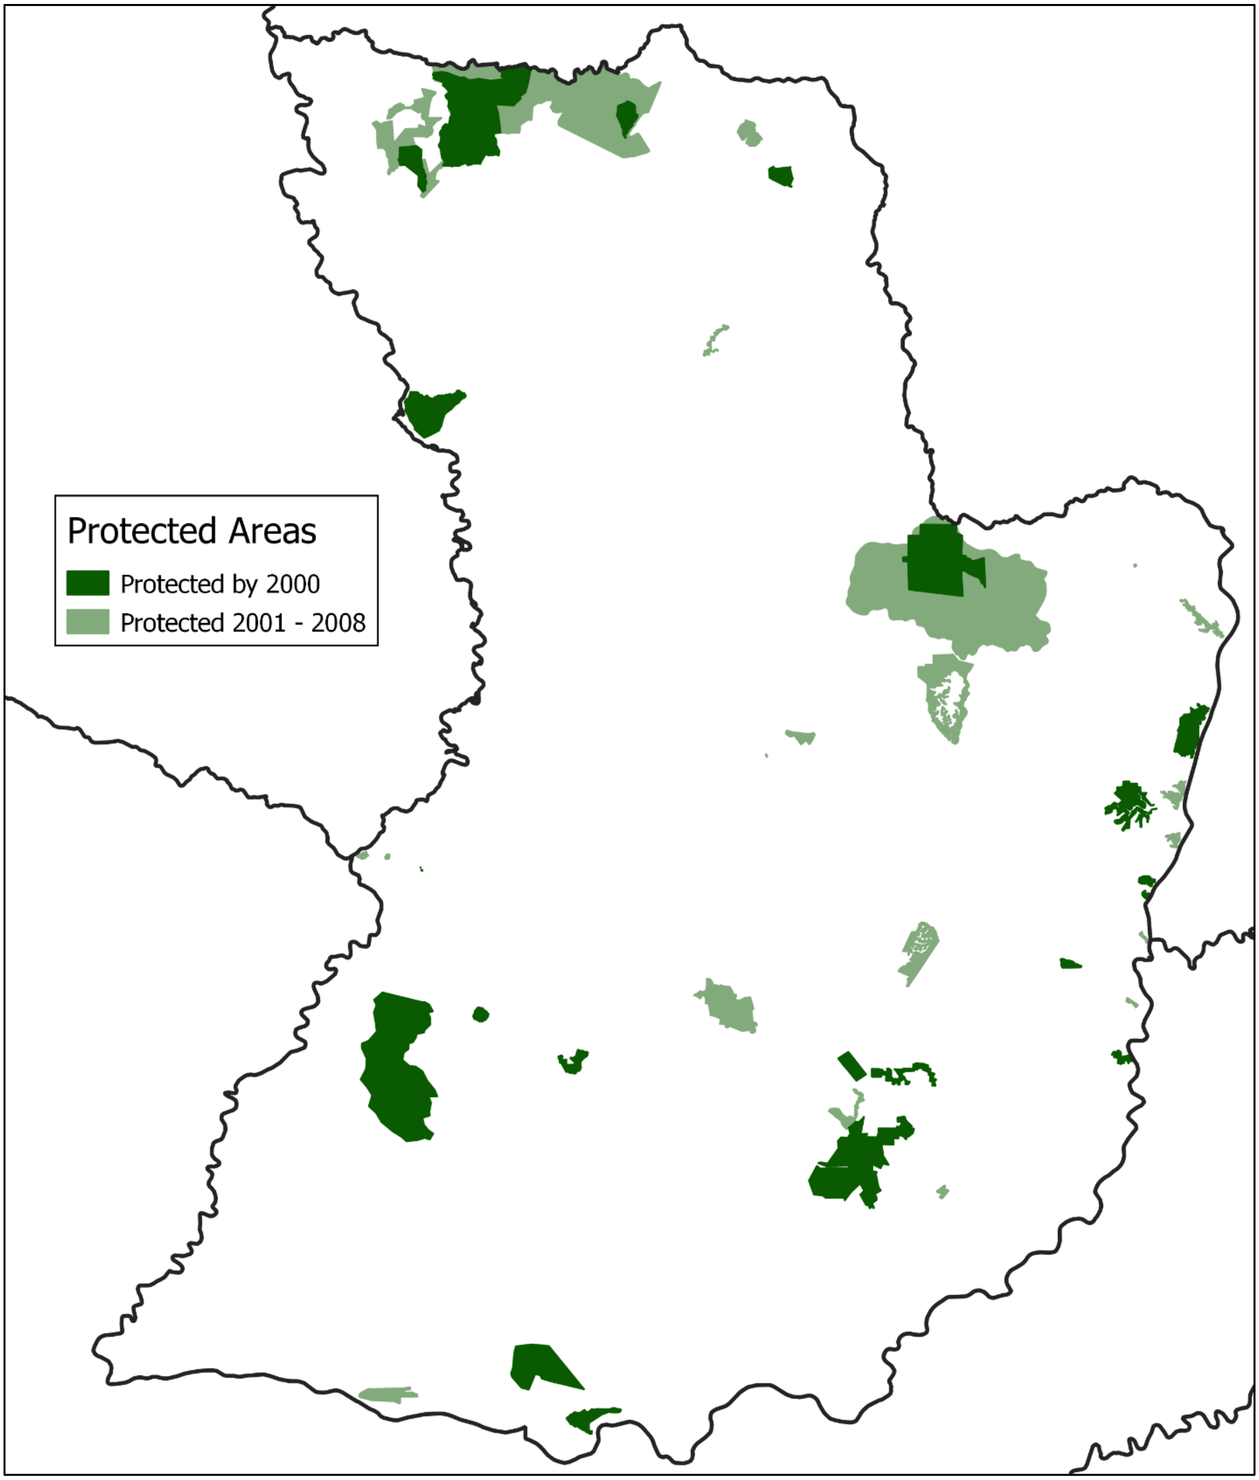

Figure 9: A map of protected areas in Eastern Paraguay (UNEP-WCMC, 2019).

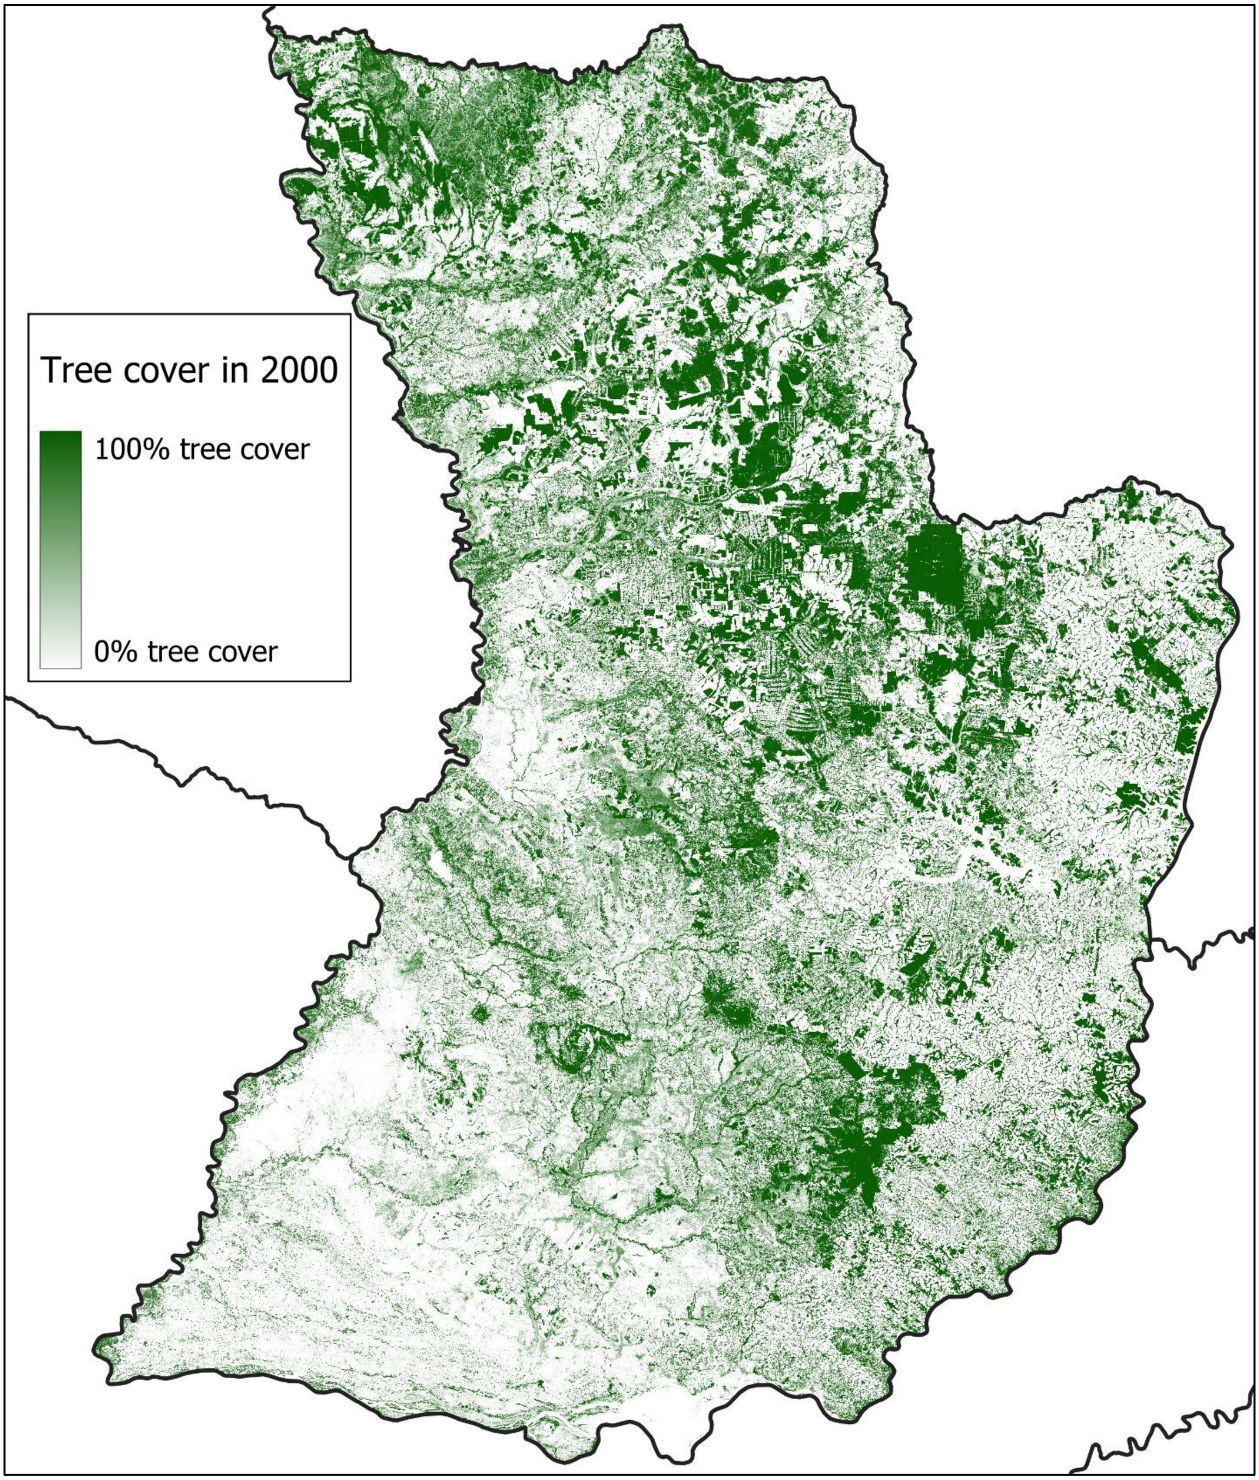

Figure 10: A map of tree cover in 2000 in Eastern Paraguay, measured as the percent of a 30m pixel under forest cover (Hansen et al., 2013).

Macroeconomic Variables

I include commodity prices and unemployment data to investigate whether macroeconomic trends explain deforestation, rather than changes in Paraguay's deforestation policy. Global commodity prices come from the International Monetary Fund (IMF), and are downloaded using an annual timestep from the Federal Reserve Bank of St. Louis (IMF, 2024a; IMF, 2024b; IMF, 2024c). I consider prices for soy, cotton, and beef, the most important market-oriented commodities for large-scale producers, small-scale producers, and ranchers respectively. Soybean price is measured in U.S. dollars per metric ton, while cotton and beef prices are measured in U.S. cents per pound. In all cases, values represent benchmark prices determined by the largest exporter of a commodity and are considered representative of the global market. I expect the global prices of these commodities to influence commodity production decisions in Paraguay. However, Paraguay's production does not make up a large portion of global supply, therefore I do not expect changes in Paraguay's production to impact global prices. All prices are deflated using a consumer price index for Paraguay with a base year of 2010, from the International Monetary Fund and downloaded from the World Bank (IMF, 2024d).

Unemployment data provides information on the state of Paraguay's economy. When unemployment is higher, I expect that more people may rely on agriculture, specifically subsistence agriculture, for their livelihoods. The unemployment rate is provided by Paraguay's National Institute of Statistics (INE, 2024).

III. Drivers of Deforestation in Eastern Paraguay

In this section I investigate how previously identified drivers of deforestation serve as predictors of deforestation in Eastern Paraguay before 2004, when deforestation was a legal activity.

Understanding the drivers of deforestation is important for managing deforestation dynamics. By 2004, Paraguay had reached a point where less deforestation was desired. Identifying which land characteristics are associated with higher deforestation makes it possible to predict where deforestation is most likely to occur next. This is valuable information for law enforcement, NGOs, or other groups working to slow deforestation and need to prioritize the use of limited resources in a way that best fits their goals.

Deforestation is a land use question, and whether to clear a plot of land or not depends on the characteristics of the land as well as the resources and motives of the person making the land use decision. The best land use choice for an individual landowner is the one that leads to the highest utility, a measure encompassing both monetary and nonmonetary benefits. This approach draws on Ricardian theory, where land rents are determined by biophysical characteristics of the land, as well as Thunian theory, which incorporates distance to markets and transport costs. Given a set of resources and incentives, the best choice may be to leave the parcel forested and enjoy benefits such as firewood, shade, and leisure activities. Alternatively, the best choice may be best to clear the parcel and engage in small-scale labor-intensive agricultural production, commercial agricultural production, or ranching, or to clear for urban development or infrastructure. Using this framework, it is possible to identify factors that may promote more or less deforestation in an area. These factors are referred to as drivers of deforestation.

A large literature has focused on identifying the drivers of deforestation (Busch and Ferretti-Gallon, 2017). These drivers can be physical characteristics of the land, such as the terrain and soil characteristics, or anthropological measures related to human activity, such as the proximity to infrastructure or population centers. Identifying the drivers of deforestation sheds light on which locations are at higher risk of deforestation. This knowledge can inform land use policies and shape land protection priorities.

Agriculture is a major driving force for deforestation. A recent report from the Food and Agriculture Organization of the United Nations found that nearly 90 percent of global deforestation between 2000 and 2018 was related to agriculture (FAO, 2022). The importance of agriculture is supported by the land characteristics that have been found to promote deforestation in prior literature. Previous analyses of physical drivers of deforestation found that deforestation is more likely to occur where land is physically better suited for agricultural production. Deforestation is less likely to occur on steep slopes and at higher elevations, and more likely where soil is more productive for local agriculture.

Regarding anthropological land characteristics, the literature again emphasizes the importance of agriculture. Deforestation is more likely closer to cities or towns, especially where local agricultural input markets are well developed (Garrett et al., 2013). Deforestation is more likely closer to roads, which facilitate access for people as well as the transport of machinery and agricultural products. Deforestation is more likely nearer to existing agriculture and previously cleared land. Deforestation is less likely to occur where it is legally restricted, particularly if law enforcement is strong.

Relationships between commonly identified drivers and deforestation vary across settings. For example, in areas dominated by agroforestry production proximity to cities and soil suitability for agriculture have been found to increase forest cover, rather than decrease it (Blackman et al., 2008). Contrary to frequently held expectations, logging activity is not consistently associated with higher deforestation. Logging activity can be a driver of deforestation, but can also be a product of managed agroforestry systems, which include planned cycles of planting and cutting, or may increase dependence on local forests leading to more sustainable forest management.

In this analysis I test whether the physical and anthropological characteristics discussed above hold in the context of Eastern Paraguay. I test seven specific hypotheses to determine if relationships between commonly identified drivers of deforestation and deforestation hold in this setting.

My results show that before the Zero Deforestation Law was implemented, physical land characteristics were significant predictors of deforestation in Eastern Paraguay. Only a few of the anthropological land characteristics were helpful in predicting where deforestation will occur during this period. Measures of nearby tree cover behaved as expected and were statistically significant predictors of deforestation. However, the hypothesized relationships with large population centers, towns, and roads did not hold, with the latter two showing no significant relationship.

This analysis contributes to the literature on drivers of deforestation by expanding the geographic area from which our knowledge is drawn. Existing deforestation studies focus on a few locations. Brazil alone represents around half of all deforestation studies conducted in Latin America (Da Ponte et al., 2015; Busch and Ferretti-Gallon, 2017). Within this, much of the research focuses on the Brazilian Amazon, which is located two thousand kilometers to the north of Eastern Paraguay (i.e. Alix-Garcia and Gibbs, 2017; Garrett et al., 2013, Macedo et al., 2012). These prior studies reveal that suitability for agriculture and proximity to infrastructure and population centers are important predictors of deforestation (Busch and Ferretti-Gallon, 2017). Widening the geographic area of study broadens our understanding of deforestation dynamics. Eastern Paraguay's different administrative, cultural, and policy environment provide a valuable opportunity to check the external validity of these trends. This is especially valuable since eastern Paraguay is dominated by similar agricultural activities as the Brazilian Amazon, and thus deforestation for agriculture might be expected to behave similarly (De Sy et al., 2015; Laso Bayes et al., 2022; Curtis et al., 2018). I show that some of the frequently identified patterns in drivers of deforestation do not hold in the Eastern Paraguayan context. While suitability for agriculture remains an important driver of deforestation, proximity to infrastructure and population centers is less predictive of deforestation. This may be because Eastern Paraguay is already heavily deforested and clearing no longer takes place along a well-defined agricultural frontier. Understanding how drivers perform as predicted or differ from previous studies in a new context, such as in Eastern Paraguay, enables better deforestation policy and action in Paraguay and in other areas that are underrepresented in deforestation research.

This analysis also contributes to knowledge of deforestation dynamics in Eastern Paraguay during the naughts. Prior research reveals that Paraguay was a top contributor to deforestation at both the regional and global levels during the 2000s (i.e., Austin et al., 2017, Hansen et al., 2013). These national-level deforestation statistics are difficult to interpret due to significant differences between Eastern and Western Paraguay. Population and agricultural patterns also vary significantly between Eastern and Western Paraguay. As of 2002, 97 percent of Paraguay's population resided in the more hospitable Eastern region which supports a wide range of agricultural activities and rangeland (DGEEC, 2004). Western Paraguay, in contrast, is sparsely populated and agriculturally is almost entirely rangeland (De Sy et al., 2015; Laso Bayes et al., 2022; Caldas et al., 2015; Fehlenberg et al., 2017). In addition, the Zero Deforestation Law, the focus of my later analyses and an important piece of deforestation legislation, only applies in Eastern Paraguay and was implemented mid-decade. This analysis provides important clarification on the deforestation patterns within the Eastern region of the country during the 2000s.

Finally, this research sets the stage for later analysis of how changing deforestation from a legal to illegal activity impacted how, where, and for what purposes deforestation was conducted. Analyzing the drivers of deforestation pre-policy sets a baseline for future work in this area.

Data and methods

The land characteristics investigated in this analysis include elevation, slope, soil group, ecoregion, the mean and standard deviation of tree cover within a 100-meter radius, the maximum nighttime lights value within ten kilometers, the distance from a town, and the distance from a main road.

Data sources for each variable are described in section II. The analysis is run on a random sample of one percent of all pixels in Eastern Paraguay due to computing limitations, and also to partially control for spatial correlation between nearby variables. From this sample, pixels that have forest cover of at least thirty percent, the cutoff used throughout this analysis to define forested areas, in 2000 are retained. This results in just under one million observations.

Table 1 provides summary statistics for the random sample and the predicted relationship between each variable and the probability of deforestation based on the prior literature. I test seven hypotheses. For physical characteristics, I hypothesize that locations with higher elevation will have more deforestation than locations with lower elevation, and that locations with steeper slopes will have less deforestation than flatter locations (Busch and Ferretti-Gallon, 2017). For anthropological characteristics, I hypothesize that locations with higher average nearby mean tree cover will have less deforestation than locations with lower mean tree cover nearby, and that locations with patchier tree cover nearby will have more deforestation than areas that are more uniformly forested. These two hypotheses stem from findings that deforestation is more likely near existing clearings, for example existing agricultural fields, which would exhibit this low mean forest cover and patchiness. I hypothesize that the maximum nighttime light value within 10 kilometers is positively correlated with deforestation, that locations closer to towns will have more deforestation than locations further from towns, and that locations closer to roads will have more deforestation than locations further from roads. These hypotheses stem from findings in previous literature that deforestation is more likely near population centers and infrastructure (Busch and Ferretti-Gallon, 2017). In addition, I expect soil groups and ecoregions that are better for agriculture to be more likely to be deforested (FAO 2022). These hypothesized relationships are summarized in the final column of Table 1.

Variable	Unit	Mean	St. Dev.	Min	Max	Expected relationship with deforestation
Elevation	meters	195.39	92.29	43	812	-
Slope	degrees	3.37	2.93	0	52.21	-
Soil group	categorical	.	.	.	.	N/A
Ecoregion	categorical	.	.	.	.	N/A
Mean tree cover within 100m	percent	63.38	24.21	0.633	95.048	-
St. dev. of tree cover within	value	28.62	7.00	4.425	50.251	+
Maximum NTL within 10km	NTL value	90.99	177.03	0	2142	+
Distance from a town	5 km bins	30.878	16.539	0	100*	-
Distance from a main road	5 km bins	15.346	10.812	0	60	-

* All locations over 95km from a town are included in '100'. This category contains less than 1% of observations

Table 1: Summary statistics for land characteristics in the random sample of forested pixels, and hypotheses on the relationship between each of these land characteristics and deforestation.

The land characteristic measures are not highly correlated, as shown in the variance covariance matrix in Table 2.

	elevation	slope	mean cover	sd cover	max NTL	dist town	dist road
elevation	1.00	0.35	0.29	0.21	-0.11	0.08	-0.01
slope		1.00	0.18	0.20	0.00	0.00	-0.02
mean tree cover within 100m			1.00	0.17	-0.15	0.18	0.07
st. dev. of tree cover within 100m				1.00	0.05	-0.10	-0.10
maximum NTL within 10km					1.00	-0.46	-0.33
distance from a town						1.00	0.54
distance from a main road							1.00

Table 2: Variance covariance matrix showing the low correlation between land use characteristics used to predict deforestation in the pre-policy period.

I test the relationship between these land characteristics and deforestation in Eastern Paraguay using a linear probability model:

$$D_i = \beta X_i + \epsilon_i$$

The outcome variable, D_i, is an indicator equal to 1 if a pixel lost forest cover between 2001 and 2004, and equal to 0 otherwise. X_i is a vector of time-invariant land characteristics that have been identified as drivers of deforestation in the literature. All continuous variables are normalized by dividing by their standard deviation and demeaned. ϵ_i is an error term, and errors are clustered at a half-degree latitude longitude grid to control for spatial correlation.

Additionally, I run specifications where the outcome is deforestation in a single year. This tests whether the relationships between potential drivers of deforestation and the probability of deforestation are consistent across time, as opposed to being driven by strong patterns from a single year. Pixels that lose forest cover before the year in question are removed.

Results

By the end of 2004, 6% of pixels that were forested in 2000 had been deforested. The rate of deforestation increased slightly throughout this period, from just under 1% of forested pixels lost in 2001 to just over 2% of forested pixels lost in 2004. This is due primarily to increased deforestation during these years, as can be seen in the annual deforestation in appendix A1, as well as to a smaller total remaining forested area in each year in the denominator.

Results from the linear probability model regressions are shown in Table 3. The first specification examines the relationship between land characteristics and deforestation that occurred between 2001 and 2004, before the policy made deforestation illegal. Some of the relationships found here differ from results in the prior literature. Physical characteristics, including elevation, slope, soil group, and ecoregion, are significant predictors of deforestation, though not always in the expected manner. Anthropological characteristics are not as predictive as the physical characteristics are. I address each of the coefficients of the potential drivers of deforestation in this specification in more detail below. The discussion refers to the first specification unless noted otherwise.

