Deforestation in the Amazon:
A Unified Framework for Estimation and Policy Analysis

EDUARDO SOUZA-RODRIGUES*

Department of Economics, University of Toronto

November 10, 2018

Abstract

Deforestation is a matter of pressing global concern, yet surprisingly little is known about the relative efficacy of various policies designed to combat it. This paper sets out a framework for measuring the cost effectiveness of alternative policies – both command-and-control and incentive-based – in the Brazilian Amazon. First, I estimate the demand for deforestation on private properties, exploiting regional variation in transportation costs as a means to recover farmers' responses to permanent policies. Here, rescaling transportation costs using local yields allows me to express changes in farmers' valuations in dollars per hectare. I then use the estimated demand to infer farmers' willingness to deforest under different counterfactual policies, such as payments to avoid deforestation and taxes on land use, along with the corresponding potential farmers' lost surpluses. The results indicate that payment programs and land-use taxes on agricultural land can be highly effective in preserving the rainforest and also be substantially less expensive than command-and-control policies (approximately 8 times less costly). A carbon tax equal to the social cost of carbon could virtually eliminate all agricultural land in the Amazon, given the low agricultural returns there.

JEL Classifications: Q2, Q57, Q58, L73, L78

KEYWORDS: deforestation, land use, carbon tax, Amazon, rainforest, quantile instrumental variables, social cost of carbon.

*I am grateful to Philip Haile, Steven Berry, Donald Andrews, and Xiaohong Chen for many helpful discussions about this research. I would also like to thank Joseph Altonji, Pryianka Anand, John Asker, Jorge Balat, Gustavo Bobonis, Oliver Bunn, Newton de Castro, Ken Gillingham, Penny Goldberg, Myrto Kalouptsidi, Chris Knittel, Matthew Kotchen, Marcos Longo, Vitor Farinha Luz, Rob McMillan, Marcio Nappo, Ariel Pakes, Abraham Parrish, Peter Phillips, Eustáquio Reis, Ted Rosenbaum, Paul Scott, Boudhayan Sen, Chris Timmins, Olga Timoshenko, Chris Udry, Ed Vytlacil, Marfisa Queiroz, the Editor and four anonymous referees, and seminar participants at Harvard, UPenn, University of Wisconsin-Madison, UBC, Boston College, Warwick, and Yale, for further comments and suggestions. Financial support from a Charles V. Hickox Fellowship, the Arthur K. Watson Fund, Yale University, the University of Toronto Mississauga, and a Kernan Brothers Environmental Fellowship is gratefully acknowledged. All remaining errors are mine. E-mail: e.souzarodrigues@utoronto.ca.

1 Introduction

Deforestation is a matter of global concern, not least because of its clear linkage to the pressing issue of climate change. One fifth of greenhouse gas emissions during the 1990s and one tenth during the 2000s have been attributed to deforestation and forest degradation (IPCC, 2007, 2013). Further, reducing deforestation is viewed as a highly cost-effective means of cutting emissions – see, for example, the Stern Review (2007). Absent global coordination, the most likely arena for implementing policy initiatives to combat deforestation is at the national level. In that context, the Brazilian government has been particularly active, both as home to the largest expanse of intact rainforest on the planet, and because of intense deforestation there over the past three decades.[1] The main focus of Brazil's federal government has been on (command-and-control) policies that limit land use on private and public land; and while recent efforts to monitor new deforestation and enforce restrictions applying to it have yielded some success, the rainforest continues to shrink.

Incentive-based interventions are potentially appealing alternatives to command-and-control policies. Examples include payment programs, carbon and land use taxes, and emissions trading.[2] In theory, they can achieve any level of protection at the lowest possible cost to society. While eliminating wasteful expenditures is an important objective by itself and can help avoid political strife, less is known in practice about the magnitudes of these costs, and how they compare to command-and-control policies. At a general level, the accurate measurement of costs of environmental policies is understood to be a challenging problem, as noted by Pizer and Kopp (2005). In a deforestation context, a key obstacle is a lack of clear framework based on credible estimation of the (often-decentralized) process that leads to deforestation, which can then be used to generate reasonable quantitative policy prescriptions. This paper aims to fill that gap.

To that end, I develop a unified approach for measuring the cost effectiveness of alternative policies – both command-and-control and incentive-related – in the Brazilian Amazon. The framework is based on a revealed preference approach, uses credible microeconometric estimates, and can be applied to both existing and yet-to-be implemented policies.