	Dependent variable:				
	lost 2001-2004	lost 2001	lost 2002	lost 2003	lost 2004
	(1)	(2)	(3)	(4)	(5)
elevation	0.034^{***}	0.005^{***}	0.010^{***}	0.010^{***}	0.011^{***}
	(0.005)	(0.001)	(0.001)	(0.002)	(0.002)
slope	-0.017^{***}	-0.003^{***}	-0.004^{***}	-0.005^{***}	-0.007^{***}
	(0.002)	(0.0003)	(0.001)	(0.001)	(0.001)
tree cover mean, 100m	-0.006^{**}	-0.003^{***}	-0.003^{***}	-0.002^{**}	0.001
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
tree cover stdv, 100m	0.018^{***}	0.004^{***}	0.004^{***}	0.006^{***}	0.005^{***}
	(0.002)	(0.001)	(0.0005)	(0.001)	(0.001)
max NTL 10km	-0.004^{**}	-0.0002	-0.001^{**}	-0.001^{*}	-0.001^{**}
	(0.001)	(0.0005)	(0.0004)	(0.001)	(0.001)
distance to a town	0.001	0.002^{**}	0.001	-0.001	-0.001
	(0.003)	(0.001)	(0.001)	(0.001)	(0.001)
distance to a main road	0.001	0.001	0.0003	0.001	-0.0002
	(0.003)	(0.001)	(0.001)	(0.001)	(0.002)
SE clusters	1/2 degree	1/2 degree	1/2 degree	1/2 degree	1/2 degree
% forest cleared	5.8%	0.9%	1.2%	1.6%	2.1%
Ecoregions	Yes	Yes	Yes	Yes	Yes
Soil groups	Yes	Yes	Yes	Yes	Yes
Observations	928,277	928,277	920,255	908,790	893,827
R^2	0.095	0.016	0.022	0.029	0.035
Adjusted R^2	0.095	0.016	0.022	0.029	0.035
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 3: Results from linear probability models relating potential drivers of deforestation to the probability of deforestation.

Between 2001 and 2004 deforestation was more likely to occur at higher elevations, contrary to previous findings. Regression results in from the first specification show that a one standard deviation increase in elevation, which is an increase of around 100 meters in Eastern Paraguay, increases the probability of deforestation between 2001 and 2004 by 0.034 percentage points. This is nearly a 60% increase from the average probability of deforestation of 5.8% between 2001 and 2004. That the probability of deforestation increases with elevation, rather than decreases as was hypothesized, may be because Eastern Paraguay contains low-lying areas that experience issues with wetness and waterlogging. Meanwhile, the highest elevation forested locations are at only around 800 meters of elevation, thus the region does not suffer from expanses of inhospitable high elevations.

Slope is negatively correlated with deforestation, as expected. A one standard deviation increase in slope leads to a 0.017 percentage point decrease in the probability of deforestation. This is expected because flatter areas are more suitable for agriculture, which is the main use of deforested land in Eastern Paraguay.

Anthropological land characteristics are not as clearly linked to deforestation in Eastern Paraguay. The literature predicts that deforestation will be more likely nearer to population centers and roads. However, the coefficients on distance from a town and distance from a main road are both insignificant. This mismatch may be because past studies often focused on deforestation that encroaches into large, uncleared areas such as the Brazilian Amazon. In Eastern Paraguay, by contrast, deforestation is widespread and current deforestation removes remaining patches rather than expanding the frontier.

While roads and towns are not significant predictors of deforestation, some anthropological land characteristics do play a role. Deforestation becomes less likely as nighttime lights within a ten-kilometer radius become brighter, indicating that deforestation is less likely in the immediate vicinity of cities or large towns. This is the opposite of the hypothesized relationship, that more deforestation occurs near population centers. This may be because agriculturally suitable land close to very dense population centers has already been cleared. However, the magnitude of the coefficient is small. My results show that a one standard deviation increase in the brightest nearby nighttime lights value, around a 200-point jump, results in a 0.004 percentage point decrease in the probability of deforestation. This corresponds to a 7% change in the probability of deforestation, from 5.8% probability of being deforested on average to a 5.4% probability of being deforested in areas where the brightest nighttime light within 10 kilometers is one standard deviation higher than average, all else constant. However, the distribution of values for the nearest nighttime lights are heavily skewed to the right. Most forested pixels have low nearby nighttime lights values, with over half the observations having a value below 26, and 86% of forested pixels falling within one standard deviation of the minimum value of 0. Therefore, while this relationship is statistically significant, its contribution to predicting where deforestation occurs is small.

Nearby tree cover metrics also significantly predict deforestation. While tree cover can be the result of natural suitability of the land for forest, in developed areas it is often the result of agriculture and other human-conducted clearing. A one standard deviation increase in the standard deviation of nearby tree cover leads to a 0.018 percentage point increase in deforestation. This is expected because areas with patchy forest due to agricultural clearings are more likely to be deforested than uniformly forested areas that have not experienced this partial clearing. A one standard deviation increase in mean tree cover within 100 meters leads to a 0.006 percentage point decrease in the probability of deforestation. This aligns with expectations that deforestation is more likely to occur where there is less nearby tree cover, potentially due to existing manmade clearings. The magnitudes of these coefficients are small, meaning that a change in their standard deviation is less predictive than, say, a one standard deviation in slope. However, they are more predictive than the brightest nearby nighttime lights, which has a similar magnitude, due to the values of these variables having a relatively less skewed distribution.

Soil groups and ecoregion are also statistically significant predictors of where deforestation occurs. The coefficients for these land characteristics can be found in appendix A3. Because soil groups and ecoregions cover large, continuous areas these controls essentially serve as additional fixed effects.

The final four specifications examine the relationship between land characteristics and deforestation year-by-year. The magnitudes of the coefficients on these annual deforestation regressions are lower to accommodate the lower area deforested in a one-year period as compared to a four-year period. For example, a one standard deviation change in elevation leads to a 0.005 percentage point increase in the probability of deforestation in 2001. This is a 55% increase from the annual average rate of 0.9%, comparable to the 57% increase in the probability of deforestation from a one standard deviation increase in elevation found in the pooled 2001-2004 model. The average rate of deforestation for each year is listed at the bottom of the table.

The direction and significance of relationships from the single-year regressions are generally consistent with the pooled regression findings. This indicates that the relationships identified above between land characteristics and deforestation are not driven by anomalies in any single year. In the annual deforestation linear probability models, elevation, slope, and the standard deviation of tree cover remain strongly significant predictors of deforestation. The mean nearby tree cover and the maximum nighttime lights value within 10 kilometers are weakly predictive of deforestation. The distance from a town and the distance from a road consistently lack predictive power.

Additional specifications in appendix A3 show that these findings are generally robust to the inclusion of additional fixed effects, variations in the clustering of standard errors, and to using a logit specification in place of the linear probability model.

Discussion

This investigation into drivers of deforestation in Eastern Paraguay provides a baseline for the patterns of deforestation that dominated in the years before the Zero Deforestation Law was implemented. Some of the land characteristics identified in the previous literature as drivers of deforestation are statistically significant predictors of deforestation during this time. However, the relationships that I found do not always match the results from previous research.

In Eastern Paraguay, land that is physically better suited for agriculture is more likely to be deforested, as expected. Anthropological drivers, however, are less helpful than physical characteristics in determining where deforestation will occur. While the literature suggests that land nearer to population centers and infrastructure is more likely to be deforested than land that is further from these, I do not find that this holds. Proximity to towns and to roads does not significantly predict deforestation. On the other hand, proximity to patchily forested areas, which can indicate proximity to existing man-made clearings such as fields, is associated with higher levels of deforestation.

This deviation from expected relationships likely results from the stage of agricultural development and deforestation in Eastern Paraguay. Much of the literature focuses on less deforested regions where deforestation is removing land on a forested frontier, such as in the arc of deforestation in the Brazilian Amazon (Kalamandeen et al., 2021). In these areas, deforestation moves outwards into the forest from established roads and settlements. However, Eastern Paraguay has been deforested to the point where there is no clear forested frontier. Deforestation instead removes patches that remain between existing developed areas, some of which are near to towns and roads while others are not.

The statistically significant drivers of deforestation show that the characteristics that underlie land's suitability for agriculture were of vital importance when predicting where deforestation will occur in Eastern Paraguay. Meanwhile, anthropological characteristics cannot be counted on as heavily to predict where deforestation is more likely to occur. This implies that understanding the accessibility of a location is not a high priority when looking to understand its risk of deforestation. This provides valuable information for groups, such as law enforcement of NGOs, who wish to influence or eliminate deforestation.

In addition to providing information that can help to enforce the Zero Deforestation Law, this understanding of pre-policy deforestation dynamics sets a baseline for understanding how deforestation changed after the Zero Deforestation Policy came into effect.

IV. The Zero Deforestation Law and Aggregate Deforestation

In this section, I investigate whether and how the Zero Deforestation Law impacted aggregate deforestation in Eastern Paraguay, and whether it impacted how drivers of deforestation predict deforestation. Theoretically, the law should have eliminated deforestation altogether. A quick look at deforestation after the Zero Deforestation Law came into effect (appendix A1) shows that this did not happen. However, even though it did not eliminate deforestation, the Zero

Deforestation Law may still have slowed it. In this section, I test whether a decrease in deforestation following the implementation of the Zero Deforestation Law occurred.

Understanding the impact of the Zero Deforestation law on deforestation and on drivers of deforestation is important because it was such a groundbreaking piece of legislation in Paraguay. The outcome of this law can be used as evidence for or against similar command-and-control style legislation to slow deforestation in similar settings. Understanding the strengths and shortcomings of the policy can help Paraguay to develop more effective future policies and future policy implementation in the attempt to quell excessive deforestation.

I use a linear probability model on a panel dataset of deforestation to investigate how deforestation changed after the Zero Deforestation Law was passed. My results show that deforestation fell significantly after the law was implemented. In addition, I find that commonly identified drivers of deforestation relate to deforestation differently in the pre- and post-policy periods. After deforestation became illegal, many variables become less predictive of deforestation. Deforestation also goes from being more common in areas with lower forest cover, to being more common in areas with higher forest cover, perhaps indicating a desire to hide the illegal action behind remaining tree cover. This analysis is not causal, but it provides strong suggestive evidence that the Zero Deforestation Law led to a decrease in deforestation. Other significant events that occurred during this time frame do not explain this decrease in deforestation.

This analysis contributes to our understanding of deforestation dynamics in Eastern Paraguay. As mentioned in section III, Paraguay was a top contributor to deforestation globally during the 2000s (i.e., Austin et al., 2017, Hansen et al., 2013). The Zero Deforestation Law was a major piece of legislation meant to reverse this trend. The impact of this groundbreaking law on deforestation has not been studied in the academic literature. This study fills this gap.

In addition, this analysis contributes to the literature on drivers of deforestation. The implementation of the Zero Deforestation Law provides a new setting in which to investigate how the relationships between land characteristics that have been identified as common drivers of deforestation and deforestation change when deforestation becomes illegal. Understanding how drivers of deforestation change when deforestation becomes illegal will enable better policies by allowing policymakers to anticipate the changes in deforestation patterns that a ban may create. It also provides valuable information for enforcement agencies acting with limited resources. Knowing how drivers of deforestation change once deforestation becomes illegal allows these organizations to update their enforcement strategies without needing to wait years for new data on deforestation and a new understanding of where deforestation will most likely occur.

Assumptions for Causality

This analysis investigates whether the Zero Deforestation Law slowed deforestation, which is difficult to test. If the law had been implemented at different times in different regions it would be possible to analyze these multiple changes over time, while other influential factors varied. However, the law was implemented uniformly and immediately. In addition, neighboring regions do not provide suitable comparison groups. Western Paraguay is physically separated from Eastern Paraguay by the Paraguay River and differs dramatically from Eastern Paraguay in population and agricultural practices. Neighboring areas in Brazil and Argentina are also unsuitable because the administrative and legal settings are significantly different. Therefore, to analyze whether the Zero Deforestation Law impacted aggregate deforestation I examine whether there was a level shift in the rate of deforestation in Eastern Paraguay before and after its implementation.

This approach requires assumptions for causality. Specifically, causality requires that no other event caused the observed changes in deforestation between the pre- and post-policy periods. I argue that there are few events of significance that occurred during the years in question, and that the potentially influential events that took place around the time of the implementation of the Zero Deforestation Law acted in the opposite direction of a deforestation ban, to promote increased deforestation in Eastern Paraguay.

Agricultural production of soy, cattle, and, historically, cotton drove deforestation in Eastern Paraguay by large-scale commodity-producing and small-scale subsistence-oriented farmers and ranchers. Therefore, influential events with regards to the production of these commodities might be expected to influence deforestation. I review potential events here.

To begin, I investigate whether fluctuations in commodity prices or unemployment rates can explain deforestation patterns across the pre- and post-policy periods. When prices for a commodity are higher the area devoted to cultivating that commodity may expand, which may cause more deforestation to increase the total amount of available agricultural land. I expect the effect to be immediate, since clearing land occurs early in the process of expanding agricultural area. There may also be a positive relationship between deforestation and lagged prices, as farmers continue to expand their cleared land in response to high prices in the previous year. When unemployment is higher, I expect more people to engage in agriculture, especially small-scale subsistence-oriented agriculture, and potentially to clear more land to do so. It is not problematic for causality if deforestation is heavily influenced by these variables. However, if there is a shift in commodity prices or unemployment that aligns temporally with the policy, then changes in deforestation may be attributable to these changes rather than to changes in the legality of deforestation.

Summary statistics suggest that such a shift did not occur. The average CPI-adjusted market price for soy remained nearly stable between the four years pre-policy and the four years post-policy, going from 262 to 263 CPI-adjusted U.S. dollars per metric ton. The average CPI-adjusted market prices for cotton and beef fell, from 92 to 75 and from 168 to 146 adjusted U.S. cents per pound respectively. Unemployment fell from 8.4% to 5.8% between these periods.

I test whether these macroeconomic variables explain changes in pre- and post-policy deforestation better than the Zero Deforestation Law using a simple regression of total annual deforestation on a post-policy indicator, annual prices, and annual unemployment rates between 2001 and 2008. The regression results are shown in Table 4. The first specification includes only a post-policy indicator. The coefficient is negative, and reveals that annual deforestation in the post-policy period is, on average, 23 thousand hectares less than in the pre-policy period. This relationship is not statistically significant. However, with only eight annual data points this is not surprising.

The second specification includes annual commodity prices and annual unemployment rates, and tests whether these variables better explain the decrease in deforestation than the post-policy indicator. The coefficient on post remains insignificant and becomes substantially more negative. Once commodity prices and unemployment have been controlled for, changes in the post-policy period such as implementation of the Zero Deforestation Law would need to decrease deforestation by 308 thousand hectares per year to explain the lower deforestation seen in the data. Rather than prices explaining the decrease in deforestation between the pre- and post-policy periods, they appear to increase deforestation. The third specification uses lagged prices, and the large negative coefficient on post also suggests that rather than explaining the decrease in deforestation, prices and unemployment may increase deforestation over this period.

	Dependent variable:		
	deforestation		
	(1)	(2)	(3)
post	−22.565	−307.890	−59.517
	(21.766)	(262.919)	(72.170)
price of soy		−0.180	
		(0.367)	
price of beef		−1.529	
		(1.549)	
price of cotton		−8.364	
		(8.440)	
price of soy, lagged			−0.064
			(0.442)
price of beef, lagged			−0.683
			(1.347)
price of cotton, lagged			−0.527
			(1.837)
unemployment		−44.234	−5.784
		(38.914)	(21.796)
Constant	93.420***	1,556.331	334.448
	(15.391)	(1,390.832)	(218.546)
Observations	8	8	8
R^2	0.152	0.551	0.543
Adjusted R^2	0.011	−0.571	−0.599
Residual Std. Error	30.781 (df = 6)	38.790 (df = 2)	39.135 (df = 2)
F Statistic	1.075 (df = 1; 6)	0.491 (df = 5; 2)	0.475 (df = 5; 2)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 4: Regression of total annual deforestation on annual commodity prices and unemployment in Paraguay between 2001 and 2008, with policy implementation at the end of 2004.

In addition to changes in the macroeconomic environment, changes in the policy environment or larger trends in commodity markets might explain deforestation during this time. Cotton, previously a profitable crop for small-scale producers, had fallen to a fraction of its peak area by the early 2000s making it no longer an influential player in land use decisions (Richards, 2011). Genetically modified soy was introduced in 1996 and legalized in 2004 (Correia, 2019; Richards, 2011). Therefore, the legalization of this new production method for soy is in the time frame of interest. However, by 2004 genetically modified soy was already common practice (Peters, 2015), and in any case would be expected to increase, not decrease, deforestation especially for large-scale agriculture.

In 2006, the Soy Moratorium severely limited markets for soy produced on recently deforested land in the Brazilian Amazon. The Brazilian Amazon is geographically distant from Paraguay, and regions including the Brazilian Cerrado stand between the Brazilian Amazon and Eastern Paraguay to absorb spillovers. These limitations on clearing in the Brazilian Amazon would again be expected to increase, not decrease, Paraguayan deforestation through potential spillovers, especially for large-scale production.

Paraguayan soy production has close ties with Brazil (Galeano, 2012), so changes in Brazilian policy may be relevant. In 2006, Brazil banned deforestation in the remaining Brazilian Atlantic Forest. This sudden change in land available for deforestation might have pushed farmers and other agriculturalists across the border to deforest Paraguay instead. However, little forested land remained in the Brazilian Atlantic Forest in 2006. A ban on deforestation on the Brazilian side of the border there would not be expected to generate large spillovers. If it did, that would again be expected to increase, not decrease, deforestation in Eastern Paraguay in the post-policy period.

Non-agricultural factors may have influenced deforestation in Eastern Paraguay. Previous work has shown that deforestation increases around elections (Ruggiero et al., 2021). Visual inspection of deforestation data suggests that deforestation in Eastern Paraguay spikes the year prior to a presidential election after deforestation becomes illegal, as shown in appendix A1. While there is an election in the analysis post-policy period, this again would work opposite the expected impact of the ban to attenuate any detected decrease in deforestation.

Another potential concern is that changes in deforestation may be caused by a shortage of forested land suitable for a certain agricultural use after 2004. After all, by 2004 Paraguay had progressed beyond clearing an empty frontier to clearing areas of remaining forest. If this is the case, deforestation may decrease after 2004 due to a lack of land, rather than due to the new policy. If a change in the availability of land explains changes in deforestation following the Zero Deforestation Law, then I would expect to see a shift in the characteristics of available land that occurred around the time of the law's implementation. I do not. Table 5 provides summary statistics for characteristics of land that remained forested at the beginning of each year. Each column shows the mean and standard deviations of each land characteristic for pixels that had forest cover at the beginning of that year and could potentially be deforested in that year.

		2001	2002	2003	2004	2005	2006	2007	2008
	N	934,944	926,917	915,451	900,482	881,394	869,839	863,943	849,084
Elevation	mean	195.1	194.7	193.9	193.0	192.0	191.5	191.3	190.7
meters	(stdv)	(92.27)	(92.27)	(92.18)	(92.19)	(92.29)	(92.49)	(92.54)	(92.76)
Slope	mean	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4
degrees	(stdv)	(2.92)	(2.93)	(2.93)	(2.95)	(2.96)	(2.97)	(2.98)	(2.99)
Mean tree cover	mean	77.2	77.2	77.0	76.9	76.7	76.5	76.4	76.1
percent	(stdv)	(23.57)	(23.60)	(23.65)	(23.71)	(23.77)	(23.81)	(23.83)	(23.88)
Mean tree cover within 100m	mean	63.1	63.1	63.1	63.0	62.9	62.7	62.7	62.4
percent	(stdv)	(24.35)	(24.35)	(24.35)	(24.36)	(24.38)	(24.38)	(24.38)	(24.39)
St. dev. of tree cover within 100m	mean	28.6	28.5	28.5	28.4	28.4	28.4	28.4	28.3
value	(stdv)	(7.05)	(7.04)	(7.03)	(7.02)	(7.02)	(7.03)	(7.04)	(7.06)
Mean tree cover within 1km	mean	40.2	40.1	40.1	40.1	40.0	39.9	39.8	39.6
percent	(stdv)	(19.05)	(19.06)	(19.07)	(19.08)	(19.10)	(19.09)	(19.08)	(19.06)
St. dev. of tree cover within 1km	mean	37.2	37.2	37.1	37.1	37.0	36.9	36.9	36.8
value	(stdv)	(7.54)	(7.55)	(7.56)	(7.57)	(7.59)	(7.60)	(7.61)	(7.63)
NTL	mean	1.7	1.7	1.7	1.7	1.8	1.8	1.8	1.8
NTL value	(stdv)	(15.55)	(15.60)	(15.69)	(15.78)	(15.92)	(15.97)	(16.00)	(16.04)
Maximum NTL within 1km	mean	5.3	5.3	5.4	5.4	5.4	5.5	5.5	5.5
NTL value	(stdv)	(31.35)	(31.45)	(31.61)	(31.80)	(32.05)	(32.19)	(32.25)	(32.41)
Maximum NTL within 10km	mean	91.2	91.4	91.9	92.1	92.6	93.0	93.1	93.6
NTL value	(stdv)	######	######	######	######	######	######	######	######
Distance from a town	mean	30.9	30.8	30.8	30.8	30.8	30.8	30.8	30.8
5 km bins	(stdv)	(16.54)	(16.55)	(16.58)	(16.64)	(16.70)	(16.72)	(16.73)	(16.76)
Distance from a main road	mean	15.3	15.3	15.3	15.3	15.4	15.4	15.4	15.4
5 km bins	(stdv)	(10.82)	(10.83)	(10.84)	(10.87)	(10.89)	(10.91)	(10.91)	(10.91)

Table 5: Mean and standard deviation of land characteristics for all pixels that remain forested at the start of each year.

While there are gradual changes in the characteristics of forested pixels between 2001 and 2008, there is no evidence of a discontinuous shift between 2004 and 2005 when the Zero Deforestation Law came into effect. For example, the mean elevation of forested pixels fell from 195 meters in 2001 to 191 meters in 2008. This matches my findings that, all else held constant, higher elevation pixels are more likely to be deforested in Eastern Paraguay. Between 2004 and

2005, however, the mean elevation of forested pixels only fell by only one meter. This is a similar magnitude as the change in mean elevation of forested pixels between 2002 and 2003, or between 2003 and 2004.

Other land characteristics also display a gradual change over the period and do not exhibit an abrupt shift between 2004 and 2005. This indicates that land that was available to be deforested in 2004 did not differ significantly from land that was available to be deforested in 2005. Abrupt changes in post-deforestation land use patterns after the Zero Deforestation Law came into effect should not, therefore, be attributed to changes in the type of land still under forest cover and potentially available to be deforested. Instead, the evidence suggests that change in post-deforestation agricultural land use seen after the implementation of the Zero Deforestation Law can be attributed to the passage of this law.

All larger trends and events considered during this time therefore would either not be expected to impact deforestation, or would be expected to increase deforestation after the Zero Deforestation Law was passed. Therefore, I argue that any significant change in deforestation between the periods immediately before and after implementation can be largely attributed to this law and its enforcement.

Data and Methods

I use a linear probability model to investigate the impact of the Zero Deforestation Law on aggregate deforestation. The model tests whether there was a level shift in the rate of deforestation before and after implementation. The analysis uses the same random sample as the analysis of the drivers of deforestation. I select one percent of the pixels in Eastern Paraguay, and retain pixels that have forest cover of at least thirty percent in 2000, resulting in just under one million observations. These observations are duplicated to create a panel with one copy of each observation before and one copy after implementation. Pixels that lost forest cover in the pre period are removed from the post period data. This duplication enables me to analyze deforestation before and after the Zero Deforestation Law is implemented simultaneously in a single model. Standard errors are clustered at a half degree latitude longitude grid to control for spatial correlation in the error term.

The outcome of the linear probability model is a deforestation indicator equal to 1 if the pixel lost forest cover between 2001 and 2004 for the pre observations and equal to 1 if the pixel lost forest cover between 2005 and 2008 in the post observations.

The variable of interest is a post-policy indicator. The indicator equals one for all location observations in the post-period, both those that remain forested throughout the post period and those that lose cover during the post period. I estimate three specifications.

The first specification includes the post-policy indicator and does not control for any additional explanatory variables. This specification tests whether there was a significant change in the amount of deforestation in the pre-policy and post-policy periods.

$$D_{it} = post_t + \epsilon_{it}$$

The second specification includes the post-policy indicator and also controls for physical and anthropological land characteristics that have been identified as drivers of deforestation in the literature. All continuous explanatory variables are demeaned and normalized by their standard deviation. This specification investigates whether any changes in deforestation between the pre-and post-policy periods can still be attributed to the law's implementation once other factors have been accounted for.

$$D_{it} = post_t + \beta X_i + \epsilon_{it}$$

The third specification adds to the second by including interactions between the physical and anthropological characteristics and the post-policy indicator. The interaction terms reveal how the relationships between land characteristics and deforestation change after the Zero Deforestation Law is implemented.

$$D_{it} = post_t + \beta X_i + post_t * \beta X_i + \epsilon_{it}$$

Additional specifications are investigated and reported in appendix A4.

Results

The results of the linear probability model investigating the impact of the Zero Deforestation Law on aggregate deforestation in Eastern Paraguay are shown in Table 6. Full results, including the coefficients on soil groups and ecoregions, can be found in appendix A4.

	Dependent variable:		
	Deforested		
	(1)	(2)	(3)
post	-0.011^{***}	-0.010^{**}	-0.010^{***}
	(0.004)	(0.004)	(0.003)
elevation		0.024^{***}	0.034^{***}
		(0.004)	(0.004)
slope		-0.015^{***}	-0.017^{***}
		(0.002)	(0.002)
tree cover mean, 100m		0.003	-0.006^{***}
		(0.002)	(0.002)
tree cover stdv, 100m		0.012^{***}	0.019^{***}
		(0.002)	(0.002)
max NTL 10km		-0.003^{**}	-0.004^{***}
		(0.001)	(0.001)
distance to a town		−0.002	−0.001
		(0.003)	(0.003)
distance to a main road		0.0001	0.001
		(0.004)	(0.004)
post x elevation			-0.021^{***}
			(0.003)
post x slope			0.005^{***}
			(0.002)
post x tree cover mean, 100m			0.019^{***}
			(0.003)
post x tree cover stdv, 100m			-0.013^{***}
			(0.001)
post x max NTL 10km			0.003^{**}
			(0.001)
post x distance to a town			−0.001
			(0.003)
post x distance to a main road			−0.001
			(0.004)
Constant	0.058^{***}	0.074^{***}	0.074^{***}
	(0.006)	(0.008)	(0.008)
Ecoregions	No	Yes	Yes
Soil groups	No	Yes	Yes
SE clusters	1/2 degree	1/2 degree	1/2 degree
Observations	1,803,027	1,803,027	1,803,027
R^2	0.001	0.030	0.034
Adjusted R^2	0.001	0.030	0.034
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table 6: Linear probability models regressing deforestation on land characteristics and a post-policy indicator.

The negative and significant coefficient on the post-policy indicator in the first specification indicates that deforestation decreased after the Zero Deforestation Law was implemented. Not controlling for any other variables, the probability of deforestation fell by 0.011 percentage points between the pre- and post-policy periods. This matches the drop in the deforestation rate from around 6% of forested pixels lost in the four years before the policy was enacted, to around 5% of forested pixels lost in the four years after. This drop in deforestation significant, implying that it is larger than would be expected if it were due to random variation rather than to a change in deforestation after the Zero Deforestation Law was implemented.

The coefficient on the post-policy indicator remains consistent in the second specification, which controls for land characteristics. This indicates that the decrease in deforestation after the policy was enacted is not due to changes in the characteristics of land available to be deforested. That is to say, the drop in deforestation is not due to a sudden shortage of land with characteristics that make it suitable for agriculture. The relationships between land characteristics used as control variables and deforestation are consistent with findings in section III. Deforestation is more likely on higher elevation, flatter land with patchy tree cover that is not near dense population centers.

The third specification allows these relationships between land characteristics and deforestation to vary before and after the Zero Deforestation Law came into effect. This is done by including interactions between the post-policy indicator and each explanatory land characteristic with the two exceptions of soil group and ecoregion. This is because these two controls cover large continuous areas, and can be seen as fixed effects that offer insufficient variation with which to estimate meaningful coefficients. Interacted land characteristics are demeaned, allowing the coefficient on the post-policy indicator to be interpreted as-is despite the interaction terms. The coefficient on the post-policy indicator is consistent with the prior specifications, showing a 0.01 percentage point decrease in the probability of deforestation after the Zero Deforestation Law was implemented.

Coefficients on these interactions identify whether and how the relationship between each land characteristics and deforestation changed after the Zero Deforestation Law was implemented. Post-implementation, deforestation shifts towards relatively steeper and relatively lower lying land than was cleared before the policy. Before implementation, deforestation was 0.034 percentage points more likely to occur on land where elevation is one standard deviation above the mean, all else constant. After implementation, the impact of a one standard deviation increase in elevation is reduced by 0.021 percentage points, meaning that locations where elevation is one standard deviation higher 0.013 percentage points more likely to be deforested. This means the probability of deforestation drops from 13.5% to 11.3% for locations that have an elevation one standard deviation above the mean when added to a baseline average deforestation rate of 10.1% between 2001 and 2008. The results are similar for slope. Before implementation, a one standard deviation increase in slope meant a 0.017 percentage point decrease in the probability of deforestation. After implementation, this falls to a 0.012 percentage point decrease in the probability of deforestation. The shifts are not large enough to reverse the patterns, and deforestation remains more likely to occur on flatter, higher elevation areas even after the ban occurs. However, the magnitude of the relationship and therefore the power of the land characteristic to predict deforestation is weakened. This means that it became more difficult to predict where deforestation will occur based on these characteristics after implementation.

Deforestation also becomes more likely in areas that are more densely forested and less patchy. Deforestation remains more likely to occur in areas of patchy cover, but the relationship between this variable and the probability of deforestation is again weakened.

The relationship between nearby forest cover and deforestation reverses in the post-policy period. After the law is implemented, deforestation is more likely to occur in areas with high nearby forest cover, whereas before the policy deforestation was more likely to occur in areas with low nearby forest cover. This could indicate that new deforestation is hidden within existing forest, or that deforestation is moving into areas that were previously less desirable and therefore left uncleared. Deforestation also becomes more likely closer to bright nighttime lights values, but the relationship remains negative and the magnitude, when taken into account with the extreme skewedness of the variable, is nearly negligible. Distance from a town and distance from a road remain insignificant in predicting deforestation after the ban is put in place.

Results are generally robust to additional specifications which can be found in appendix A4. Additional specifications include varying the level of clustering for the standard errors, including alternate fixed effects, using a logit model, and removing deforestation in 2004 and 2005 from the panel. Most relationships remain consistent. The logit specification suggests that distance from towns and roads may be significant, however the predictive power of this relationship is low. Removing the years directly before and after implementation provide evidence that the magnitude of the decrease in deforestation may be lower than found on the main specification, but do not discredit the main results. Additional causal research will be needed to confirm how much of the amount of the decrease in deforestation detected between the pre- and post-policy period can truly be attributed to the Zero Deforestation Law.

Discussion

After the Zero Deforestation Law was implemented at the end of 2004, deforestation became illegal. Although it did not stop, the rate of deforestation decreased. This decrease cannot be attributed to other significant events that took place at the end of 2004, or in the years directly proceeding or following this time. In fact, many events that might have impacted deforestation during this time would be expected to increase deforestation after the ban, leading to a potential attenuation of the measured impact of the ban.

This decrease in deforestation also cannot be explained by a change in the characteristics of land available to be deforested in the pre- and post-policy period. There is no abrupt change in the characteristics of land available for deforestation between 2004 and 2005. In addition, regressions of deforestation on a post-policy indicator and land characteristics reveal a significant decrease in deforestation after the Zero Deforestation Law is implemented. This decrease in deforestation is robust to controlling for land characteristics, and also to including interactions between these land characteristics and the post-policy period. This second specification with interaction terms reveals that in addition to the rate of deforestation slowing after deforestation becomes illegal, the types of land that are deforested change. After the ban is enacted, land that is relatively lower-lying, steeper, less patchy, and nearer to bright nighttime lights is more likely to be cleared. The change is not enough to reverse the relationships between these variables and deforestation. However, the magnitudes and therefore the predictive power of these variables is decreased.

After the ban, the relationship between nearby mean forest cover and deforestation does reverse, with more deforestation occurring in areas with higher mean forest cover within 100 meters. This may indicate that the deforestation that occurs after the ban is implemented is conducted in areas hidden by existing forest cover. It may also indicate that, once deforestation is illegal, areas that were previously less desirable and left densely forested are now more attractive for clearing.

The evidence suggests that this command-and-control policy had a statistically significant impact on deforestation. However, deforestation continues. If Paraguay wants to further decrease deforestation, then additional policies, or additional means of implementing this policy, must be considered. These future approaches should take into account the ways in which deforestation patterns changed post-policy to proactively plan appropriate enforcement or policy actions.

V. Heterogeneous Policy Impacts Across Groups of Farmers

In the previous section I showed that deforestation in Eastern Paraguay decreased after the Zero Deforestation Law was implemented. However, the amount that deforestation fell may have varied across different land uses. In this section, I investigate whether the extent of the decrease in deforestation varied across three types of farming systems: small-scale subsistence-oriented farming, large-scale commodity-oriented farming, and ranching following the implementation of the Zero Deforestation Law.

The Zero Deforestation Law altered the decision-making process of whether to clear land or to leave it forested by making deforestation illegal. For some land managers the magnitude of the potential punishment, which could be 500 to 2000 daily wages, a common measure of the pay for a day of work, or 3 to 8 years' incarceration, in combination with the low enforcement rate were insufficient to deter deforestation when weighed against the benefits of clearing. For others, the possibility of punishment was enough to halt clearing. The decision varied depending on characteristics of the land parcel in question and characteristics of the land manager.

The distribution of decreases in deforestation across agricultural groups has important implications for environmental, equity, and economic impacts of the Zero Deforestation Law. The three post-deforestation land uses considered here are associated with different environmental outcomes, socioeconomic groups, and GDP-generating potential. Understanding the heterogenous impacts of a policy within its area of implementation is important to avoid unintended or unanticipated side effects.

It is difficult to analyze how small-scale farmers, large-scale farmers, and ranchers were impacted by the Zero Deforestation Law because the use of land in Eastern Paraguay after it is deforested is not systematically tracked. Therefore, to answer this question, I generate a dataset classifying deforested areas into post-deforestation agricultural use categories. I generate this data by combining the rich Global Forest Watch data on where deforestation occurs with widely available physical and anthropological attributes of these locations. The process has three basic steps. First, I manually label the post-deforestation land use for a sample of deforested locations as either small-scale agriculture, large-scale agriculture, rangeland, or other.

Second, I use this sample to train random forest models that classify the type of agriculture most likely to occur on deforested locations based on land characteristics. I expect that land that is more likely to be deforested due to a specific characteristic, such as having low slope, is also most likely to be deforested for a specific use. Flatter parcels may be more likely to be deforested, and also more likely to be deforested for large-scale mechanized soy production which cannot be conducted on steeply sloping parcels. Even land characteristics that do not significantly impact the probability of deforestation may significantly impact the type of agriculture practiced post-clearing. In my earlier analyses, I find no significant link between proximity to towns or roads and deforestation. However, parcels near towns may be more likely to be deforested for small-scale agriculture than for large-scale agriculture or for rangeland because small-scale producers typically cultivate fields within easy walking or driving distance from their homes. Meanwhile, locations far from population centers or far from roads may be more likely to be deforested for rangeland or large-scale agricultural because these require less frequent local labor.

Finally, I predict the post-deforestation agricultural uses for all locations deforested around the time the Zero Deforestation Law was implemented. I find that less area is deforested for large-scale agricultural use post-policy than is deforested for large-scale agricultural use pre-policy. This change includes a composition effect and a land use effect. Under the composition effect, different locations are selected for deforestation after deforestation becomes illegal. Some parcels that would have been cleared in a world where deforestation is allowed are no longer cleared when deforestation is illegal, and vice versa. In net, less area is cleared in the post-policy period than in the pre-policy period. Under the land use effect, the use of deforested land changes after the law came into effect. I find that land that is deforested post-policy is less likely to be used for large-scale agriculture and more likely to be used for small-scale agriculture or rangeland than that same parcel had it been deforested pre-policy. For rangeland and small-scale uses, the increase in deforestation from the composition effect makes up for the decrease in deforestation due to the land use effect. When the two effects are considered, total deforestation for rangeland and small-scale use remains nearly stable pre- and post-policy. Meanwhile both the land use effect and the composition effect lead to a decrease in deforestation for large-scale agriculture. This leads to an overall decrease in clearing for large-scale agriculture post-policy as compared to pre-policy.

This research contributes to the literature on the distribution of deforestation policy impacts across local populations. This is in contrast to a more traditional focus on a deforestation policy's efficiency in slowing aggregate deforestation, which I investigate in section IV. Heterogeneous policy impacts can take many forms. For example, heterogenous impacts can show up as different policy outcomes across groups, across space, or across both groups and space (Kazungu et al., 2021; Carvalho et al., 2017; Assuncao et al., 2017). In this analysis I focus on heterogeneous impacts across agricultural groups. The distribution of these impacts is important to understand because different outcomes may exacerbate any existing economic inequalities between groups (Nagel, 1999; Kovacic and Viteri, 2017). While implications of deforestation policies for specific groups of farmers are often discussed (i.e. Grabs et al., 2021), little empirical research focuses on this topic. This analysis contributes to this gap by showing that the Zero Deforestation Law primarily slowed deforestation for large-scale agricultural land use, but did not lead to meaningful changes in total deforestation for small-scale or rangeland uses. This finding provides valuable information for future work on the environmental, equity, and economic implications from the distribution of impacts of the Zero Deforestation Law.

This analysis also contributes to our understanding of drivers of deforestation. As mentioned in previous sections, a rich literature exists on the land characteristics that drive deforestation or that slow deforestation. In the first two analyses I investigate how these drivers relate to deforestation in Eastern Paraguay, both before and after the Zero Deforestation Law. In this section, I further this understanding of drivers of deforestation by showing that they can be used not only to predict deforestation, but also to predict the use for which land is deforested. While anthropological characteristics are not very useful in predicting deforestation in Eastern Paraguay, they are useful in determining the type of land use post-deforestation. Physical characteristics are significant predictors both of whether deforestation will occur, and what the post-deforestation land use will be. I find that measures of nearby characteristics are often better predictors of what the post-deforestation land use will be than measures of land characteristics at a location. Understanding what type of deforestation is most likely to occur in locations with different characteristics is valuable information for enforcement agencies attempting to manage deforestation.

Finally, this analysis contributes a method to classify post-deforestation land use that can be replicated in other data-scarce regions. Paraguay, like many less developed regions, lacks data on post-deforestation land use. The three-step method outlined in this analysis relies on manually generated training data and globally available, open-source satellite-derived data to classify land use. This method can be used in other regions where there is limited information on the use of land. This method vastly expands the regions for which this type of analysis can be undertaken, and can help to broaden the focus of existing deforestation research.

Generating Post-Deforestation Land Use Training Data

The first step I take to generate data on agricultural land uses after deforestation is to generate labeled training data. I use stratified random sampling to select a sample of pixels from the set of all pixels in Eastern Paraguay that have at least thirty percent forest cover in 2000 and that lost forest cover between 2001 and 2010. Sampling is stratified by year of loss, with the number of observations in each year proportional to the total deforestation between 2001 and 2010 that took place in that year.

I use Google Earth Pro images to manually categorize the post-deforestation agricultural land use for each sampled location as small-scale, large-scale, rangeland or other. Categorization is based on a few key characteristics of the satellite images. The first is clearing size. Clearing size is the area, in hectares, of the pasture, field, or other cleared area that contains the sampled deforested location, as delineated by visible field boundaries or boundaries between forest and cleared areas. The clearing size is assessed using built-in measurement tools. The second key characteristic is signs of mechanized agriculture. This appears in the satellite images as uniform areas with linear vegetation patterns left behind by tractors or other farm equipment. The final key characteristic is settlement patterns. In Eastern Paraguay settlement patterns most commonly display as houses along a road with property boundaries extending outwards from the residences. Within these property boundaries are numerous small fields and pasture areas. In some cases, settlement patterns are circular with fields radiating from a central group of houses, or less structured with houses and small fields interspersed in a less planned manner.

Recently deforested pixels that are in a clearing of less than five hectares and located within settlement patterns are categorized as having been cleared for small-scale agriculture. Recently deforested pixels that are in a clearing larger than five hectares, display patterns of mechanized agriculture, and are not located within a settlement pattern are categorized as having been cleared for large-scale agriculture. Recently deforested pixels that are in a clearing greater than five hectares, do not show evidence of mechanized agriculture, and are not located within a settlement pattern are categorized as rangeland. Figure 3 shows examples of recently deforested locations that I sorted into small-scale agriculture, large-scale agriculture, and rangeland, respectively.

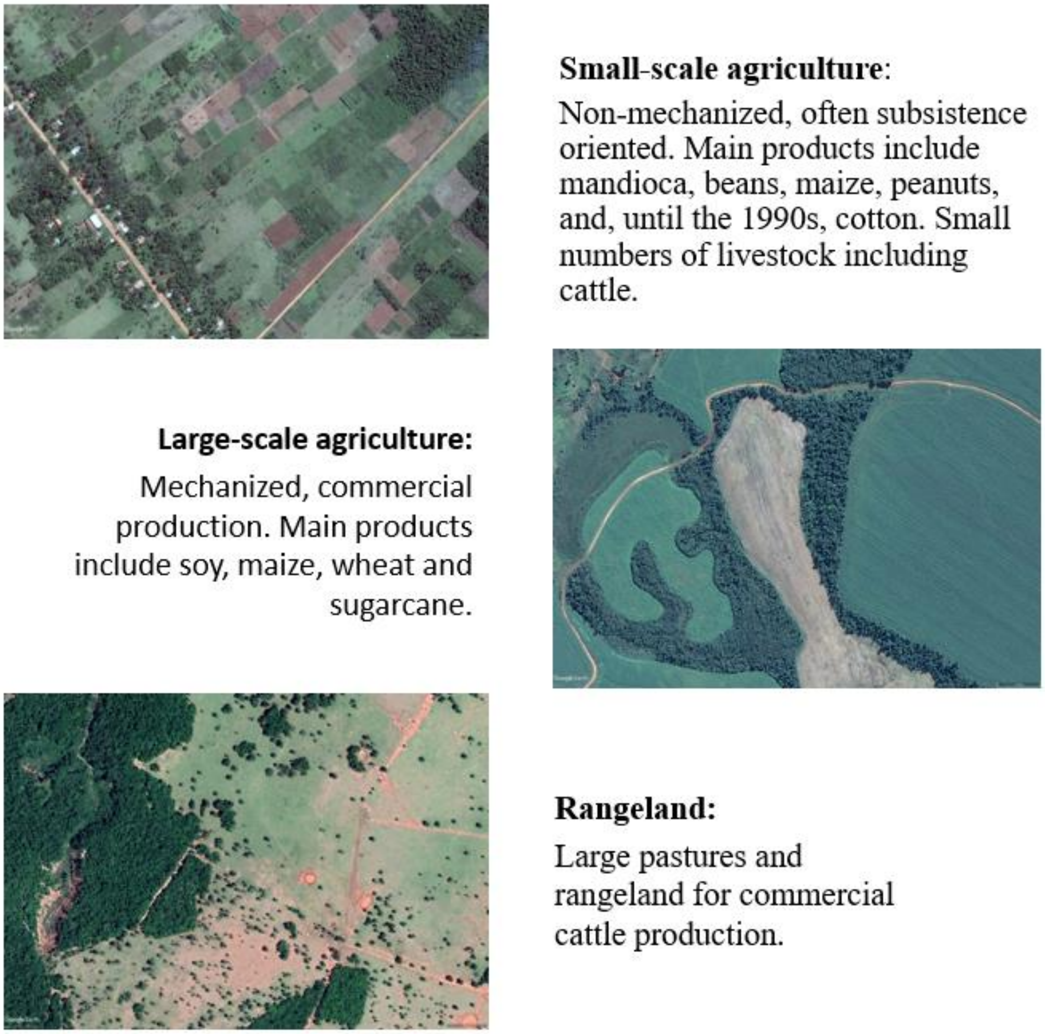

Figure 11: Examples of each post-deforestation agricultural land use category in training location satellite imagery from Google Earth Pro.

Recently cleared pixels that do not fit into one of these categories are classified individually based on information from the satellite image. For example, a 10-hectare mechanized field that is adjacent to a settlement would be classified as large-scale agriculture despite the proximity to houses that is typical of small-scale clearings. Deforestation for non-agricultural purposes, such as deforestation for a road or for urban development, are classified as other. The final training sample includes 136 observations of small-scale agriculture, 145 observations of large-scale agriculture, and 191 observations of rangeland. 29 observations, or 6% of the sample, fall into the other category. The distribution of post-deforestation agricultural land uses in the training data by land use category and year can be found in appendix A5.

Classifying Agricultural Land Use

I use the sample of data labeled with post-deforestation use to train random forest classification models that predict whether a pixel is most likely to be used for small-scale agriculture, large-scale agriculture, rangeland, or other uses after being cleared. Random forest is a decision tree-based machine learning algorithm that generates predictions by averaging the predicted outcome across many decision trees. The prediction is the class that is chosen most frequently across all trees. The number of trees is set to 500, and at each branch a split is made based on one of four characteristics randomly selected from the list of physical and anthropological variables.

These characteristics on which the predictions are based are physical or anthropological land characteristics that have been identified as drivers of deforestation in the literature. I expect that a characteristic that makes land more likely to be deforested also makes land more likely to be deforested for a specific reason. For example, I found in previous sections that land with a lower slope is more likely to be deforested than land with a higher slope. It may also be more likely to be deforested for large-scale, mechanized soy production since this production practice cannot be conducted in steeply sloped areas. Meanwhile, steeper areas may be more likely to be deforested for small-scale agriculture, which can more easily alter fields to deal with higher grades. Similar patterns may exist with respect to other characteristics, including those that are not associated with more or less deforestation. For example, my analyses of the drivers of deforestation found that distance from a town has no significant relationship with deforestation. However, I expect that land closer to a town will be more likely to be deforested for small-scale agriculture practiced by residents of that town. Land that is further from a town may be more likely to be deforested for large-scale agriculture or rangeland uses, which do not require local labor.

The training data for which I generate labels is representative of the full set of deforested locations, which is the set of pixels for which I will generate agricultural land use predictions. Table 7 summarizes the means and standard deviations of land characteristics that are used as predictive variables. The first three columns summarize land characteristics for pixels in the training data, all together and separated into pixels deforested pre- and post-policy. The final three columns summarize land characteristics for all pixels deforested in Eastern Paraguay during the same period, again all together and separated into pixels deforested pre- and post-policy. The means and standard deviations are similar between the training and full datasets, indicating that the training data represent the full dataset well.

Variable	unit		training 2001-2010 N = 501	training pre 2001-2004 N = 255	training post 2005-2010 N = 246	all 2001-2010 N = 10,500,507	all pre 2001-2004 N = 5,335,085	all post 2005-2010 N = 5,165,422
Elevation	meters	mean	238.435	252.949	223.39	235.0	247.5	222.1
		(stdv)	(82.62)	(84.28)	(78.24)	(74.91)	(73.98)	(73.66)
Slope	degrees	mean	3.157	3.199	3.113	3.1	3.2	3.0
		(stdv)	(2.10)	(2.23)	(1.97)	(2.11)	(2.13)	(2.09)
Soil group	categorical	mean	.	.	.	.	.	.
		(stdv)	.	.	.	.	.	.
Ecoregion	categorical	mean	.	.	.	.	.	.
		(stdv)	.	.	.	.	.	.
Mean tree cover	percent	mean	87.519	86.588	88.484	87.8	86.8	88.9
		(stdv)	(17.68)	(18.10)	(17.23)	(17.04)	(17.28)	(16.73)
Mean tree cover within 100m	percent	mean	69.155	65.945	72.483	69.7	66.7	72.8
		(stdv)	(23.11)	(23.95)	(21.75)	(22.99)	(23.60)	(21.91)
St. dev. of tree cover within 100m	value	mean	30.691	31.333	30.026	30.5	31.3	29.6
		(stdv)	(7.18)	(7.60)	(6.68)	(6.52)	(6.88)	(6.00)
Mean tree cover within 1km	percent	mean	44.757	40.855	48.801	45.2	42.9	47.7
		(stdv)	(19.11)	(18.89)	(18.53)	(18.15)	(17.90)	(18.09)
St. dev. of tree cover within 1km	value	mean	39.956	39.456	40.475	40.5	40.4	40.5
		(stdv)	(6.16)	(6.35)	(5.92)	(5.72)	(5.87)	(5.56)
NTL	NTL value	mean	1.312	0.481	2.174	1.0	0.9	1.2
		(stdv)	(16.28)	(1.45)	(23.18)	(10.21)	(7.75)	(12.24)
Maximum NTL within 1km	NTL value	mean	3.52	2.468	4.611	3.8	3.4	4.2
		(stdv)	(21.77)	(7.16)	(30.19)	(19.71)	(16.45)	(22.57)
Maximum NTL within 10km	NTL value	mean	59.517	66.652	52.12	67.8	66.6	69.0
		(stdv)	(106.78)	(125.31)	(82.96)	(124.40)	(119.80)	(128.97)
Distance from a town	5 km bins	mean	31.537	30.98	32.114	31.4	31.3	31.5
		(stdv)	(14.19)	(13.41)	(14.96)	(14.28)	(13.67)	(14.88)
Distance from a main road	5 km bins	mean	15.729	15.569	15.894	15.1	15.1	15.1
		(stdv)	(10.77)	(10.41)	(11.15)	(9.99)	(9.59)	(10.40)
Protected when deforested	indicator	mean	0.06	0.02	0.09	0.06	0.05	0.07
		(stdv)	(0.24)	(0.17)	(0.30)	(0.24)	(0.21)	(0.26)

Table 7: Summary statistics of explanatory variable values for all pixels deforested in Eastern Paraguay, and for pixel observations in the training dataset for the random forest models.

I train two random forest models. A pre model is trained using observations that were deforested before the Zero Deforestation Law came into effect (before 2005) and reflects patterns in deforestation that dominated when deforestation was legal. A *post* model is trained using observations that were deforested after the Zero Deforestation Law came into effect (after 2004) and reflects patterns in deforestation that dominated when deforestation was illegal. Statistics on the models' accuracies are in appendix A6.

Some land characteristics are more influential in predicting the post-deforestation agricultural land use than others. This is reflected in the importance metric, which measures the mean decrease in the Gini coefficient of the model's classification that can be attributed to that variable. A variable with a higher importance value is more informative for splitting observations into the correct categories than a variable with a lower importance value. The importance of explanatory variables for predicting post-deforestation land use for the pre and post models ordered by importance is shown in Figure 12.

The order of importance is similar between the pre and post models. In both models, variables related to human activity are ranked higher in importance. Influential human activity-related variables include the maximum nighttime lights values within one and ten kilometers, the distance to the nearest town, and the distance to a main road. Elevation and soil group are influential predictors as well, especially in the pre model. Another pattern that emerges is that statistics of nearby areas are more influential than values at the pixel location. This can be seen in the low importance of measures of nighttime lights and measures of tree cover at the pixel location.

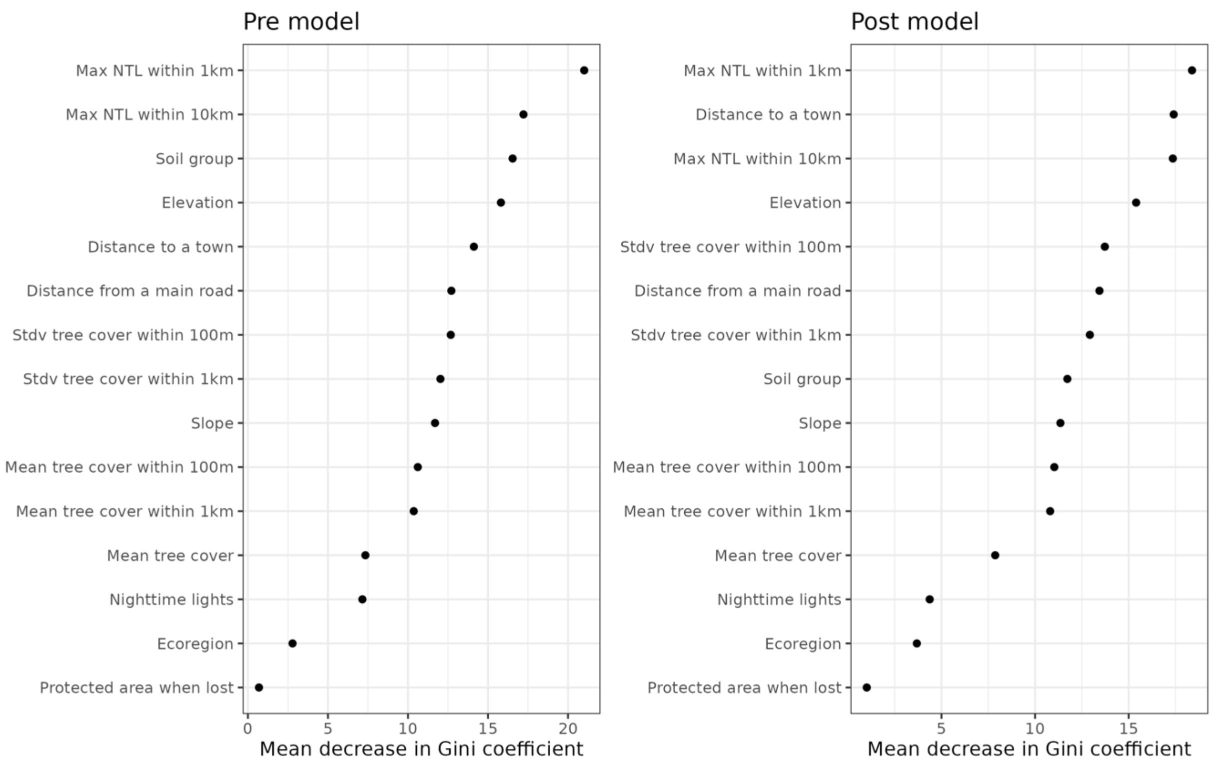

Figure 12: Variable importance plots for the pre and post random forest models that predict post-deforestation agricultural land use.

The output of the model is predicted agricultural land use. I generate a land use prediction for every pixel that lost forest cover in Eastern Paraguay in the four years before and after the Zero Deforestation Law's implementation. Predictions generated using only subsets of explanatory variables in the final model can be mapped to visualize how the model builds complexity with the addition of relevant variables. This is illustrated in Figure 13. Panels A through C predict post-deforestation land use with pre-policy deforestation observations, using models trained on progressively more explanatory variables. Panel D predicts post-deforestation land use with post-policy deforestation observations. The area lost is amplified for visibility, and therefore greatly overestimates the amount of deforestation. The actual amount of deforestation that occurred between 2001 and 2008 is mapped in Figure 2.

The predictions in Panel A are generated based only on elevation. In general, the western side of Eastern Paraguay has lower elevation, and the eastern side of Eastern Paraguay has higher elevation. This as can be seen in the map of elevation in Figure 3. The results of this simple prediction show that, not controlling for any other variables, low-lying land is more likely to be used for small-scale agriculture after deforestation and higher elevation land is more likely to be used for large-scale agriculture. However, I expect that other variables are also important predictors of deforestation and should be accounted for in the prediction.

Panel B predicts the land use category using all of the physical characteristics from the full model: elevation, slope, soil group, and ecoregion. The distribution of these land characteristics are mapped in figures 3 through 6. The pattern of small-scale agriculture in low-lying areas and large-scale agriculture at higher elevations is still evident but it is less pronounced. Clear divisions now emerge between the land use classes along soil group boundaries. For example, on the western side of the region a distinct break appears between acrisols, which are predominantly predicted to go to large-scale agriculture, and leptosols, which have a more even mix of predicted agricultural classes.

Panel C illustrates predictions from the full pre model, which incorporates all physical and anthropological explanatory variables summarized in Table 7. Maps of the data underlying these anthropological variables can be found in figures 7 through 9. Patterns revealed using only physical characteristics remain. However, the predicted boundaries in land use classes are less exact once human patterns are considered. In addition, areas near towns are roads become much more likely to be small-scale agriculture once these characteristics have been added to the model. For example, land surrounding the capital city of Asuncion, around the midpoint of the western border, is almost entirely predicted to go to small-scale agriculture post-deforestation. In addition, there is a linear pattern of small-scale predictions along the highway connecting Asuncion with Cuidad del Este on the eastern border. Panel D again uses all explanatory variables in the prediction, but this time the model is trained using deforestation that took place post-policy, rather than pre-policy. Similar patterns are evident overall, including less large-scale agriculture predicted at lower elevations, different classes dominating predictions in different soil groups, and a dominance of small-scale predictions surrounding population centers and important roads. However, there are minor but visible differences between the post-policy prediction in panel D and the pre-policy prediction in panel C. For example, in the post-policy prediction the dominance of large-scale agriculture in the acrisols soil group is diminished compared to the pre-policy prediction.

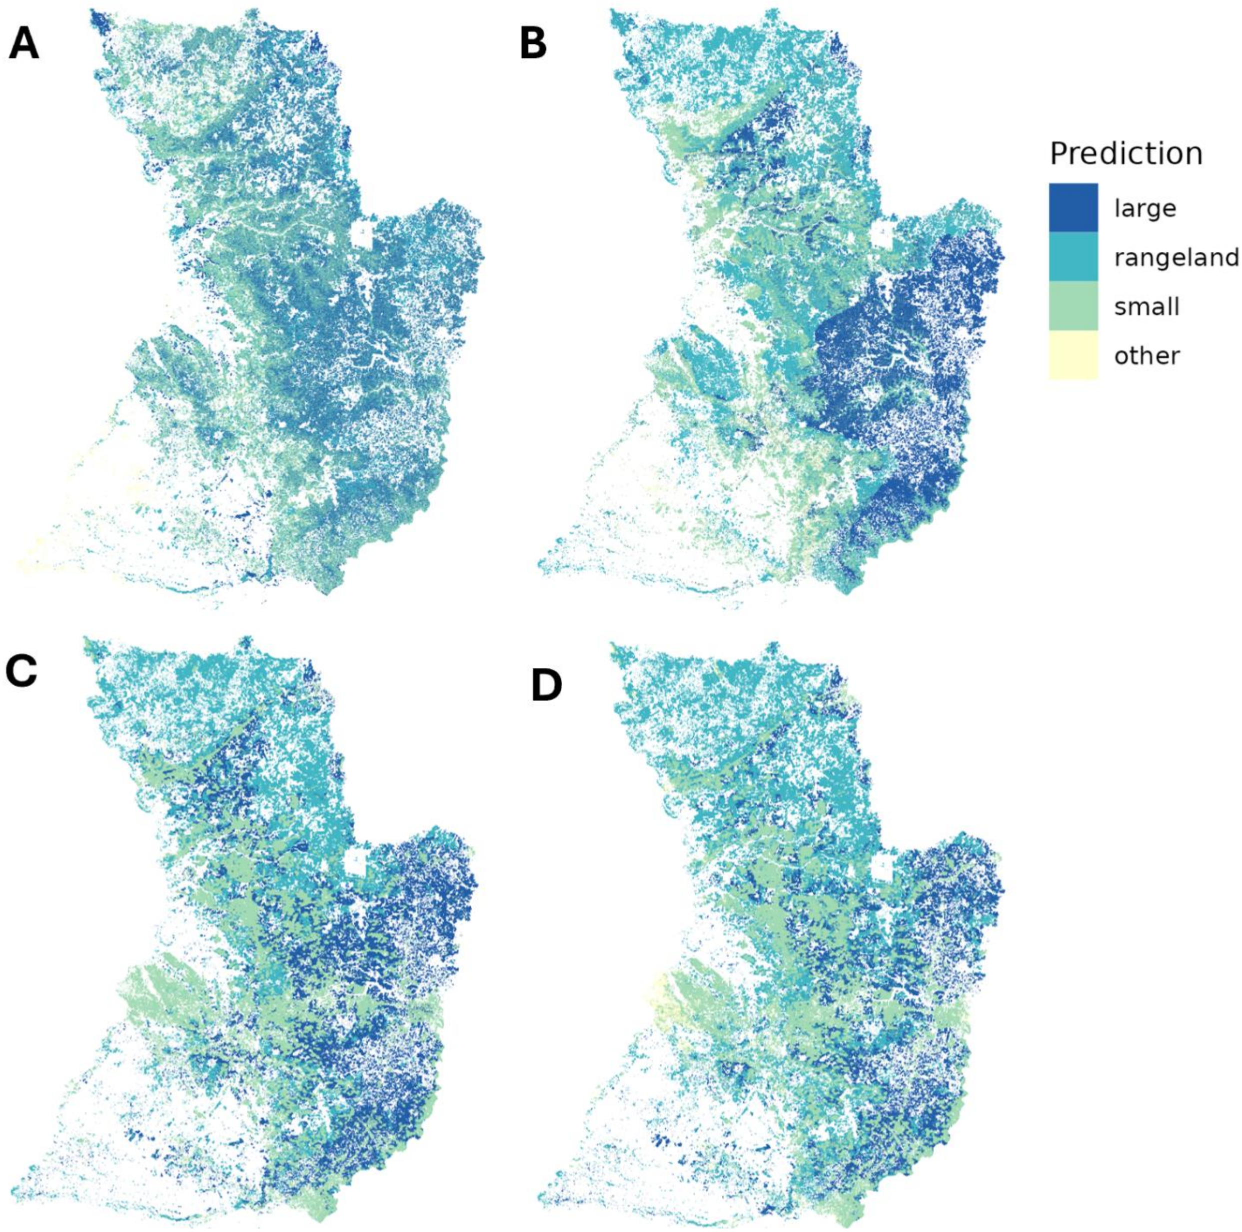

Figure 13: Maps of post-deforestation land use predictions illustrate how predicted lan use distributions are influenced by the addition of explanatory variables. Deforested area is magnified for visibility. Panel A predicts pre-policy land use using only elevation. Panel B predicts pre-policy land use using all physical land characteristics. Panel C predicts pre-policy land use using all physical and anthropological characteristics. Panel D predicts post-policy land use using all physical and anthropological characteristics.

Partial plots of the relationship between each variable and the agricultural land use classes reveal how the probability of predicting each class changes as the value of a land characteristic varies, holding all else constant. Because they are run separately, magnitudes should not be compared between the pre and post models. Some interesting patterns emerge. The correlation between elevation and land use noted in the maps is again evident in the partial plot for elevation, in Figure 14. Small-scale use is more likely to be predicted at moderately low elevations, and large-scale use is more likely to be predicted at higher elevations. There is a weaker predictive pattern for rangeland.

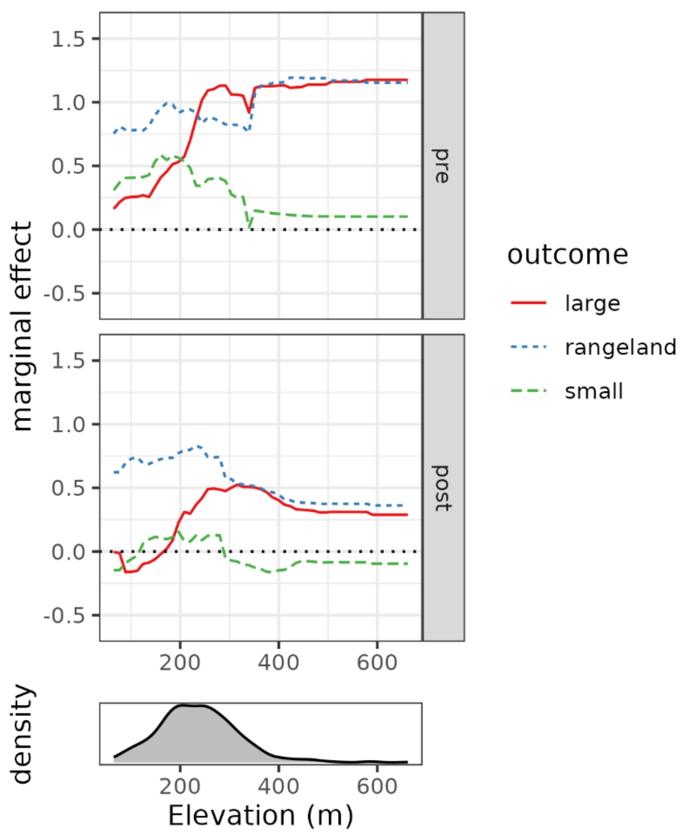

Figure 14: Partial plots for elevation for large-scale, rangeland, and small-scale post-deforestation uses. Partial plots of the relationship between each variable and the agricultural

land use classes reveal how the probability of predicting each class changes as the value of a land characteristic varies, holding all else constant.

The maximum nighttime lights values within one and within ten kilometers had among the highest importance values of all explanatory variables. The partial plots for these two variables are shown in Figure 15. The plots do not include the full range of nighttime lights values to emphasize variation at lower values where the training data is denser.

The patterns track reasonable expectations about the locations of the different types of agriculture. In Eastern Paraguay towns and settlements have nighttime lights values ranging from the double to triple digits. Pixels within one kilometer of towns or settlements are much more likely to be converted to small-scale agriculture than to rangeland or large-scale use after deforestation, all else constant. It is intuitive that pixels in the direct vicinity of settlements will be used by residents of these settlements for small-scale agriculture. Having a town or settlement within ten kilometers, however, does not bias the prediction towards small-scale use. Meanwhile pixels near no or very low nighttime lights, indicating no or very low economic activity, are less likely to be small-scale and more likely to be rangeland, all else held constant. This matches reasonable expectations since rangelands do not require heavy human capital inputs and therefore can be located further from population centers. Rangeland is a viable use of land with lower soil fertility, meaning rangeland can be more easily established in areas that agricultural populations did not prioritize. Partial plots for the remaining variables can be found in appendix A7.

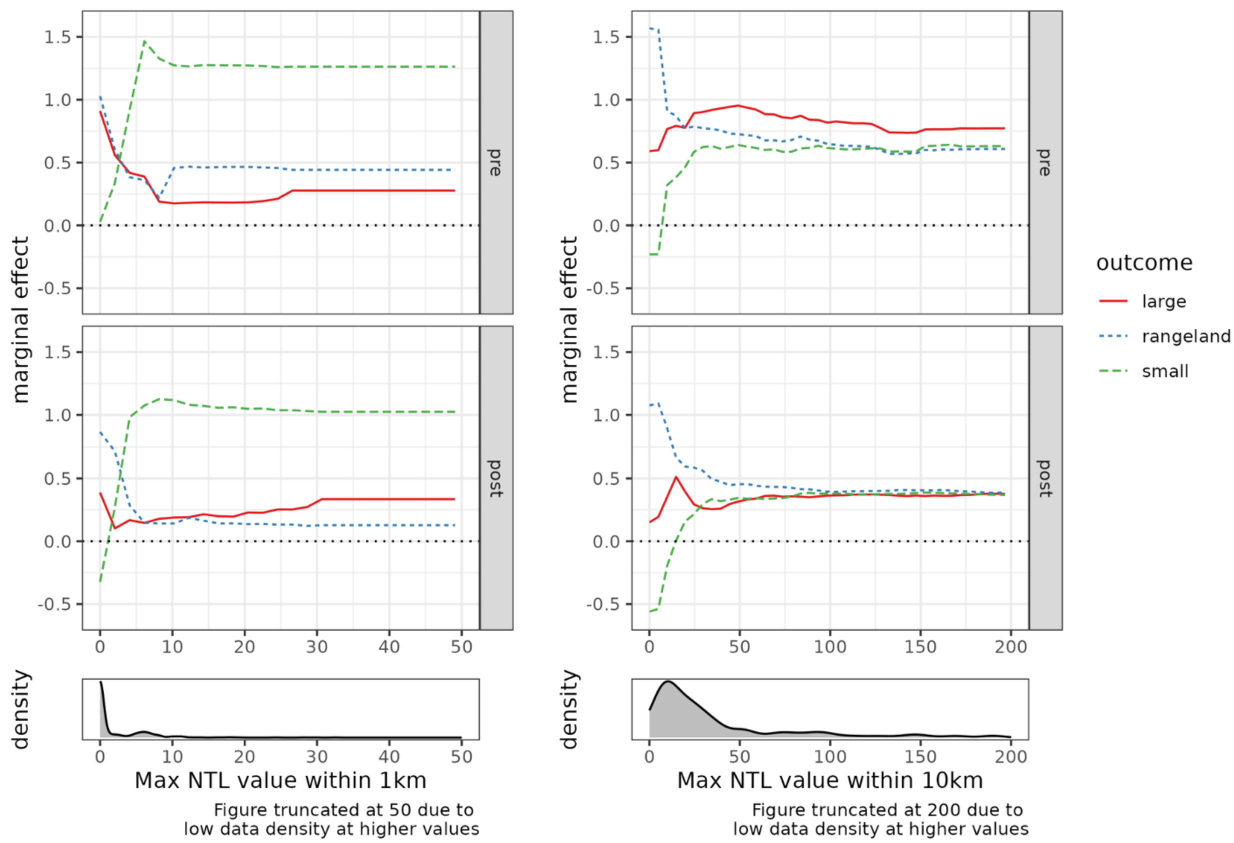

Figure 15: Partial plots for the maximum nighttime lights value within 1 kilometer and the maximum nighttime lights value within 10 kilometers for large-scale, rangeland, and small-scale post-deforestation uses. Partial plots of the relationship between each variable and the agricultural land use classes reveal how the probability of predicting each class changes as the value of a land characteristic varies, holding all else constant.

Post-Deforestation Agricultural Land Use

The goal of this section is to analyze changes in the agricultural use of newly cleared land before and after the Zero Deforestation Law was implemented. Data on post-deforestation land use with which to do this do not exist, but the models generated above can be used to predict this data.

This prediction is the final step in this three-step process to generate post-deforestation land use data.

Predictions are generated for all pixels that lost forest cover in Eastern Paraguay in the four years before and after implementation using the pre and post models. The best prediction for a pixel deforested before the Zero Deforestation Law was implemented is made with the pre model. The best prediction for a pixel deforested after the Zero Deforestation Law was implemented is made with the post model. I use "best" to describe these predictions to differentiate between these predictions and other possible predictions discussed later.

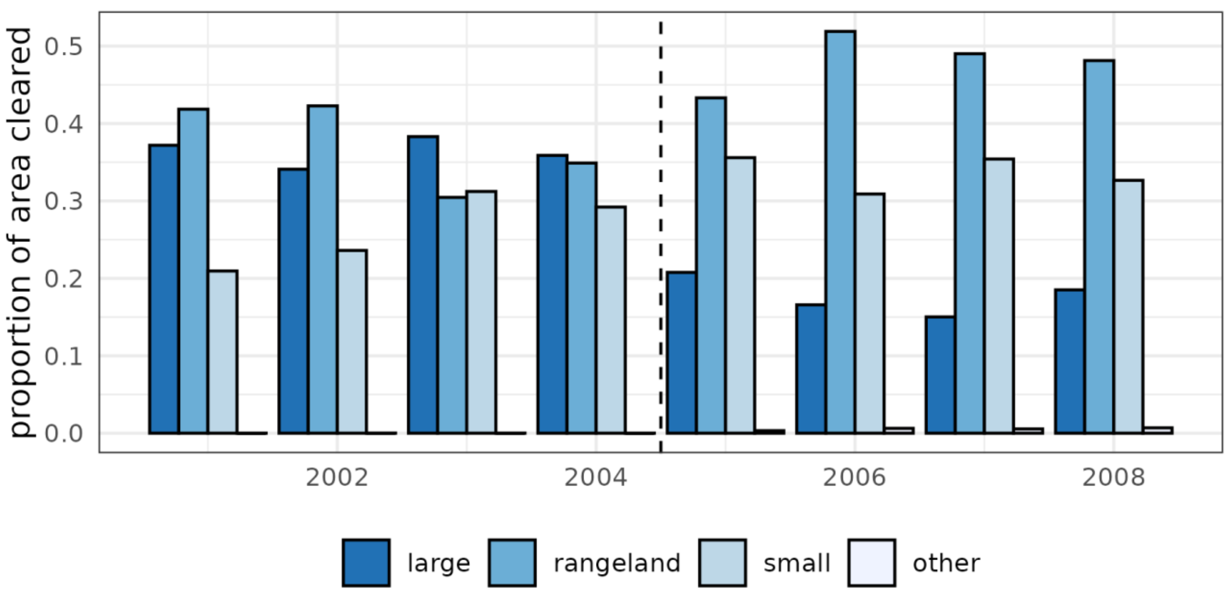

Figure 16: The proportion of the area deforested in each year that went into large-scale, rangeland, or small-scale agricultural use, or other uses. Deforestation is for the calendar year, from January 1st to December 31st of the stated year.

Figure 16 summarizes the proportion of the area that was deforested in each year by land use, using the best predictions. Each year includes the deforestation that occurred during that calendar year, from January 1st and December 31^{st}. Implementation of the Zero Deforestation Law is indicated by the dashed line between 2004 and 2005. The results reveal a drop in the proportion of deforested area going into large-scale agricultural use after the Zero Deforestation Law is passed, and a corresponding increase in the proportion of deforested area going into rangeland and small-scale uses. A negligible proportion of the total area cleared is predicted to go to other uses in both the pre- and post-policy periods.

Figure 17 plots the magnitude of the area that was deforested in each year by land use. The best predictions are shown in solid lines. Variation in the total amount of deforestation between years is due to variation in the total deforestation that occurred in each year, rather than due to model outcomes. A decrease in the area going to large-scale use after the Zero Deforestation Law is implemented is again evident when considering the magnitude of deforestation predicted for each land use, rather than the proportion. Addition years of predictions are available in appendix A8, though it should be noted that predictions become less reliable as time from the training data increases.

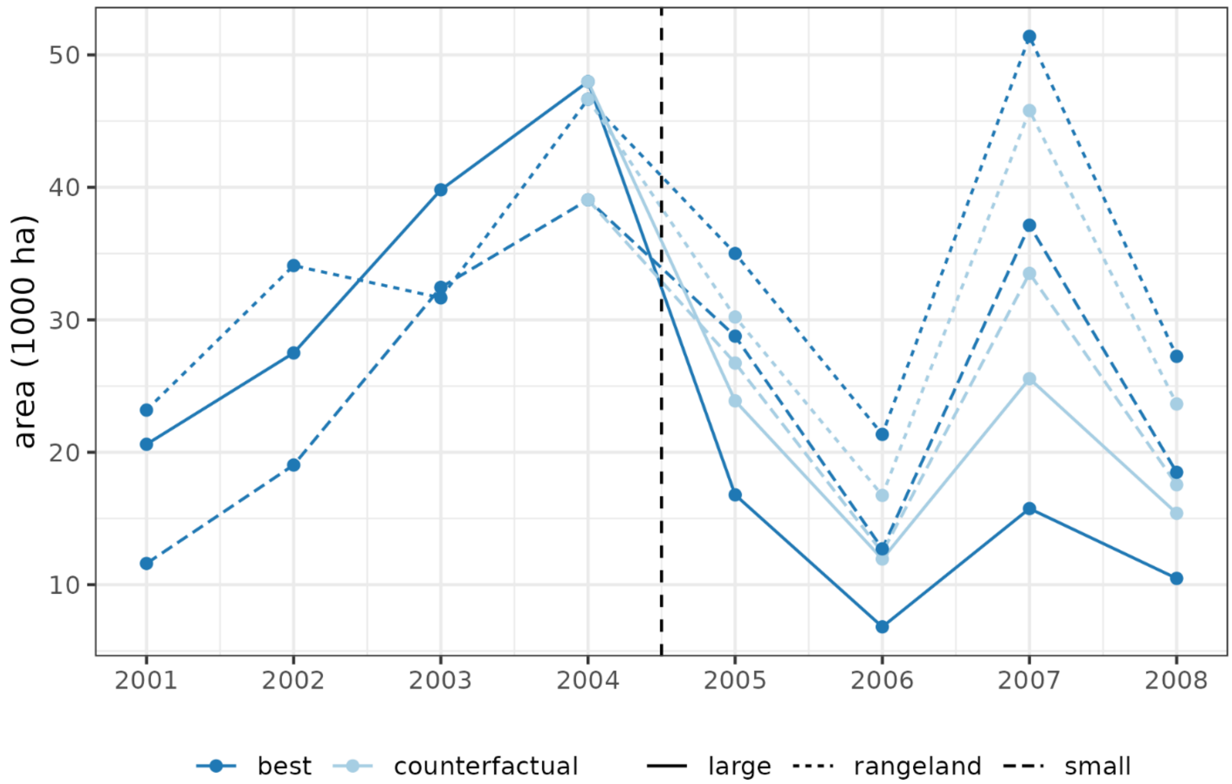

Figure 17: Agricultural land use predictions after deforestation for pixels deforested between 2001 and 2019. Solid lines are best predictions generated using the model corresponding to the deforestation year, the pre model until 2004 and the post model after 2004. The dashed lines are counterfactual land use predictions generated using the alternative model.

In addition to the best predictions, counterfactual land use predictions can be made for deforested pixels. These counterfactual predictions are made on pixels that are actually deforested. They do not account for changes in the location or amount of deforestation between the pre and post policy periods. Rather, the counterfactual predicts the land use that a pixel would be put to if it had been deforested under the counterfactual policy scenario. These counterfactual predictions are made by predicting land use for deforested pixels with the model that does not match the period in which deforestation occurred. For pixels deforested after the Zero Deforestation Law was passed, the counterfactual prediction shows the use that the pixel would have gone to if it had been deforested before the law was passed, when deforestation was legal. For pixels deforested before the law was passed, the counterfactual prediction shows the use that the pixel would have gone to if it were deforested post-policy implementation, when deforestation is illegal.

This sheds light on how the decrease in deforestation for large-scale agriculture revealed by the best predictions was achieved. The counterfactual predictions for the post-policy period are shown in dashed lines in Figure 17. The counterfactual predictions for the pre-policy period are not shown in Figure 17 in order have a cleaner figure for the discussion. The mean values, however, are provided in Table 8. Table 8 provides, for both the best predictions and the counterfactual predictions, the magnitude and the average percent of total deforestation by land use class for the four years before and the four years after the Zero Deforestation Law.

Predicted Area for Each Land Use

	2001-2004 pre model		2005-2008 post model	
large	36%	135,893 ha	18%	49,834 ha
rangeland	36%	135,594 ha	48%	134,998 ha
small	27%	102,153 ha	32%	90,228 ha
other	0%	39 ha	1%	1,488 ha
total	100%	373,678 ha	100%	283,419 ha

Counterfactual Predicted Area for Each Land Use

	2001-2004 post model		2005-2008 pre model	
large	24%	89,896 ha	27%	76,772 ha
rangeland	46%	170,268 ha	41%	116,396 ha
small	30%	112,594 ha	32%	90,228 ha
other	0%	921 ha	0%	23 ha
total	100%	373,678 ha	100%	283,419 ha

Table 8: The average percent of total deforested area and the hectares of deforested area predicted go into each land use post-deforestation are summarized for the four years before and the four years after the Zero Deforestation Law was implemented. The predicted area, in the top portion, is the best prediction. This uses the pre model in the pre-policy period and the post model in the post-policy period. The counterfactual predicted area below is the counterfactual land use prediction, which uses the post model in the pre-policy period and the pre model in the post-policy period.

Investigating the best and the counterfactual predictions together reveals that changes in the area deforested for each agricultural land are due to both different locations being cleared, and to the same locations being put to different uses post-deforestation. I refer to these as a composition effect and a land use effect.

Under the composition effect, the locations that are deforested change when the policy is implemented. Areas that would have been deforested in the post-policy period if the Zero Deforestation Law never came into effect may no longer be deforested once the law is implemented, and vice versa.

The composition effect is most clearly visible in Figure 17 when comparing the best prediction for large-scale agricultural use in the pre period with the counterfactual prediction for large-scale agricultural use in the post period. Both predictions are generated using the pre model for pixels that were deforested in the Global Forest Watch data. If there was no change in the locations or amount of deforestation after the law was implemented, I would expect the predictions to remain relatively stable across periods since they are generated with the same model. Instead, there is a notable decrease between the pre and post periods. In total, 136 thousand hectares were deforested for large-scale use in the pre-policy period. If the pre-policy use patterns had continued for pixels deforested in the post period, only 77 thousand hectares would have gone into large-scale agricultural use post-policy. This suggests that some land that would have gone into large-scale agriculture based on pre-policy patterns was not cleared at all in the post-policy period, and is therefore not in the data on which these counterfactual predictions are generated. The composition effect can also be seen in the proportion of total deforestation that is predicted to be used for large-scale agriculture, rather than the magnitude. 36% of total deforestation went to large-scale use in the pre-policy period. If the pre-policy use patterns had continued for pixels cleared post-policy, only 27% of deforested land would have gone into large-scale use. This drop is due to the composition effect shifting where deforestation occurs. While the composition effect is most evident for large-scale agriculture, all three agricultural classes were impacted. 136 thousand hectares and 102 thousand hectares were deforested for rangeland and small-scale use, respectively, in the pre-policy period. If the pre-policy use patterns had continued for pixels cleared in the post-policy period, only 116 thousand hectares and 90 thousand hectares would have gone into rangeland and small-scale use, respectively. In the post-policy period, land that would have been used for each land use class, but especially land that would have gone into large-scale agriculture in the pre-policy period is less likely to be deforested at all.

The same deforested locations are also sometimes put to a different use after the policy was implemented than they would have been before the policy. This is the land use effect. For example, land that would have gone into large-scale agricultural use may be used for rangeland instead after the Zero Deforestation Law is implemented.

The land use effect is visible when comparing counterfactual predictions and best predictions in the post-policy period in Figure 17. Both predictions are made on the same set of pixels, all pixels deforested in the post-policy period. A decrease is visible in the deforested area going to large-scale agriculture between the counterfactual prediction and the best prediction. Under the counterfactual prediction, 77 thousand hectares or 27% of the total deforested area go to large-scale agriculture between 2005 and 2008. Under the best prediction, this drops to 50 thousand hectares or 18% of the total deforested area. The land that no longer goes to large-scale agriculture primarily goes into rangeland use. Under the counterfactual prediction 116 thousand hectares or 41% of the total deforested area go to rangeland between 2005 and 2008. Under the best prediction, this increases to 135 thousand hectares or 48% of the total deforested area. Model confusion matrices, which provide estimates for correct and incorrect land use classifications by class and are provided in appendix A6, suggest that this movement from large-scale production to rangeland post-policy may even be larger. There is no change in the total area predicted to go into small-scale use using the pre or post model for 2005 through 2008. This does not exclude a land use effect, as a similar area may have switched out of and switched into being deforested for small-scale use after the Zero Deforestation Law was passed. Overall, it is evident under the land use effect that in some cases the same locations are used for different purposes in the pre- and post-policy periods.

When both the composition and land use effects are considered, there is nearly no change in the amount of land cleared for rangeland in the four years before and after the Zero Deforestation Law was implemented. Using the best predictions, deforestation for rangeland use changes by less than one thousand hectares between the pre and post periods. Deforestation for small-scale use drops by around 11 thousand hectares. Meanwhile, deforestation for large-scale agriculture experienced the biggest impact from the law. This decreased by 63%, from 136 thousand hectares between 2001 and 2004 to 50 thousand hectares between 2005 and 2008. Deforestation for small-scale agriculture decreased marginally, from 102 thousand hectares between 2001 and 2004 to 90 thousand hectares between 2005 and 2008, a decrease of only 12%.

The random forest model uses a winner-take-all approach when predicting the outcome. The predicted post-deforestation land use is the land use with the highest probability for that pixel, regardless of whether this probability is 99% or 50%. Another approach is to examine the average probability, across all pixels, of each land use class in each period to see how this average probability changes. A decrease in the average probability of a land class indicates that this land use was less likely to be predicted across pixels, regardless of whether this decrease in probability was enough to alter the winner-take-all algorithm.

In the pre-policy deforestation best predictions, there is a 34% probability, on average, that a pixel went to large-scale agriculture after being cleared. In the post-policy deforestation best predictions this probability, on average, falls to 25%. Meanwhile the mean probability of a small-scale prediction does not change much, from 27% before to 29% after. The mean probability of a rangeland prediction behaves similarly, increasing slightly from 34% before to 37% after. This supports the findings above that the Zero Deforestation Law primarily decreased deforestation for large-scale agriculture.

Causality of Predictions

Patterns in the total area predicted for each agricultural land use class reveal how deforestation for each class changed following the implementation of the Zero Deforestation Law. Attributing the changes to the law assumes that these changes are not explained by other events that took place simultaneously or in the years immediately surrounding its implementation. A discussion of events that could have impacted deforestation and land use can be found in the causality discussion in section IV. These include commodity prices, unemployment rates, policy changes, election cycles, and changes in land availability. I argue that there are few other events with significance to deforestation that occurred during the years included in my study, and the events that did take place acted in the opposite direction of a deforestation ban, to promote deforestation. In particular, many events that took place during these years promoted more, rather than less, deforestation for large-scale agricultural production and soy production specifically.

I extend the discussion on causality from section IV here by investigating whether changes in commodity prices and unemployment rates explain changes in patterns of post-deforestation land use in the post-policy period. I test this by regressing the annual area deforested for each agricultural land use on a post-policy indicator, current and lagged annual commodity prices, and unemployment rates. The results are shown in Table 9. Specifications 1 through 3 investigate clearing for large-scale production, specifications 4 through 6 investigate clearing for rangeland production, and specifications 7 through 9 investigate clearing for small-scale production.

The post-policy indicator is only significant in the first specification and indicates a decrease of around 22 thousand hectares per year of deforestation for large-scale production after the policy. The second specification additionally controls for annual commodity prices and unemployment, and the third specification additionally controls for lagged commodity prices and unemployment. These additional controls do not explain the decrease in deforestation for large-scale production in the post-policy period. In fact, the impact of the post-policy indicator on deforestation would need to be even larger once these factors are accounted for.

Specifications 4 and 7 reveal that there is nearly no change, significant or not, in the average area deforested pre- and post-policy for rangeland or for small-scale agriculture. However, once commodity prices and unemployment are controlled for, the coefficient on the post-policy indicator grows in magnitude. This indicates that the patterns detected between deforestation and these macroeconomic variables indicate that deforestation should have increased post-policy, while the observed data shows a decrease. This again provides suggestive evidence that the Zero Deforestation Law decreased deforestation in Eastern Paraguay, and that this change cannot be explained by commodity prices or by unemployment rates.

	Dependent variable:								
	Area cleared for large-scale production			Area cleared for rangeland production			Area cleared for small-scale production		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
post	−21.540**	−78.915	−38.339	−0.061	−131.903	−0.288	−0.964	−97.073	−20.890
	(6.533)	(84.633)	(18.574)	(8.235)	(89.746)	(28.569)	(8.238)	(101.451)	(26.892)
price of soy		−0.038			−0.071			−0.071	
		(0.118)			(0.125)			(0.142)	
price of beef		−0.381			−0.572			−0.576	
		(0.499)			(0.529)			(0.598)	
price of cotton		−1.468			−4.236			−2.660	
		(2.717)			(2.881)			(3.257)	
price of soy, lagged			0.053			−0.101			−0.015
			(0.114)			(0.175)			(0.165)
price of beef, lagged			−0.143			−0.402			−0.138
			(0.347)			(0.533)			(0.502)
price of cotton, lagged			−0.397			0.211			−0.342
			(0.473)			(0.727)			(0.685)
unemployment		−9.634	−2.710		−19.203	1.044		−15.397	−4.118
		(12.526)	(5.609)		(13.283)	(8.628)		(15.016)	(8.122)
Constant	34.407***	328.102	102.712	33.389***	706.188	108.272	25.624***	522.041	123.465
	(4.619)	(447.703)	(56.245)	(5.823)	(474.753)	(86.513)	(5.825)	(536.671)	(81.434)
Observations	8	8	8	8	8	8	8	8	8
R^2	0.644	0.783	0.859	0.00001	0.569	0.410	0.002	0.451	0.479
Adjusted R^2	0.585	0.242	0.507	−0.167	−0.508	−1.064	−0.164	−0.921	−0.824
Residual Std. Error	9.239 (df = 6)	12.486 (df = 2)	10.072 (df = 2)	11.647 (df = 6)	13.241 (df = 2)	15.492 (df = 2)	11.650 (df = 6)	14.968 (df = 2)	14.582 (df = 2)
F Statistic	10.871** (df = 1; 6)	1.447 (df = 5; 2)	2.439 (df = 5; 2)	0.0001 (df = 1; 6)	0.528 (df = 5; 2)	0.278 (df = 5; 2)	0.014 (df = 1; 6)	0.329 (df = 5; 2)	0.368 (df = 5; 2)
Note:									*p<0.1; **p<0.05; ***p<0.01

Table 9: Regressions of area deforested for each land use on a post-policy indicator, annual commodity prices, and unemployment. These price and unemployment controls do not explain the decrease in large-scale use better than the Zero Deforestation Law.

I therefore argue that the decrease in deforestation for large-scale agriculture and the relative lack of a decrease in deforestation for small-scale agriculture and rangeland can be attributed to the Zero Deforestation Law. By this same logic, I also argue that the composition and land use effects discussed above can also be attributed to the policy.

Discussion

Understanding how impacts of deforestation policies are distributed is important because these policies can have significant environmental, social, and economic implications. Evidence presented here suggests that the Zero Deforestation Law in Paraguay successfully decreased deforestation by around 25%. Nearly all deforestation in this region is conducted for agricultural expansion, therefore this decrease can be expected to have had an impact on farming groups. In this section, I investigated how the impacts of the Zero Deforestation Ban on deforestation patterns were distributed across agricultural groups.

Data on the land use of deforested locations is not available, and so I generated this data using a replicable three step process. First, I manually identified the post-deforestation land use for a subset of deforested pixels. Second, I use this data to train random forest models that predict post-deforestation agricultural use. Third, I use the models to generate data on the post-deforestation agricultural use of all deforested pixels in Eastern Paraguay.

My results reveal that deforestation primarily decreased for large-scale crop production post-policy. In the four years before the law was implemented, an average of 34 thousand hectares were cleared for large-scale agricultural production each year. In the four years after the law was implemented this fell to an average of 12 thousand hectares per year, more than a 60% decrease.

Small-scale and rangeland production were not as impacted by the law, with the average clearing for these uses remaining close to constant between the pre- and post-policy periods.

This change in the amount of deforestation for large-scale agriculture can be attributed to two effects. The Zero Deforestation Law changed where deforestation occurred, called the composition effect. Less land was cleared in total post-policy, and this led to decreases in clearing land that would have gone to all three types of agriculture had the law never been implemented. The second effect is the land use effect. Under the land use effect, the same location would be put to different uses depending on if it was cleared pre-policy or post-policy. After the Zero Deforestation Law was implemented, much of the land that would have been used for large-scale agriculture is instead used for rangeland or small-scale agriculture. Taken together, these effects led to the noted decrease in deforestation for large-scale agriculture after the policy was enacted.

The findings from this analysis expand understanding of the wider impacts of the Zero Deforestation Law. From an environmental perspective, the fact that the decrease in deforestation came primarily from a decrease in deforestation for large-scale systems is a potential benefit for biodiversity. Large-scale systems are often less biodiverse than rangeland and small-scale systems, due to low crop diversity in conventional crop production, and to having fewer field borders and uncultivated areas which can host additional flora and fauna (Fahrig et al., 2014). It should be noted, however, that any deforestation is a setback for maintaining the biodiversity habitat and other ecosystem services that forests provide. This biodiversity benefit is not absolute, but rather relative to other possible distributions of deforestation reductions across land uses.

This distribution of impacts across types of agriculture also may represent an equity gain relative to other possible outcomes. This is because the decrease in deforestation did not come at the expense of the more economically vulnerable small-scale farmers. In Paraguay, and in general, large-scale producers and ranchers have access to more capital, land, and other resources than subsistence farmers. Whether or not it is desirable for the subsistence farmers to remain in agriculture, and potentially continue deforestation, is a topic open to debate (Dorward et al., 2009; DFID, 2015). Regardless, a deforestation law that would push small-scale farmers out of their farming livelihood as a side effect, without offering a safety net or alternative income path, would increase the socioeconomic gap between subsistence farmers and other agricultural groups. That the Zero Deforestation Law decreased deforestation without changing overall patterns in deforestation for these subsistence farmers is, therefore, a potential equity win.

Finally, the disproportionate impact of the Zero Deforestation Law on deforestation for large-scale agriculture may have negatively impacted Paraguay's economy and economic growth. Agriculture is an important part of Paraguay's economy, and large-scale agriculture plays an important role in Paraguay's economic development that continues today (Weisskoff, 1992).

This investigation sets the stage for future research. I present suggestive evidence for a causal relationship between implementation of the Zero Deforestation Law and patterns in deforestation pre- and post-policy. Future research can verify this causal link. I also present evidence that the Zero Deforestation Law decreased deforestation for large-scale agriculture. However, this is not the same as saying that the Zero Deforestation Law decreased expansion of large-scale agriculture. It is possible that large-scale agriculture continued to expand by moving into land that was already cleared for other uses. My data cannot address this question. Future research on potential spillovers will be important to fully understand the implications of this distribution of the laws impacts on the different agricultural groups.

VI. Conclusion

In this dissertation I delve into the dynamics of deforestation in Eastern Paraguay and how these dynamics change after the Zero Deforestation Law was implemented by addressing three research questions. My findings contribute to a better understanding of deforestation in this understudied region, as well as contributing a method to investigate similar questions in other data-scarce regions of the world.

In the first analysis, I investigate how commonly identified drivers of deforestation relate to deforestation in Eastern Paraguay before the Zero Deforestation Law is implemented. I test seven hypotheses on the direction of the relationship between these variables and deforestation. My results show that not all patterns I hypothesized based on prior literature hold in the context of Eastern Paraguay. Physical land characteristics are predictive of where deforestation will occur, as expected. However, one of the physical characteristics, elevation, is significantly related to deforestation in the opposite direction from my expectation. This may be due to Paraguay's geography, which includes expansive low-lying waterlogged areas.

Anthropological land characteristics are not as useful as physical characteristics in predicting where deforestation will occur. Previous literature suggests that proximity to population centers and infrastructure increases the probability of deforestation. This does not hold in Eastern Paraguay. This divergence from the findings in other regions highlights the importance of understanding deforestation dynamics for specific settings. Some anthropological characteristics, however, are predictive. Proximity to existing clearings, such as agricultural fields, increases the probability of deforestation as hypothesized.

This analysis provides important information for law enforcement, nongovernmental organizations, and other organizations that would like to manage forests or impact deforestation. Understanding the types of land on which deforestation is most likely to occur enables these groups to prioritize their actions and resources, for example by prioritizing protection of land that has high ecological value and high deforestation risk. This analysis also lays the foundation for the second two analyses in this dissertation by establishing a baseline of deforestation patterns.

The second analysis investigates how deforestation changed after the implementation of the Zero Deforestation Law, and how relationships between drivers of deforestation and deforestation changed in this post-policy period. My results suggest that the Zero Deforestation Law successfully reduced the rate of deforestation. I argue that these changes were caused by the Zero Deforestation Law because other events that occurred around this time cannot explain the decrease in deforestation. In addition, the relationships between commonly identified drivers of deforestation and deforestation were altered in the post-policy period. The magnitude of the relationship between drivers of deforestation and deforestation decreases for many land characteristics. After the law is implemented, land that is higher elevation and flatter remains more likely to be deforested than low-lying, steep land, however the probability of deforestation for land with these characteristics is lower post-policy. I also find that the relationship between mean nearby tree cover on the probability of deforestation reverses after the law is implemented. Before the Zero Deforestation Law came into effect, pixels with low nearby mean forest cover were more likely to be deforested. After the law, pixels with high nearby mean forest cover are more likely to be deforested. This may imply that once deforestation becomes illegal, newly cleared plots are hidden in dense forest. It might also imply that forested locations that were initially passed over are now reinvestigated and sometimes cleared.

One implication of the diminished relationships between drivers of deforestation and deforestation in the post-policy period is that deforestation becomes more difficult to predict. Deforestation is now more widely dispersed across land with characteristics that previously had been useful in predicting deforestation, including elevation, slope, and that maximum nearby nighttime lights. This means that organizations wishing to manage forests or impact deforestation can no longer predict as accurately where deforestation will occur. This, in combination with evidence that plots cleared post-deforestation may be actively hidden under remaining forest cover, makes the Zero Deforestation Law more difficult to enforce.

In the final analysis I investigate the distribution of impacts of the Zero Deforestation Law across types of agriculture. Results reveal that deforestation was the most changed for large-scale agriculture after implementation of the Zero Deforestation Law. Meanwhile, deforestation for small-scale and rangeland systems was relatively unaffected. The changes in deforestation are due both to changes in the locations cleared (composition effect), and changes in the use of cleared areas (land use effect). The composition effect shows that, after the law was passed, deforestation decreased across land that would have gone into each of the three post-deforestation use categories. The largest decrease in deforestation is seen on land that was more likely to go to large-scale agriculture before deforestation became illegal. The land use effect shows that this decrease in deforestation from the composition effect was accentuated by a shift in the agricultural use of deforested land away from large-scale agricultural use and into, primarily, rangeland use. The net impact of these two effects was a large decrease in deforestation for large-scale agricultural use, from 138 thousand hectares to 52 thousand hectares in the four years before and after the policy was implemented. Meanwhile, relatively minor changes in deforestation were observed for small-scale agricultural are rangeland uses pre- and post-policy.

This third analysis sets the stage for a more thorough investigation into the implications of the Zero Deforestation Law. In addition to slowing deforestation in Eastern Paraguay, this law altered the patterns of agricultural expansion. Post-policy, less forest is cleared for large-scale agricultural expansion. This outcome has important environmental, equity, and economic implications relative to other potential outcomes, due to the different environmental, and socioeconomic characteristics of the different systems. Future research can investigate the full implications of the distribution of impacts identified here.

This dissertation uncovered novel knowledge on the patterns of deforestation in Eastern Paraguay and how these patterns changed with the implementation of the Zero Deforestation Law. The incomplete impact of the law and the uneven distribution of impacts reveals that additional enforcement methods or laws may be needed to further decrease deforestation, and especially to reduce deforestation for small-scale and rangeland land use. Before the law was implemented, deforestation for small-scale agriculture or rangeland made up 60% of all deforestation for agricultural expansion. After implementation, these two groups made up over 80% of deforestation for agricultural expansion. Therefore, these groups must be considered for Paraguay to decrease overall deforestation substantially. While the law is an example of success in that it lowered deforestation notably, there is still substantial work to be done if deforestation is to be eliminated as the law's name suggests.

Citations

Alix-Garcia, J., & Gibbs, H. K. (2017). Forest conservation effects of Brazil's zero deforestation cattle agreements undermined by leakage. Global Environmental Change, 47(November), 201–217. https://doi.org/10.1016/j.gloenvcha.2017.08.009

Assunção, J., Gandour, C., Pessoa, P., & Rocha, R. (2017). Property-level assessment of change in forest clearing patterns: The need for tailoring policy in the Amazon. Land Use Policy, 66(January 2016), 18–27. https://doi.org/10.1016/j.landusepol.2017.04.022

Austin, K. G., González-Roglich, M., Schaffer-Smith, D., Schwantes, A. M., & Swenson, J. J. (2017). Trends in size of tropical deforestation events signal increasing dominance of industrial-scale drivers. Environmental Research Letters, 12(5). https://doi.org/10.1088/1748-9326/aa6a88

Blackman, A., Albers, H. J., Ávalos-Sartorio, B., & Murphy, L. C. (2008). Land cover in a managed forest ecosystem: Mexican shade coffee. American Journal of Agricultural Economics, 90(1), 216–231. https://doi.org/10.1111/j.1467-8276.2007.01060.x

Blanc, J. (2015). Enclaves of inequality: Brasiguaios and the transformation of the Brazil-Paraguay borderlands. Journal of Peasant Studies, 42(1), 145–158. https://doi.org/10.1080/03066150.2014.967685

Busch, J., & Ferretti-Gallon, K. (2017). What drives deforestation and what stops it? A meta-analysis. Review of Environmental Economics and Policy, 11(1), 3–23. https://doi.org/10.1093/reep/rew013

Caldas, M. M., Goodin, D., Sherwood, S., Campos Krauer, J. M., & Wisely, S. M. (2015). Land-cover change in the Paraguayan Chaco: 2000–2011. Journal of Land Use Science, 10(1), 1–18. https://doi.org/10.1080/1747423X.2013.807314

Carvalho, T. S., Domingues, E. P., & Horridge, J. M. (2017). Controlling deforestation in the Brazilian Amazon: Regional economic impacts and land-use change. Land Use Policy, 64(2017), 327–341. https://doi.org/10.1016/j.landusepol.2017.03.001

Correia, J. E. (2019). Soy states: resource politics, violent environments and soybean territorialization in Paraguay. Journal of Peasant Studies, 46(2), 316–336. https://doi.org/10.1080/03066150.2017.1384726

Curtis, P. G., Slay, C. M., Harris, N. L., Tyukavina, A., & Hansen, M. C. (2018). Classifying drivers of global forest loss. Science, 361(6407), 1108–1111. https://doi.org/10.1126/science.aau3445

Da Ponte, E., Fleckenstein, M., Leinenkugel, P., Parker, A., Oppelt, N., & Kuenzer, C. (2015). Tropical forest cover dynamics for Latin America using Earth observation data: a review covering the continental, regional, and local scale. International Journal of Remote Sensing, 36(12), 3196–3242. https://doi.org/10.1080/01431161.2015.1058539

De Sy, V., Herold, M., Achard, F., Beuchle, R., Clevers, J. G. P. W., Lindquist, E., & Verchot, L. (2015). Land use patterns and related carbon losses following deforestation in South America. Environmental Research Letters, 10(12). https://doi.org/10.1088/1748-9326/10/12/124004

DFID/UK AID. (2015). DFID’s Conceptual Framework on Agriculture. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/472999/Conceptual-Framework-Agriculture2.pdf

Dirección General de Estadística, Encuestas y Censos (DGEEC). (2004). Paraguay: Resultados Finales: Censo Nacional de Población y Viviendas Año 2002 - Total País.

Dorward, A., Anderson, S., Bernal, Y. N., Vera, E. S., Rushton, J., Pattison, J., & Paz, R. (2009). Hanging in, stepping up and stepping out: Livelihood aspirations and strategies of the poor. Development in Practice, 19(2), 240–247. https://doi.org/10.1080/09614520802689535

Driessen, P., J. Deckers, O. Spaargaren, and F. Nachtergaele. (2001). World Soil Reports 94: Lecture notes on the major soils of the world. Available at: http://www.fao.org/3/y1899e/y1899e.pdf.

FAO. (2022). The State of the World’s Forests 2022. Forest pathways for green recovery and building inclusive, resilient and sustainable economies. Rome, FAO. https://doi.org/10.4060/cb9360en

Fahrig, L., Girard, J., Duro, D., Pasher, J., Smith, A., Javorek, S., King, D., Lindsay, K. F., Mitchell, S., & Tischendorf, L. (2015). Farmlands with smaller crop fields have higher within-field biodiversity. Agriculture, Ecosystems and Environment, 200, 219–234. https://doi.org/10.1016/j.agee.2014.11.018

Fehlenberg, V., Baumann, M., Gasparri, N. I., Piquer-Rodriguez, M., Gavier-Pizarro, G., & Kuemmerle, T. (2017). The role of soybean production as an underlying driver of deforestation in the South American Chaco. Global Environmental Change, 45(May), 24–34. https://doi.org/10.1016/j.gloenvcha.2017.05.001

Finer, M., Novoa, S., Weisse, M. J., Petersen, R., Mascaro, J., Souto, T., Stearns, F., & Martinez, R. G. (2018). Combating deforestation: From satellite to intervention. Science, 360(6395), 1303–1305. https://doi.org/10.1126/science.aat1203

Galeano, L. A. (2012). Paraguay and the expansion of Brazilian and Argentinian agribusiness frontiers. Canadian Journal of Development Studies, 33(4), 458–470. https://doi.org/10.1080/02255189.2012.744301

Garrett, R. D., Lambin, E. F., & Naylor, R. L. (2013). The new economic geography of land use change: Supply chain configurations and land use in the Brazilian Amazon. Land Use Policy, 34, 265–275. https://doi.org/10.1016/j.landusepol.2013.03.011

Grabs, J., Cammelli, F., Levy, S. A., & Garrett, R. D. (2021). Designing effective and equitable zero-deforestation supply chain policies. Global Environmental Change, 70, 102357. https://doi.org/10.1016/j.gloenvcha.2021.102357

Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., Kommareddy, A., Egorov, A., Chini, L., Justice, C. O., & Townshend, J. R. G. (2013). High-resolution global maps of 21st-century forest cover change. Science, 342(6160), 850–853. https://doi.org/10.1126/science.1244693

Huang, C., Kim, S., Altstatt, A., Townshend, J. R. G., Davis, P., Song, K., Tucker, C. J., Rodas, O., Yanosky, A., Clay, R., & Musinsky, J. (2007). Rapid loss of Paraguay's Atlantic forest and the status of protected areas - A Landsat assessment. Remote Sensing of Environment, 106(4), 460–466. https://doi.org/10.1016/j.rse.2006.09.016

Huang, C., Kim, S., Song, K., Townshend, J. R. G., Davis, P., Altstatt, A., Rodas, O., Yanosky, A., Clay, R., Tucker, C. J., & Musinsky, J. (2009). Assessment of Paraguay's forest cover change using Landsat observations. Global and Planetary Change, 67(1–2), 1–12. https://doi.org/10.1016/j.gloplacha.2008.12.009

IMF. (2024a.) Global price of Beef [PBEEFUSDA], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/PBEEFUSDA, January 11, 2024.

IMF. (2024b). Global price of Cotton [PCOTTINDUSDA], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/PCOTTINDUSDA, January 11, 2024.

IMF. (2024c). Global price of Soybeans [PSOYBUSDA], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/PSOYBUSDA, January 11, 2024.

IMF, 2024d. Consumer price index (2010 = 100) – Paraguay, retrieved from the World Bank; https://data.worldbank.org/indicator/FP.CPI.TOTL?end=2022&locations=PY&start=1960&view=chart, February 13, 2024.

INE. (2024). "Principales indicadores de empleo por año de la encuesta, según área de residencia, 1997/98 - 2021. Serie comparable". Retrieved from https://www.ine.gov.py/publicacion/3/empleo, January 22, 2024.

ISRIC. (2005). "Soil and Terrain Database (SOTER) for Latin America and the Caribbean (SOTERLAC), version 2.0." Available at:

https://data.isric.org/geonetwork/srv/eng/catalog.search#/metadata/436bd4b0-7ffc-4272-be57-686b7d7eea7d.

Kalamandeen, M., Gloor, E., Mitchard, E., Quincey, D., Ziv, G., Spracklen, D., Spracklen, B., Adami, M., Aragaõ, L. E. O. C., & Galbraith, D. (2018). Pervasive Rise of Small-scale Deforestation in Amazonia. Scientific Reports, 8(1), 1–10. https://doi.org/10.1038/s41598-018-19358-2

Kazungu, M., Ferrer Velasco, R., Zhunusova, E., Lippe, M., Kabwe, G., Gumbo, D. J., & Günter, S. (2021). Effects of household-level attributes and agricultural land-use on deforestation patterns along a forest transition gradient in the Miombo landscapes, Zambia. Ecological Economics, 186(April). https://doi.org/10.1016/j.ecolecon.2021.107070

Kovacic, Z., & Viteri Salazar, O. (2017). The lose-lose predicament of deforestation through subsistence farming: Unpacking agricultural expansion in the Ecuadorian Amazon. Journal of Rural Studies, 51(2017), 105–114. https://doi.org/10.1016/j.jrurstud.2017.02.002

Laso Bayas, J. C., See, L., Georgieva, I., Schepaschenko, D., Danylo, O., Dürauer, M., Bartl, H., Ho, F., Za, R., Burianchuk, M., Sirbu, F., Magori, B., Blyshchyk, K., Blyshchyk, V., Rabia, A. H., Pawe, C. K., Su, Y., Ahme, M., Panging, K., … Fritz, S. (2022). Drivers of tropical forest loss between 2008 and 2019. Scientific Data, 9(146), 1–8. https://doi.org/10.1038/s41597-022-01227-3

Macedo, M. N., DeFries, R. S., Morton, D. C., Stickler, C. M., Galford, G. L., & Shimabukuro, Y. E. (2012). Decoupling of deforestation and soy production in the southern Amazon during the late 2000s. Proceedings of the National Academy of Sciences of the United States of America, 109(4), 1341–1346. https://doi.org/10.1073/pnas.1111374109

Ministerio de Agricultura y Ganaderia (MAG). (2010). Produccion Agropecuaria: Sintesis estadisticas ano Agricola 2009/2010. Paraguay.

Ministerio de Agricultura y Ganaderia (MAG). (2009). Censo Agropecuario Nacional 2008. San Lorenzo, Paraguay.

Nagel, B. Y. (1999). "Unleashing the fury": The cultural discourse of rural violence and land rights in Paraguay. Comparative Studies in Society and History, 41(1), 148–181. https://doi.org/10.1017/s0010417599001905

Nickson, A. (1981). Brazilian Colonization of the Eastern Border Region of Paraguay. Journal of Latin American studies, 1981, Vol.13 (1), p.111-131

Olson, D. M., Dinerstein, E., Wikramanayake, E. D., Burgess, N. D., Powell, G. V. N., Underwood, E. C., D'Amico, J. A., Itoua, I., Strand, H. E., Morrison, J. C., Loucks, C. J., Allnutt, T. F., Ricketts, T. H., Kura, Y., Lamoreux, J. F., Wettengel, W. W., Hedao, P., & Kassem, K. R. (2001). Terrestrial ecoregions of the world: a new map of life on Earth. Bioscience, 51(11), 933–938.

OpenStreetMap. (2021). "OpenStreetMap Data in Layered GIS Format." Available at: http://download.geofabrik.de/.

Peters, K. (2015). Changing Relationship with the Land and One Another in Pindoyu, Paraguay [The University of New Mexico Albuquerque].

https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.837.4102&rep=rep1&type=pdf

Potapov, P., Hansen, M.C., Kommareddy, I., Kommareddy, A., Turubanova, S., Pickens, A., Adusei, B., Tyukavina A., and Ying, Q., 2020. Landsat analysis ready data for global land cover and land cover change mapping. Remote Sens. (2020), 12, 426; doi:10.3390/rs12030426

Richards, P. D. (2011). Soy, cotton, and the final Atlantic Forest frontier. Professional Geographer, 63(3), 343–363. https://doi.org/10.1080/00330124.2011.566516

Roman, M., Wang, Z., Shrestha, R., Yao, T., Kalb V. (2021). Black Marble VNP46A4 15 arc second annual global data. https://blackmarble.gsfc.nasa.gov/

Ruggiero, P. G. C., Pfaff, A., Nichols, E., Rosa, M., & Metzger, J. P. (2021). Election cycles affect deforestation within Brazil's Atlantic Forest. Conservation Letters, 14(5), 1–11. https://doi.org/10.1111/conl.12818

Taye, F. A., Folkersen, M. V., Fleming, C. M., Buckwell, A., Mackey, B., Diwakar, K. C., Le, D., Hasan, S., & Ange, C. Saint. (2021). The economic values of global forest ecosystem services: A meta-analysis. Ecological Economics, 189(April), 107145. https://doi.org/10.1016/j.ecolecon.2021.107145

UNEP-WCMC (2019). User Manual for the World Database on Protected Areas and world database on other effective area-based conservation measures: 1.6. UNEP-WCMC: Cambridge, UK. Available at: http://wcmc.io/WDPA_Manual

UN FAO, 2020. Evaluacion de los recursos forestales mundiales 2020 informe: Paraguay [pdf]. Presentation for Informe FRA. (2020). Rome. https://www.fao.org/3/cb0111es/cb0111es.pdf

Vancutsem, C., Achard, F., Pekel, J. F., Vieilledent, G., Carboni, S., Simonetti, D., Gallego, J., Aragão, L. E. O. C., & Nasi, R. (2021). Long-term (1990–2019) monitoring of forest cover changes in the humid tropics. Science Advances, 7(10). https://doi.org/10.1126/sciadv.abe1603

Weisskoff, R. (1992). The Paraguayan agro-export model of development. World Development, 20(10), 1531–1540. https://doi.org/10.1016/0305-750X(92)90072-4

Appendix

A1. Additional Deforestation Data

Annual and cumulative deforestation is plotted by year in Eastern Paraguay using the Global Forest Watch (Hansen et al., 2013) and Tropical Moist Forests (Vancutsem et al., 2021) datasets. Global Forest Watch deforestation is restricted to pixels with forest cover of at least 30% in 2000. The Tropical Moist Forest data excludes a low forest cover region in the southwest of Paraguay. The two datasets display similar patterns in deforestation during the time of interest, despite differences in the data generating methods. A visual inspection of the data suggests an increase in deforestation the year prior to presidential elections (2003, 2008, 2013, 2018) after deforestation becomes illegal.

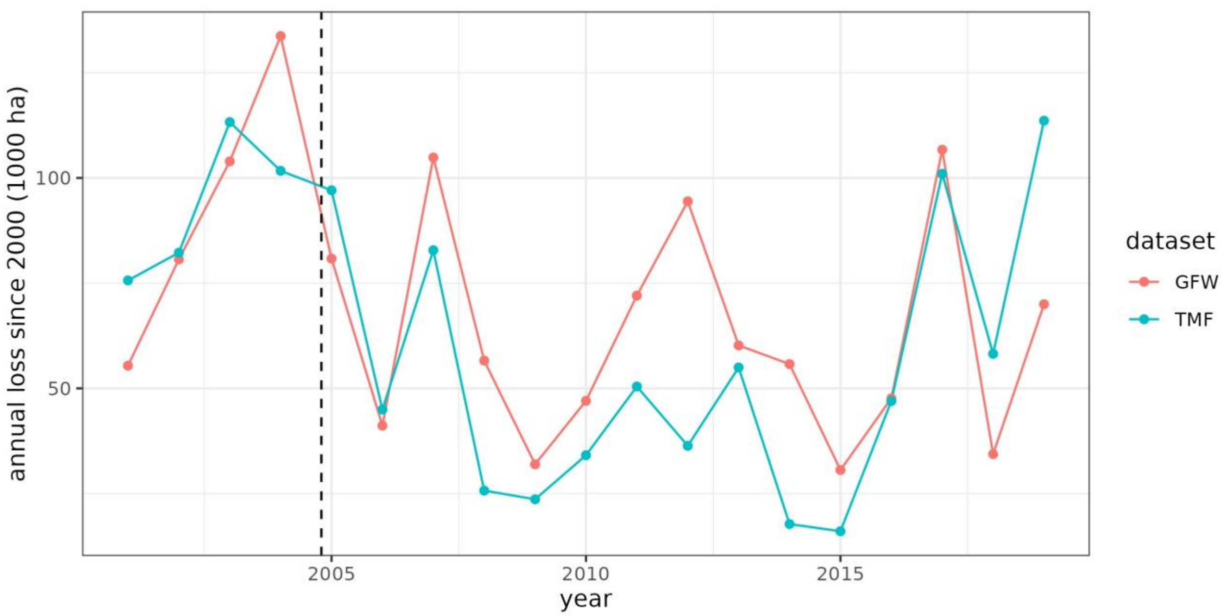

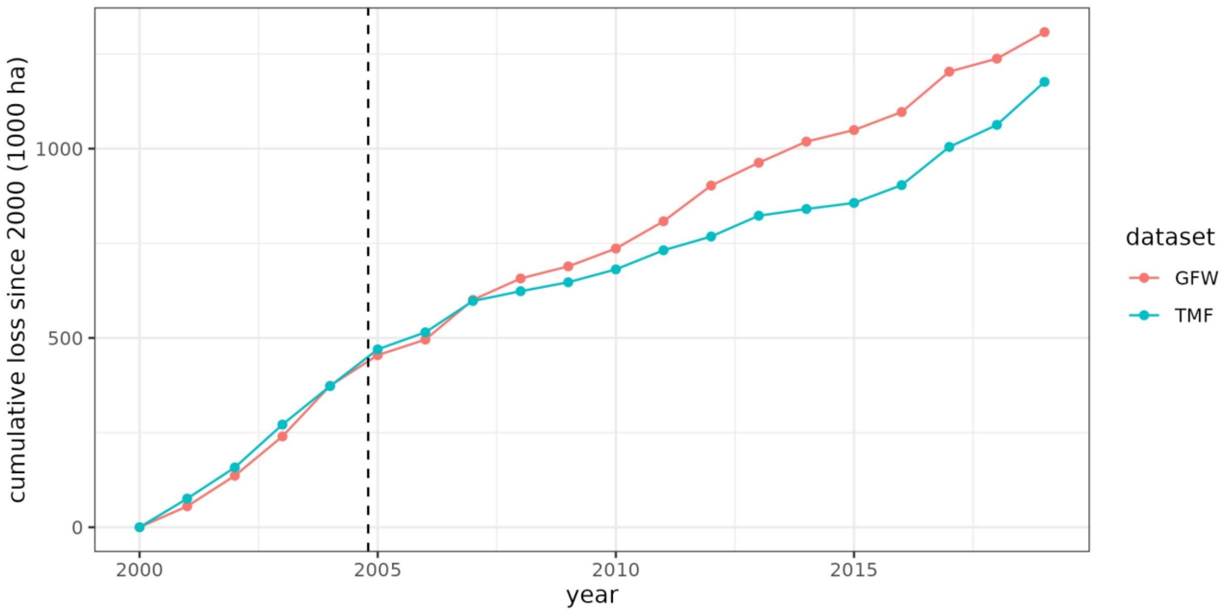
cumulative loss since 2000 (1000 ha)
0
500
1000
2000
2005
2010
2015
year
dataset
GFW
TMF

A2. Soil Group Descriptions

Soil group	Description	Connotation
Acrisols	characterized by accumulation of low activity clays in an argic subsurface horizon and by a low base saturation level.	strongly weathered acid soils with low base saturation; from L. acris, very acid.
Arenosols	sandy soils, both soils developed in residual sands, in situ after weathering of old, usually quartz-rich soil material or rock, and soils developed in recently deposited sands as occur in deserts and beach lands.	sandy soils
Chernozems	soils with a thick black surface layer rich in organic matter.	black soils rich in organic matter
Fluvisols	genetically young, azonal soils in alluvial deposits.	soils developed in alluvial deposits
Ferralsols	the `classical', deeply weathered, red or yellow soils of the humid tropics. These soils have diffuse horizon boundaries, a clay assemblage dominated by low activity clays (mainly kaolinite) and a high content of sesquioxides.	red and yellow tropical soils with a high content of sesquioxides
Gleysols	wetland soils that, unless drained, are saturated with groundwater for long enough periods to develop a characteristic "gleyic colour pattern". This pattern is essentially made up of reddish, brownish or yellowish colours at ped surfaces and/or in the upper soil layer(s), in combination with greyish/bluish colours inside the peds and/or deeper in the soil.	soils with clear signs of excess wetness
Leptosols	accommodates very shallow soils over hard rock or highly calcareous material but also deeper soils that are extremely gravelly and/or stony. Leptosols are azonal soils with an incomplete solum and/or without clearly expressed morphological features. They are particularly common in mountain regions.	shallow soils
Lixisols	strongly weathered soils in which clay has washed out of an eluvial horizon (L. lixivia is washed-out substances) down to an argic subsurface horizon that has low	strongly weathered soils in which clay is washed down from the surface soil to an

	activity clays and a moderate to high base saturation level.	accumulation horizon at some depth
Nitisols	deep, well-drained, red, tropical soils with diffuse horizon boundaries and a subsurface horizon with more than 30 percent clay and moderate to strong angular blocky structure elements that easily fall apart into characteristic shiny, polyhedric (`nutty') elements. Nitisols are strongly weathered soils but far more productive than most other red tropical soils.	deep, red, well-drained tropical soils with a clayey `nitic' subsurface horizon that has typical `nutty', polyhedric, blocky structure elements with shiny ped faces
Planosols	soils with bleached, light-coloured, eluvial surface horizons that show signs of periodic water stagnation and abruptly overly dense, slowly permeable subsoil with significantly more clay than the surface horizon.	soils with a degraded, eluvial surface horizon abruptly over dense subsoil, typically in seasonally waterlogged flat lands
Regosols	a taxonomic rest group containing all soils that could not be accommodated in any of the other Reference Soil Groups. In practice, Regosols are very weakly developed mineral soils in unconsolidated materials that have only an ochric surface horizon and that are not very shallow (Leptosols), sandy (Arenosols) or with fluvic properties (Fluvisols). Regosols are extensive in eroding lands, in particular in arid and semi-arid areas and in mountain regions.	soils in the weathered shell of the earth
Solonetz	soils with a dense, strongly structured, clay illuviation horizon that has a high proportion of adsorbed sodium and/or magnesium ions.	Soils with a high content of exchangeable sodium and/or magnesium ions

Source: Driessen et al. 2001

A3. Additional specifications of the Deforestation Model

I. Main specification, showing soil group and ecoregion coefficients

The following table shows the full regression table for the linear probability model regressions from the main text, of deforestation on demeaned and normalized land characteristics that have been identified as drivers of deforestation in the literature. Nearly all soil groups and nearly all ecoregions are significant predictors of deforestation in each period tested. All coefficients on all groups are negative, to adjust for positive coefficients elsewhere. However, the magnitudes of the coefficients reveal how the probability of deforestation varies between soil groups. No soil group is omitted, because there is no constant in the regression.

The Atlantic Forest is the omitted category for ecoregions. All other ecoregions have negative and significant coefficients. This indicates that deforestation was most likely to occur in the Atlantic Forest Ecoregion.

	Dependent variable:				
	lost 2001-2004	lost 2001	lost 2002	lost 2003	lost 2004
	(1)	(2)	(3)	(4)	(5)
elevation	0.034***	0.005***	0.010***	0.010***	0.011***
	(0.005)	(0.001)	(0.001)	(0.002)	(0.002)
slope	−0.017***	−0.003***	−0.004***	−0.005***	−0.007***
	(0.002)	(0.0003)	(0.001)	(0.001)	(0.001)
tree cover mean, 100m	−0.006**	−0.003***	−0.003***	−0.002**	0.001
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
tree cover stdv, 100m	0.018***	0.004***	0.004***	0.006***	0.005***
	(0.002)	(0.001)	(0.0005)	(0.001)	(0.001)
max NTL 10km	−0.004**	−0.0002	−0.001**	−0.001*	−0.001**
	(0.001)	(0.0005)	(0.0004)	(0.001)	(0.001)
distance to a town	0.001	0.002**	0.001	−0.001	−0.001
	(0.003)	(0.001)	(0.001)	(0.001)	(0.001)
distance to a main road	0.001	0.001	0.0003	0.001	−0.0002
	(0.003)	(0.001)	(0.001)	(0.001)	(0.002)
acrisols	−0.034**	−0.013***	−0.015***	−0.007	−0.003
	(0.015)	(0.004)	(0.003)	(0.005)	(0.006)
arenosols	−0.029**	−0.008**	−0.012***	−0.011**	−0.002
	(0.013)	(0.004)	(0.003)	(0.004)	(0.005)
gleysols	−0.040***	−0.009***	−0.015***	−0.012***	−0.008
	(0.012)	(0.003)	(0.003)	(0.004)	(0.005)
leptosols	−0.071***	−0.015***	−0.022***	−0.021***	−0.020***
	(0.017)	(0.004)	(0.004)	(0.006)	(0.007)
nitisols	−0.075***	−0.018***	−0.022***	−0.017***	−0.024***
	(0.016)	(0.005)	(0.004)	(0.006)	(0.007)
other soil group	−0.041***	−0.010**	−0.015***	−0.011***	−0.008
	(0.013)	(0.004)	(0.003)	(0.004)	(0.006)
planosols	−0.056***	−0.013***	−0.019***	−0.014***	−0.015***
	(0.013)	(0.004)	(0.003)	(0.004)	(0.005)
regosols	−0.055***	−0.014***	−0.019***	−0.016***	−0.012*
	(0.017)	(0.005)	(0.004)	(0.005)	(0.007)
Cerrado	−0.038***	−0.009***	−0.003	−0.012***	−0.016***
	(0.008)	(0.002)	(0.005)	(0.002)	(0.005)
Humid Chaco	−0.026***	−0.005**	−0.002	−0.008***	−0.012***
	(0.006)	(0.002)	(0.001)	(0.002)	(0.003)
other ecoregion	−0.026***	−0.007***	−0.003	−0.004	−0.013***
	(0.010)	(0.002)	(0.002)	(0.005)	(0.005)
SE clusters	1/2 degree	1/2 degree	1/2 degree	1/2 degree	1/2 degree
% forest cleared	5.8%	0.9%	1.2%	1.6%	2.1%
Observations	928,277	928,277	920,255	908,790	893,827
R^2	0.095	0.016	0.022	0.029	0.035
Adjusted R^2	0.095	0.016	0.022	0.029	0.035

Note: *p<0.1; **p<0.05; ***p<0.01

II. Robustness to clustering standard errors

All specifications in the table below are identical except for the level of clustering of the standard errors. The first specification does not cluster standard errors. The second specification clusters standard errors at the one-degree latitude longitude grid. The third specification is identical to the main results, clustering at the half degree level. The fourth specification clusters standard errors at the department level.

The regression results are largely robust to clustering standard errors at different levels, though the increase in significance when no clustering of standard errors is included on this spatial dataset suggests that clustering at some level is important. The increase in significance when there is no clustering of standard errors is due in part to spatial correlation. When standard errors are clustered across space this partially controls for similarities in outcomes for nearby pixels that do not represent a truly i.i.d. sample.

Clustering standard errors at any of the levels tested yields similar results.

	Dependent variable:			
	lost 2001-2004			
	(1)	(2)	(3)	(4)
elevation	0.034***	0.034***	0.034***	0.034***
	(0.0004)	(0.006)	(0.005)	(0.007)
slope	−0.017***	−0.017***	−0.017***	−0.017***
	(0.0002)	(0.002)	(0.002)	(0.002)
tree cover mean, 100m	−0.006***	−0.006*	−0.006**	−0.006**
	(0.0003)	(0.003)	(0.002)	(0.002)
tree cover stdv, 100m	0.018***	0.018***	0.018***	0.018***
	(0.0003)	(0.003)	(0.002)	(0.003)
max NTL 10km	−0.004***	−0.004**	−0.004**	−0.004*
	(0.0002)	(0.002)	(0.001)	(0.002)
distance to a town	0.001***	0.001	0.001	0.001
	(0.0003)	(0.003)	(0.003)	(0.002)
distance to a main road	0.001***	0.001	0.001	0.001
	(0.0003)	(0.004)	(0.003)	(0.003)
SE clusters	none	1 degree	1/2 degree	department
Ecoregions	Yes	Yes	Yes	Yes
Soil groups	Yes	Yes	Yes	Yes
Observations	928,277	928,277	928,277	928,277
R^2	0.095	0.095	0.095	0.095
Adjusted R^2	0.095	0.095	0.095	0.095

Note: *p<0.1; **p<0.05; ***p<0.01

III. Robustness to the inclusion of additional fixed effects

The regression results are generally robust to the inclusion of additional fixed effects. The main specification reported in the results section includes soil group and ecoregion fixed effects. In this regression table, the first specification does not include any fixed effects, including ecoregion and soil group fixed effects. The second specification includes soil group and ecoregion fixed effects, and in addition includes a department fixed effect. The third specification includes soil group and ecoregion fixed effects, and in addition includes a half degree latitude longitude grid fixed effect. In all cases, the sign of the coefficient is the same as in the main specification. The magnitudes are also similar when alternative fixed effects are employed. The significance changes in some cases, but occurs for the same variables that sometimes lacked significance in single-year regressions in the main specification, so this is not altogether surprising.

	Dependent variable:		
	lost 2001-2004		
	(1)	(2)	(3)
elevation	0.036***	0.036***	0.039***
	(0.006)	(0.005)	(0.007)
slope	−0.020***	−0.015***	−0.015***
	(0.004)	(0.002)	(0.002)
tree cover mean, 100m	−0.003	−0.008***	−0.008***
	(0.003)	(0.002)	(0.002)
tree cover stdv, 100m	0.019***	0.017***	0.017***
	(0.003)	(0.002)	(0.002)
max NTL 10km	−0.006*	−0.004**	−0.002
	(0.004)	(0.002)	(0.002)
distance to a town	−0.002	0.002	0.003
	(0.007)	(0.003)	(0.003)
distance to a main road	0.0001	0.0005	0.003
	(0.006)	(0.003)	(0.004)
SE clusters	1/2 degree	1/2 degree	1/2 degree
Ecoregions	No	Yes	Yes
Soil groups	No	Yes	Yes
Additional FE	None	Department	1/2 degree
Observations	928,277	928,277	928,277
R^2	0.029	0.098	0.105
Adjusted R^2	0.029	0.098	0.104
Note:	*p<0.1; **p<0.05; ***p<0.01		

IV. Robustness to a Logit specification

The main results from this section, and every alternative specification displayed in this appendix section, has used a linear probability model. An alternative option is to model the relationship between land characteristics and deforestation using a logit model. The functional form of this model may be more appropriate for a probabilistic outcome, such as the likelihood of deforestation.

I find that the results are robust to the logit specification. The magnitudes of the coefficients in the logit model cannot be directly compared to those in the linear probability model because different functional forms are used. However, the logit model supports the conclusions from the primary model. In all cases, the signs of the coefficients are the same between this logit model and the main linear probability model specification. In addition, the magnitude and significance are similar between the models for all coefficients.

	Dependent variable:
	lost 2001-2004
elevation	0.494***
	(0.006)
slope	−0.348***
	(0.006)
tree cover mean, 100m	−0.033***
	(0.005)
tree cover stdv, 100m	0.343***
	(0.005)
max NTL 10km	−0.117***
	(0.007)
distance to a town	0.008
	(0.007)
distance to a main road	0.008
	(0.006)
SE clusters	1/2 degree
Ecoregions	Yes
Soil groups	Yes
Observations	928,277
Log Likelihood	−185,795.400
Akaike Inf. Crit.	371,626.900
Note:	*p<0.1; **p<0.05; ***p<0.01

A4. Additional Specifications of the Post-Policy Impacts Model

I. Main specification, with soil group and ecoregion coefficients

The following table shows the full results of the linear probability model investigating the change in aggregate deforestation and changes in patterns of deforestation after implementation of the Zero Deforestation Law. All specifications are the same as Table 4 in the main dissertation, but here the coefficients for soil groups and ecoregions are displayed. Acrisols is the omitted soil group and the Atlantic Forest is the omitted ecoregion. There are significant differences in the probabilities of deforestation between soil groups, all else held constant. The Atlantic Forest ecoregion is more likely to be deforested than the others, all else held constant.

	Dependent variable:		
	Deforested		
	(1)	(2)	(3)
post	−0.011***	−0.010**	−0.010***
	(0.004)	(0.004)	(0.003)
elevation		0.024***	0.034***
		(0.004)	(0.004)
slope		−0.015***	−0.017***
		(0.002)	(0.002)
tree cover mean, 100m		0.003	−0.006***
		(0.002)	(0.002)
tree cover stdv, 100m		0.012***	0.019***
		(0.002)	(0.002)
max NTL 10km		−0.003**	−0.004***
		(0.001)	(0.001)
distance to a town		−0.002	−0.001
		(0.003)	(0.003)
distance to a main road		0.0001	0.001
		(0.004)	(0.004)
arenosols		0.016	0.016
		(0.012)	(0.012)
gleysols		0.001	0.001
		(0.008)	(0.008)
leptosols		−0.029***	−0.029***
		(0.007)	(0.007)
nitisols		−0.043***	−0.042***
		(0.010)	(0.010)
other soil group		0.002	0.001
		(0.008)	(0.008)
planosols		−0.017**	−0.017**
		(0.008)	(0.008)
regosols		−0.005	−0.005
		(0.008)	(0.008)
Cerrado		−0.036***	−0.036***
		(0.007)	(0.007)
Humid Chaco		−0.027***	−0.027***
		(0.006)	(0.006)
other ecoregion		−0.038***	−0.038***
		(0.009)	(0.009)
post x elevation			−0.021***
			(0.003)
post x slope			0.005***
			(0.002)
post x tree cover mean, 100m			0.019***
			(0.003)
post x tree cover stdv, 100m			−0.013***
			(0.001)
post x max NTL 10km			0.003**
			(0.001)
post x distance to a town			−0.001
			(0.003)
post x distance to a main road			−0.001
			(0.004)
Constant	0.058***	0.074***	0.074***
	(0.006)	(0.008)	(0.008)
Ecoregions	No	Yes	Yes
Soil groups	No	Yes	Yes
SE clusters	1/2 degree	1/2 degree	1/2 degree
Observations	1,803,027	1,803,027	1,803,027
R^2	0.001	0.030	0.034
Adjusted R^2	0.001	0.030	0.034

Note: *p<0.1; **p<0.05; ***p<0.01

II. Robustness to clustering standard errors

The regression results are largely robust to clustering standard errors at different levels, though clustering at some level is important. All specifications are identical, and the same as the third specification in Table 4 in the main text, except for the level of clustering of the standard errors. The first specification does not cluster standard errors. The second specification clusters standard errors at the one-degree latitude longitude grid. The third specification is identical to the main results, clustering at the half degree level. The fourth specification clusters standard errors at the department level.

The increase in significance when there is no clustering of standard errors is due in part to spatial correlation. When standard errors are clustered across space this partially controls for similarities in outcomes for nearby pixels that do not represent a truly i.i.d. sample.

Clustering standard errors at any of the levels tested yields similar results, though with a limited quantity of differences in levels of significance.

	Dependent variable:			
	lost 2001-2004			
	(1)	(2)	(3)	(4)
post	-0.010^{***}	-0.010^{***}	-0.010^{***}	-0.010^{**}
	(0.0003)	(0.004)	(0.003)	(0.004)
elevation	0.034^{***}	0.034^{***}	0.034^{***}	0.034^{***}
	(0.0003)	(0.006)	(0.004)	(0.007)
slope	-0.017^{***}	-0.017^{***}	-0.017^{***}	-0.017^{***}
	(0.0002)	(0.003)	(0.002)	(0.002)
tree cover mean, 100m	-0.006^{***}	-0.006^{*}	-0.006^{***}	-0.006^{***}
	(0.0003)	(0.003)	(0.002)	(0.002)
tree cover stdv, 100m	0.019^{***}	0.019^{***}	0.019^{***}	0.019^{***}
	(0.0003)	(0.003)	(0.002)	(0.003)
max NTL 10km	-0.004^{***}	-0.004^{**}	-0.004^{***}	-0.004^{*}
	(0.0002)	(0.002)	(0.001)	(0.002)
distance to a town	-0.001^{***}	-0.001	-0.001	-0.001
	(0.0003)	(0.003)	(0.003)	(0.002)
distance to a main road	0.001^{***}	0.001	0.001	0.001
	(0.0003)	(0.004)	(0.004)	(0.003)
post x elevation	-0.021^{***}	-0.021^{***}	-0.021^{***}	-0.021^{***}
	(0.0004)	(0.003)	(0.003)	(0.003)
post x slope	0.005^{***}	0.005^{**}	0.005^{***}	0.005
	(0.0003)	(0.002)	(0.002)	(0.003)
post x tree cover mean, 100m	0.019^{***}	0.019^{***}	0.019^{***}	0.019^{***}
	(0.0004)	(0.004)	(0.003)	(0.005)
post x tree cover stdv, 100m	-0.013^{***}	-0.013^{***}	-0.013^{***}	-0.013^{***}
	(0.0003)	(0.002)	(0.001)	(0.002)
post x max NTL 10km	0.003^{***}	0.003^{*}	0.003^{**}	0.003^{**}
	(0.0003)	(0.001)	(0.001)	(0.001)
post x distance to a town	-0.001^{**}	-0.001	-0.001	-0.001
	(0.0004)	(0.004)	(0.003)	(0.002)
post x distance to a main road	-0.001^{***}	-0.001	-0.001	-0.001
	(0.0004)	(0.004)	(0.004)	(0.002)
Constant	0.074^{***}	0.074^{***}	0.074^{***}	0.074^{***}
	(0.0004)	(0.012)	(0.008)	(0.009)
SE clusters	none	1 degree	1/2 degree	department
Ecoregions	Yes	Yes	Yes	Yes
Soil groups	Yes	Yes	Yes	Yes
Observations	1,803,027	1,803,027	1,803,027	1,803,027
R^2	0.034	0.034	0.034	0.034
Adjusted R^2	0.034	0.034	0.034	0.034
Note:				*p<0.1; **p<0.05; ***p<0.01

III. Robustness to the inclusion of additional fixed effects

The regression results are generally robust to the inclusion of additional fixed effects. The main specification reported in Table 4 includes soil group and ecoregion fixed effects. In this regression table, the first specification does not include any fixed effects, including ecoregion and soil group fixed effects. The second specification includes soil group and ecoregion fixed effects, and in addition includes a department fixed effect. The third specification includes soil group and ecoregion fixed effects, and in addition includes a half degree latitude longitude grid fixed effect. In all cases, the sign of the coefficient is the same as in the main specification. The magnitudes are also similar when using alternative fixed effects. There are limited cases where the significance of a variable changes across specifications.

	Dependent variable:		
	lost 2001-2004		
	(1)	(2)	(3)
post	-0.011^{***}	-0.010^{***}	-0.010^{***}
	(0.003)	(0.003)	(0.003)
elevation	0.036^{***}	0.038^{***}	0.041^{***}
	(0.004)	(0.004)	(0.006)
slope	-0.020^{***}	-0.015^{***}	-0.015^{***}
	(0.002)	(0.002)	(0.002)
tree cover mean, 100m	-0.003	-0.008^{***}	-0.009^{***}
	(0.002)	(0.002)	(0.002)
tree cover stdv, 100m	0.019^{***}	0.018^{***}	0.017^{***}
	(0.002)	(0.002)	(0.002)
max NTL 10km	-0.006^{***}	-0.004^{**}	-0.003^{*}
	(0.002)	(0.002)	(0.002)
distance to a town	-0.002	0.00003	-0.0003
	(0.004)	(0.003)	(0.003)
distance to a main road	0.0001	0.0002	0.002
	(0.003)	(0.003)	(0.004)
post x elevation	-0.021^{***}	-0.021^{***}	-0.021^{***}
	(0.003)	(0.003)	(0.003)
post x slope	0.005^{***}	0.005^{**}	0.005^{**}
	(0.002)	(0.002)	(0.002)
post x tree cover mean, 100m	0.019^{***}	0.019^{***}	0.019^{***}
	(0.003)	(0.003)	(0.003)
post x tree cover stdv, 100m	-0.014^{***}	-0.013^{***}	-0.013^{***}
	(0.001)	(0.001)	(0.001)
post x max NTL 10km	0.003^{**}	0.003^{**}	0.002^{**}
	(0.001)	(0.001)	(0.001)
post x distance to a town	-0.001	-0.001	-0.001
	(0.003)	(0.004)	(0.003)
post x distance to a main road	-0.001	-0.001	-0.001
	(0.004)	(0.004)	(0.004)
Constant	0.058^{***}	0.078^{***}	0.028^{**}
	(0.004)	(0.012)	(0.014)
SE clusters	1/2 degree	1/2 degree	1/2 degree
Ecoregions	No	Yes	Yes
Soil groups	No	Yes	Yes
Additional FE	None	Department	1/2 degree
Observations	1,803,027	1,803,027	1,803,027
R^2	0.025	0.039	0.046
Adjusted R^2	0.025	0.039	0.046
Note:			$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

IV. Robustness to a Logit Specification

The table below uses a logit specification to investigate changes in aggregate deforestation and drivers of deforestation after the Zero Deforestation Law is implemented. Most relationships are consistent with findings in the main text.

However, the distance from a town becomes significant, with deforestation more likely closer to a road. This matches expectations from the literature. In addition, distance to a town and distance to a main road become significant in the post-policy period, suggesting that deforestation may be moving closer to infrastructure once deforestation becomes illegal, all else constant. The magnitudes of these coefficients of the normalized variables are small relative to other variables, however, so a one standard deviation change in these variables remains less predictive than, for example, a one standard deviation change in elevation or slope, all else held constant.

	Dependent variable:
	Deforested
post	−0.134***
	(0.008)
elevation	0.539***
	(0.005)
slope	−0.348***
	(0.006)
tree cover mean, 100m	−0.033***
	(0.005)
tree cover stdv, 100m	0.360***
	(0.005)
max NTL 10km	−0.120***
	(0.007)
distance to a town	−0.035***
	(0.007)
distance to a main road	−0.003
	(0.006)
post x elevation	−0.331***
	(0.008)
post x slope	0.006
	(0.010)
post x tree cover mean, 100m	0.416***
	(0.008)
post x tree cover stdv, 100m	−0.148***
	(0.008)
post x max NTL 10km	0.057***
	(0.011)
post x distance to a town	−0.041***
	(0.009)
post x distance to a main road	−0.027***
	(0.009)
Constant	−2.700***
	(0.008)
Ecoregions	Yes
Soil groups	Yes
SE clusters	1/2 degree
Observations	1,803,027
Log Likelihood	−338,101.000
Akaike Inf. Crit.	676,254.100
Note:	*p<0.1; **p<0.05; ***p<0.01

V. Remove the years before and after implementation

There may be concern that the detected post-policy decrease in deforestation is driven by a ramp-up in deforestation pre-policy. Land managers may have learned that the law would pass, and decided to clear while deforestation was still legal. This would increase deforestation in the pre-policy period. Similarly, delayed awareness of the law may have delayed a decrease in deforestation. This could increase deforestation in the post-policy period. Either of these scenarios could be problematic. To investigate the impact of deforestation trends directly before and after implementation, I remove one year before and after implementation from my panel.

The interpretation for land characteristics and interaction terms are similar to the main specification. The coefficient on the post-policy indicator remains negative, but is no longer significant. The magnitudes suggest a decrease in the rate of deforestation in the range of 3.5% to 7% from a pre-deforestation rate of 5.8%. This is lower than the nearly 20% reduction in the rate of deforestation implied from the full results. This suggests that the Zero Deforestation Law may have impacted deforestation less than suggested there. However, this estimation relies on only three years in the pre-policy and three years in the post-policy period. Deforestation in one of the post-policy years, 2007, may be inflated by pre-election-year trends that are not controlled for in this regression, as discussed in section IV. This would attenuate a decrease in deforestation, both here and in the main regression. These results suggest that the decrease in deforestation post-policy may have been less than the main specification suggests, however these results do not discredit the findings above. Further causal analyses are needed to determine how much of the detected decrease in deforestation can be attributed to the Zero Deforestation Law.

	Dependent variable:		
	Deforested		
	(1)	(2)	(3)
post	−0.004	−0.002	−0.003
	(0.003)	(0.003)	(0.003)
elevation		0.017***	0.025***
		(0.002)	(0.003)
slope		−0.010***	−0.011***
		(0.001)	(0.001)
tree cover mean, 100m		0.001	−0.007***
		(0.002)	(0.002)
tree cover stdv, 100m		0.009***	0.014***
		(0.001)	(0.001)
max NTL 10km		−0.001	−0.002**
		(0.001)	(0.001)
distance to a town		−0.0004	0.0001
		(0.002)	(0.002)
distance to a main road		0.0004	0.001
		(0.003)	(0.002)
post x elevation			−0.015***
			(0.002)
post x slope			0.002
			(0.002)
post x tree cover mean, 100m			0.017***
			(0.002)
post x tree cover stdv, 100m			−0.010***
			(0.001)
post x max NTL 10km			0.002*
			(0.001)
post x distance to a town			−0.001
			(0.003)
post x distance to a main road			−0.002
			(0.003)
Constant	0.037***	0.047***	0.047***
	(0.004)	(0.006)	(0.005)
Ecoregions	No	Yes	Yes
Soil groups	No	Yes	Yes
SE clusters	1/2 degree	1/2 degree	1/2 degree
Observations	1,791,476	1,791,476	1,791,476
R^2	0.0001	0.020	0.023
Adjusted R^2	0.0001	0.020	0.023
Note:	*p<0.1; **p<0.05; ***p<0.01		

A5. Training Data Summary

The 501 post-deforestation land use observations that are classified as one of the three agricultural land uses or other use are plotted here, summarized here by type of land use and year of deforestation. These observations were sampled using a stratified random sampling methodology. These data are used to train the random forest models that predict land use class. The vertical dashed line indicates the implementation of the Zero Deforestation Law at the end of 2004.

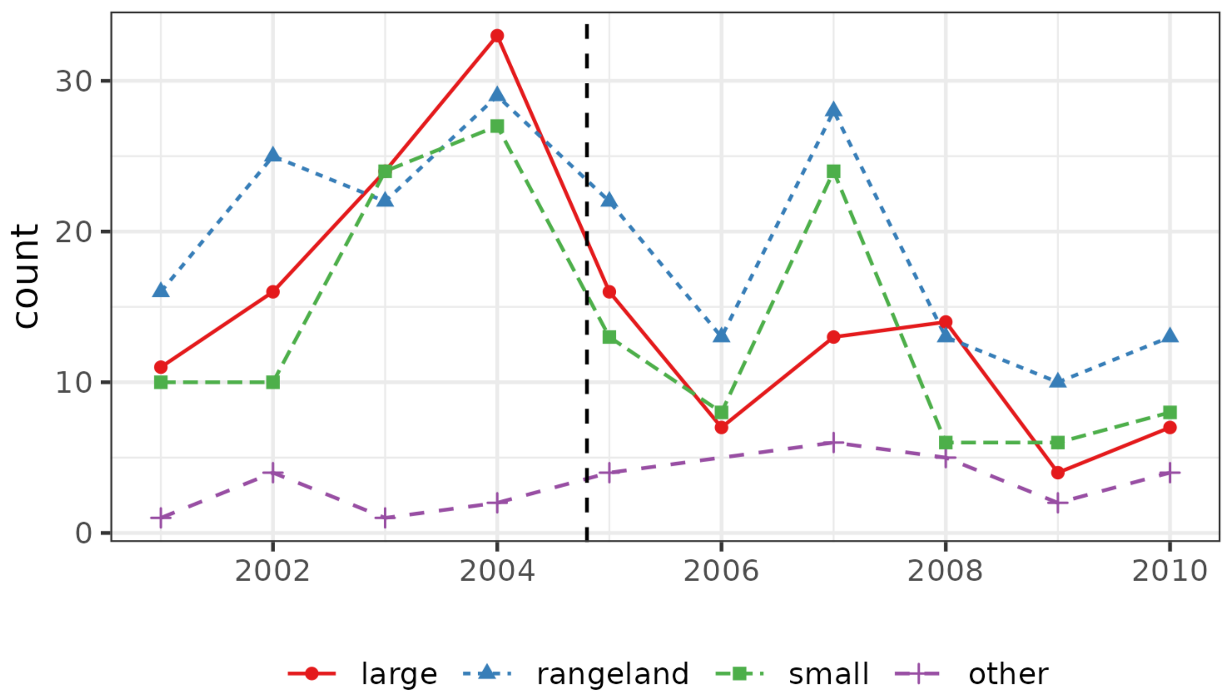

A6. Random Forest Model Statistics

Each decision tree in the random forest model is calculated using a random subset of the training dataset. The random forest algorithm uses the remaining training data to test the tree, similar to a leave-one-out method. The out of box (OOB) error summarizes the error in the model revealed through this testing. The confusion matrix reports the number of correct and incorrect classifications by category. The OOB error and confusion matrices for the pre and post random forest models are displayed in the table below.

		pre model predicted class					post model predicted class				
		large	rangeland	small	other	class error	large	rangeland	small	other	class error
class	large	56	18	10	0	33%	30	22	9	0	51%
	rangeland	18	63	11	0	32%	12	74	11	2	25%
	small	14	12	45	0	37%	9	14	42	0	35%
	other	2	4	2	0	100%	5	3	3	0	100%
		OOB error:	36%				**OOB error:**	41%			

The OOB error rate shows that around 35% to 40% of pixels are misclassified. This rate is substantially lower than would be expected from a random classification of observations into four classes, indicating that the model does have predictive power. The pre model distributes these errors approximately equally across the agricultural land use classes, with large-scale pixels predicted incorrectly 33% of the time, rangeland pixels predicted incorrectly 32% of the time, and small-scale pixels predicted incorrectly 37% of the time.

The distribution of errors is less even across agricultural land uses in the post model predictions. Large-scale pixels are predicted incorrectly far more often than the other classes, at 51% of the time. Large-scale pixels are most likely to be misclassified as rangeland, which adds strength to my result that the largest decrease in in deforestation post-policy came from a decrease in clearing for large-scale agriculture. In the training dataset, a full 36% of pixels with large-scale use post-clearing were misclassified into rangeland. Meanwhile, only 12% of pixels with rangeland use post-clearing were misclassified into large. Thus, the errors suggest that in reality there may be even less large-scale use and more rangeland use post-clearing than the predictions reveal.

Other is consistently predicted incorrectly in both models, though this is not altogether surprising since this category encompasses a variety of unrelated land uses.

In addition to the OOB errors, I conducted a manual test of prediction accuracy using training and testing datasets. These should, and do, reveal similar results as the OOB error. I randomly selected 100 observations to be withheld from the training dataset. I then generated pre and post models using the remaining data and used these to predict the land use category for the withheld test observations. The pre model predicted observations in the pre-policy period with 62% accuracy, while the post model predicted observations in the post-policy period with 52% accuracy. These imply an error rate of 38% and 48%, in line with OOB error estimates obtained above. The prediction results for the withheld test dataset are summarized below.

	pre model			**post model**		
	overall	**pre data**	**post data**	**overall**	**pre data**	**post data**
% error	45%	38%	52%	44%	40%	48%
count	100	48	52	100	48	52
correct	55	30	25	56	29	27

A7. Random Forest Partial plots

Partial plots are generated using the pre and post random forest models that predict post-deforestation land uses. These reveal how the relationship between a land characteristic and the probability of predicting each post-deforestation land use class changes as the value of a land characteristic varies, holding all else constant. Because the pre and post models are run separately, magnitudes should not be compared between the pre and post models.

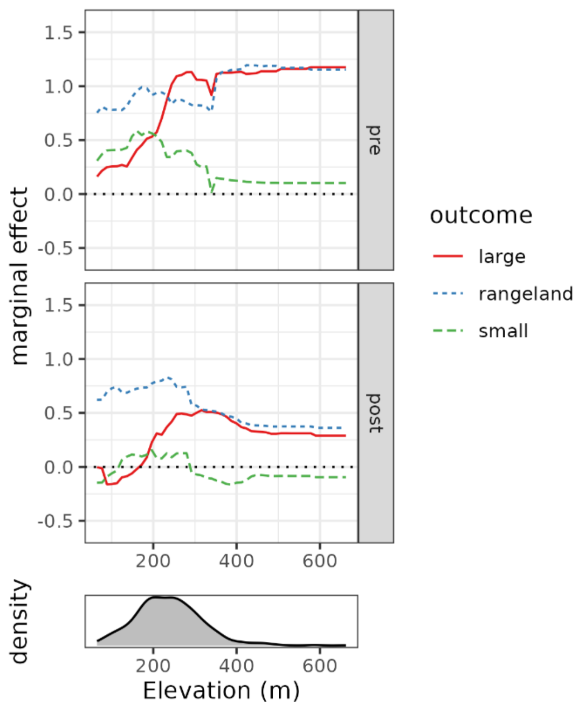

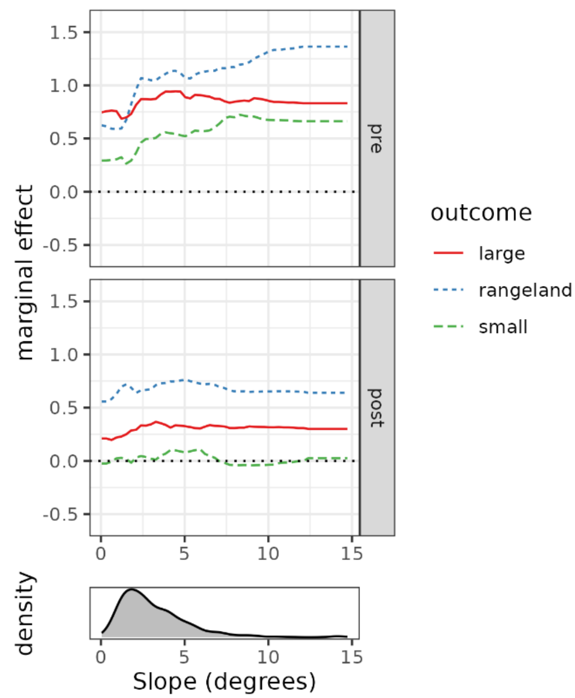

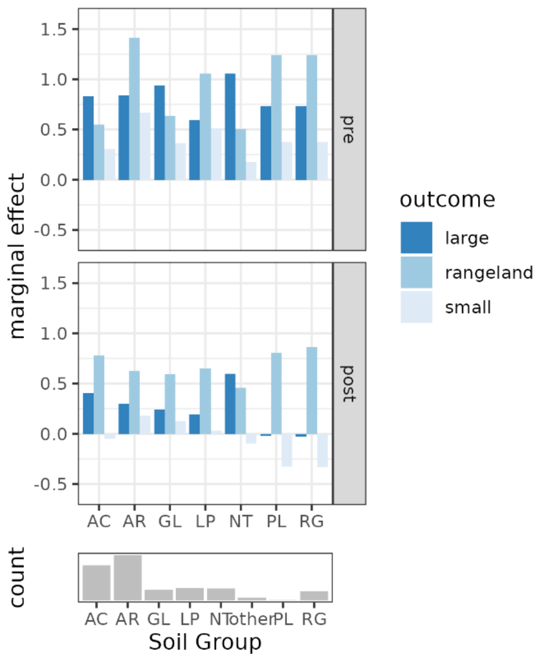

marginal effect
pre
post
count
AC AR GL LP NT PL RG
AC AR GL LP NTotherPL RG
Soil Group
outcome
large
rangeland
small

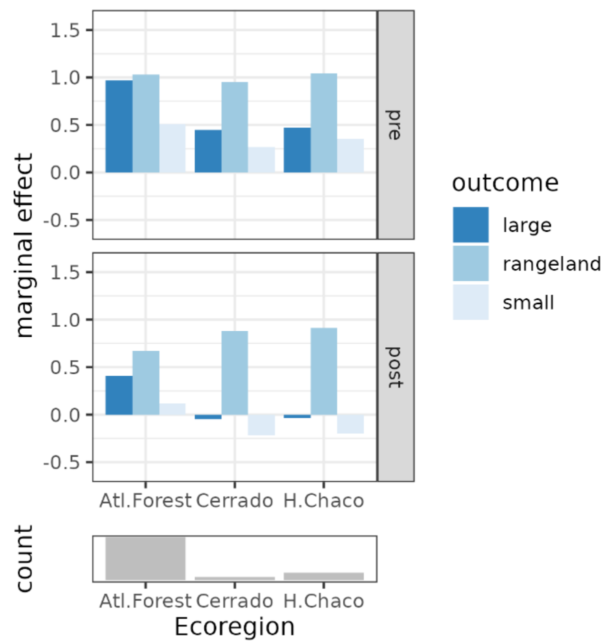

marginal effect
pre
post
count
Atl.Forest Cerrado H.Chaco
Ecoregion
outcome
large
rangeland
small

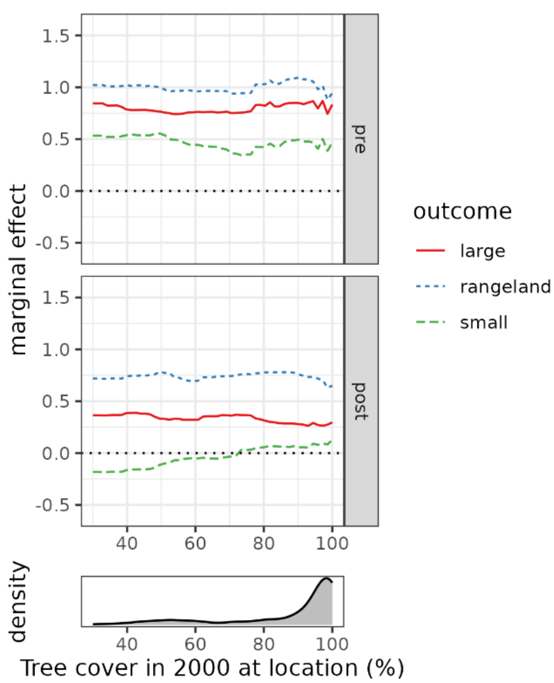

marginal effect
pre
post
density
Tree cover in 2000 at location (%)
outcome
large
rangeland
small

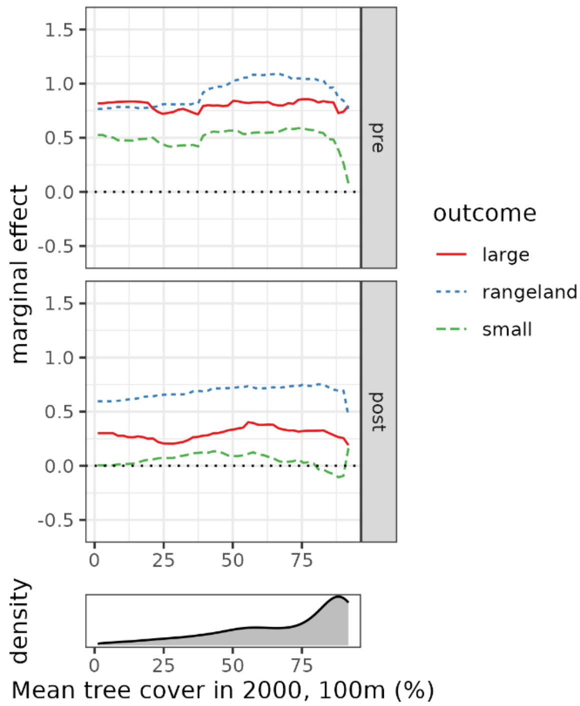
marginal effect
pre
post
outcome
large
rangeland
small
density
Mean tree cover in 2000, 100m (%)

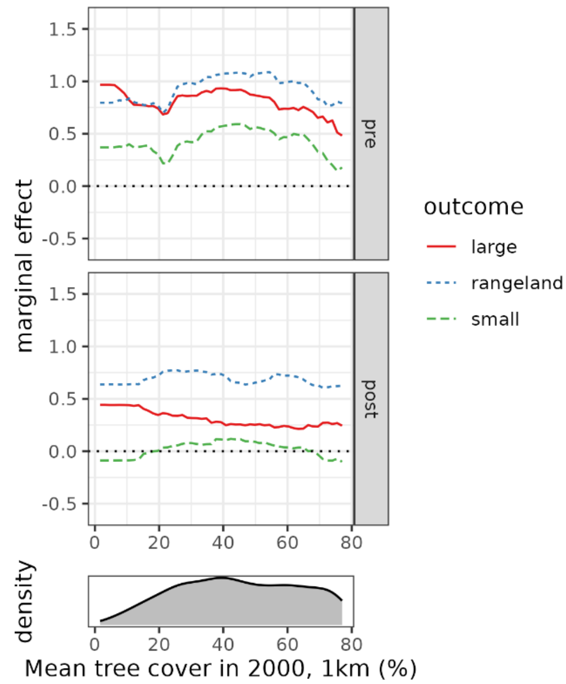
marginal effect
pre
post
outcome
large
rangeland
small
density
Mean tree cover in 2000, 1km (%)

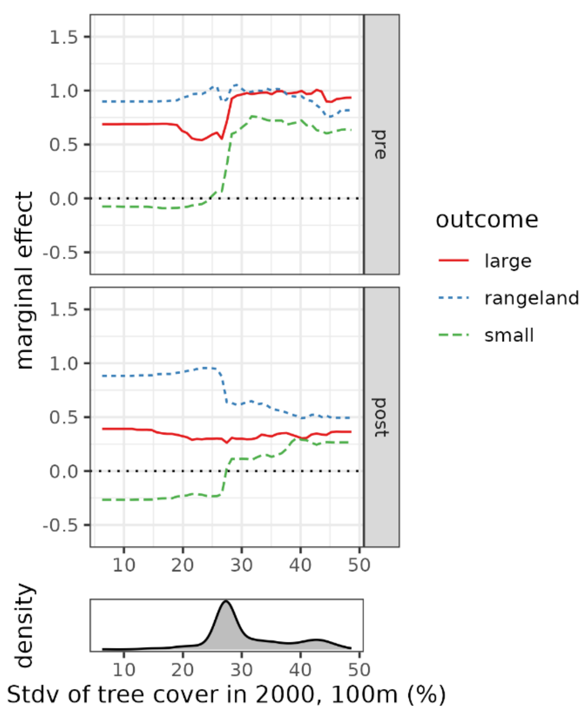
marginal effect
pre
post
outcome
large
rangeland
small
density
Stdv of tree cover in 2000, 100m (%)

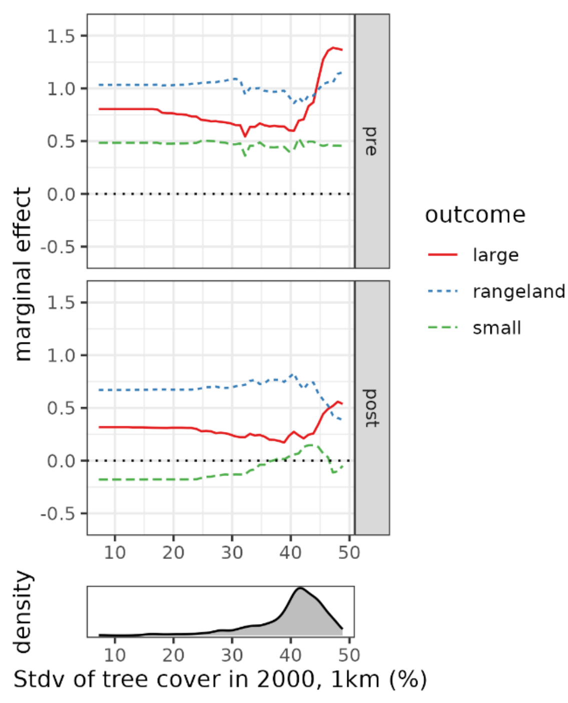
marginal effect
pre
post
outcome
large
rangeland
small
density
Stdv of tree cover in 2000, 1km (%)

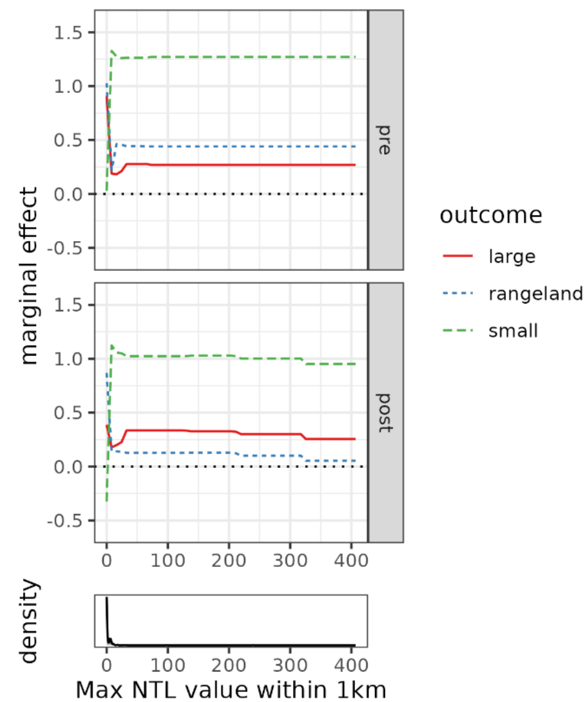
marginal effect
pre
post
outcome
large
rangeland
small
density
Max NTL value within 1km

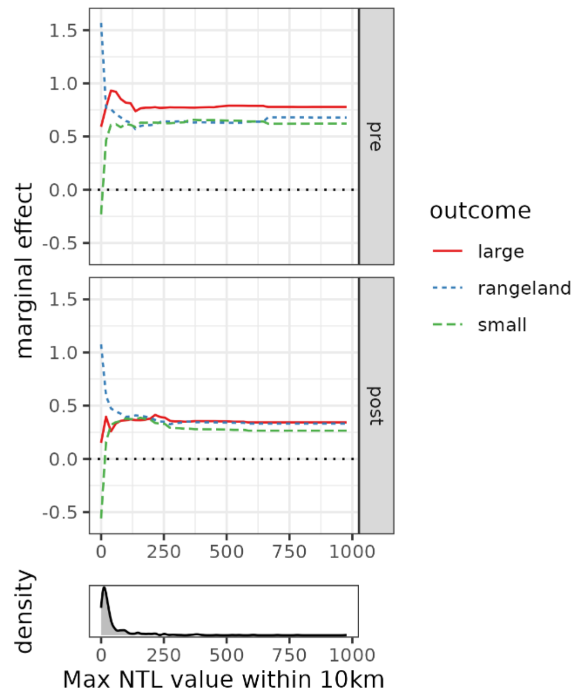
marginal effect
pre
post
outcome
large
rangeland
small
density
Max NTL value within 10km

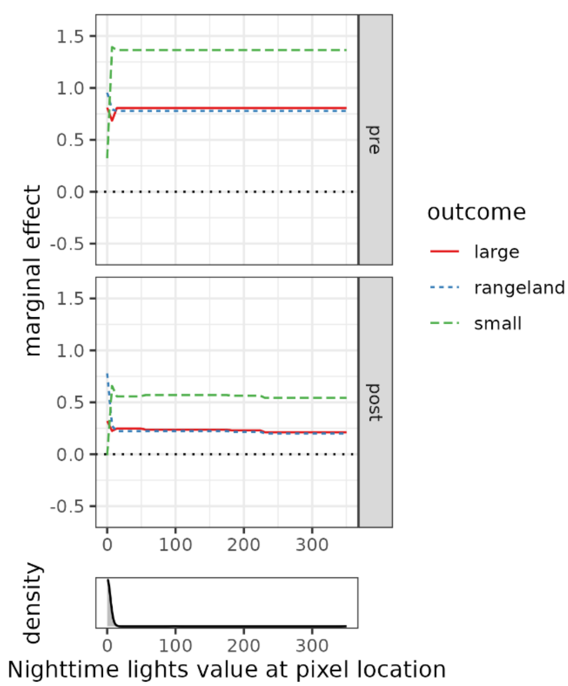
marginal effect
pre
post
outcome
large
rangeland
small
density
Nighttime lights value at pixel location

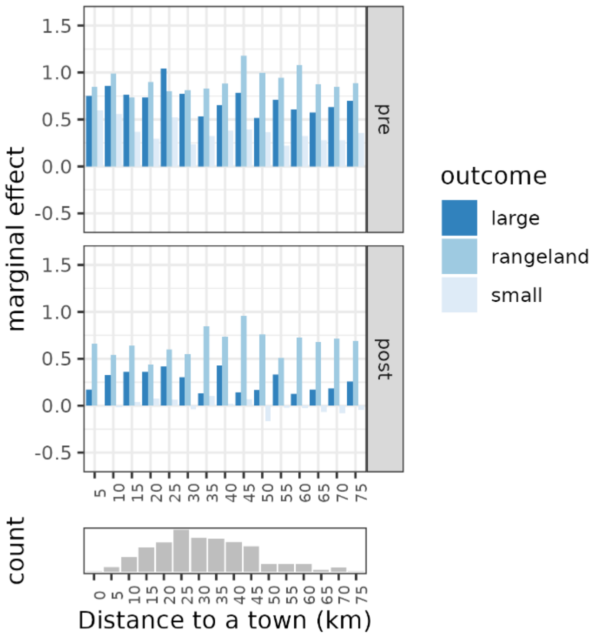
marginal effect
1.5
1.0
0.5
0.0
-0.5
pre
post
5 10 15 20 25 30 35 40 45 50 55 60 65 70 75
count
0 5 10 15 20 25 30 35 40 45 50 55 60 65 70 75
Distance to a town (km)
outcome
large
rangeland
small

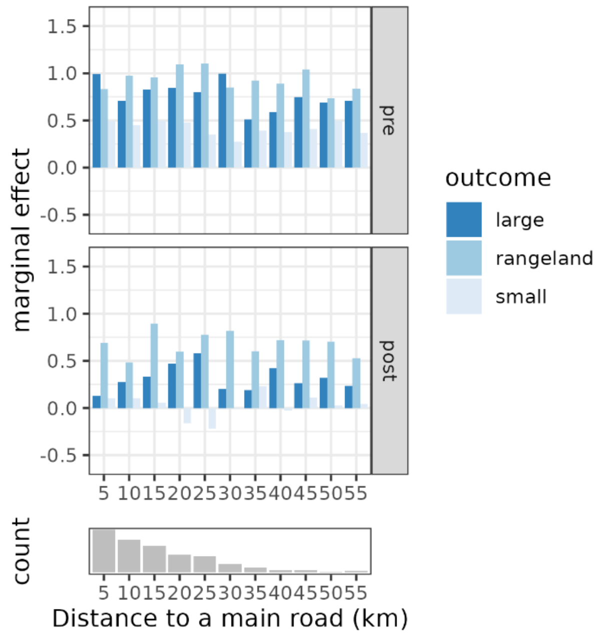
marginal effect
1.5
1.0
0.5
0.0
-0.5
pre
post
5 10152025303540455055
count
5 10152025303540455055
Distance to a main road (km)
outcome
large
rangeland
small

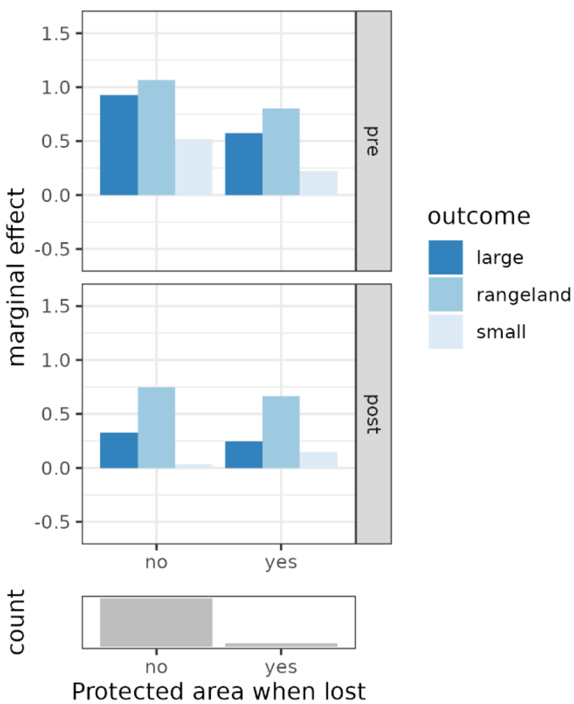
marginal effect
1.5
1.0
0.5
0.0
-0.5
pre
post
no
yes
count
no
yes
Protected area when lost
outcome
large
rangeland
small

A8. Agricultural Land Use Predictions Through 2019

The predicted post-deforestation agricultural land use for all pixels that lost forest cover in Eastern Paraguay between 2001 and 2019 is summarized by year below. The random forest models are generated using training data from 2001 through 2010. Predictions after this period may not be as reliable. The 'best' prediction is generated using the model that best reflects the time period, which is to say the pre model before the Zero Deforestation Law was passed at the end of 2004, and the post model after the law was passed. The counterfactual, shown in dashed lines, uses the pre model to predict post-agricultural land use in the post period. This reflects what land that was actually deforested when deforestation was already illegal would most likely have gone to when deforestation was still legal.