A central component of the framework is the demand for deforestation on private properties. This is defined as the function relating the amount of deforested area to the difference between the

[1] The accumulated deforestation in the Brazilian Amazon between 1988 and 2017 was approximately 42.8 million hectares – an area larger than California (INPE, 2017).

[2] Payment programs pay suppliers directly to provide environmental services (Pattanyak et al., 2010), REDD+ being the most prominent example, negotiated by the United Nations Framework Convention on Climate Change (UNFCCC). Environmental taxes are taxes levied on the production of negative externalities – the carbon tax is the most well-known example. When the measurement of the externalities is difficult, payments and taxes can focus instead on land use.

private value of agricultural and forested land.[3] It provides farmers' willingness to deforest at different levels of private costs (affected by potential taxes and payments), as well as the corresponding potential surpluses.

To identify demand, I exploit the fact that regional variation in transportation costs can be used to infer the value of agricultural land relative to forested land. To gain intuition, consider one farm located close to a major port and another that is far away. *Ceteris paribus*, as transportation costs increase, both the values of agricultural and forested land should decrease. Yet if the value of the agricultural land is disproportionately affected by transportation costs, its *relative* value should also fall, and as a result, one would expect less deforestation in the more distant farm. In this way, variation in transportation costs can be exploited to infer how farmers will respond to changes in private costs. Then, by rescaling the transportation costs using local yields, I am able to value the difference between these two forms of land use in dollars per hectare.

Following this logic, the strategy I propose is divided into two steps: first, I estimate the effects of transportation costs on deforestation, and second, I rescale these costs using local yields to recover the demand function.

My focus is on permanent policies and on their effects, as opposed to transitional dynamics. Because deforesting is costly, farmers are more likely to respond to persistent changes in private values than to temporary changes. Policy interventions will thus have substantial impacts on deforestation if they are put into effect for a long period of time. In this vein, exploiting regional variation in transportation costs is appropriate because persistent changes in private values are likely to be captured by differences in transportation costs in a geographical cross-section.[4]

To estimate the model, I combine the Brazilian transportation network of 2006 with the Agricultural Census of 2006, which is the most recent and comprehensive data set available for the agricultural sector in the country.[5] I also supplement these data with detailed spatial information relating to important determinants of land use, such as soil quality, topography, temperature and rainfall. The estimates show significant negative impacts of transportation costs on deforestation, as might be expected.

After estimating the model, I investigate the effects of three policy interventions: (a) payments

[3]The vast majority of the deforested areas in the Amazon are used for agriculture – mostly pasture and cropland (see Section 3).

[4]See Berry (2011) for a discussion of the importance of distinguishing the short-run and the long-run land-use elasticities for biofuels policies, and Scott (2013) for a fully dynamic model of land-use for the US. Important contributions to the literature developing empirical land use models include Stavins and Jaffee (1990), Stavins (1999), Pfaff (1999), and Mason and Platinga (2013), among others. Section 2 relates the current analysis to the literature.

[5]The 2017 Agricultural Census was not available at the time of the analysis.

for ecological services (PES), (b) taxes on agricultural land, and (c) quantitative limits on the deforestation allowed on private properties. Large-scale payment programs and land-use taxes have not yet been adopted in the Brazilian Amazon; instead, the federal government has relied on quantitative limits. By law, landowners in the Amazon are obligated to keep 80 percent of their properties as native forest (the '80 percent rule'). Yet there is ample evidence that this rule has not been perfectly enforced: in the data, forest coverage on private properties is approximately 44 percent (see Subsections 3.4 and 5.3).[6]

Exploring these policy interventions using my framework yields three main findings. First, taxes can be effective in avoiding deforestation. In response to a perfectly-enforced tax of US$ 42.5 per hectare per year on agricultural land, farmers would be willing to maintain 80 percent forest coverage on private properties as opposed to the 44 percent forest coverage observed in the data.[7] The 36 percent difference corresponds to approximately 27.7 million hectares, which is about 5.5 years of the worldwide net forest loss observed over the past decade. Because farmers' gross revenue per hectare in 2006 was US$ 120/ha on average (with a standard deviation of US$ 560/ha), it should be no surprise that many farmers would not be willing to use land for agriculture under such a tax. In addition, policies that *only* target small landholders are not able to promote substantial conservation. The extremely unequal distribution of land in the Amazon suggests that payment programs are unlikely to significantly reduce local poverty and deforestation simultaneously.

Second, the existing legislation (in the form of the '80 percent rule') would be expensive for local farmers if it were perfectly enforced, resulting in at least US$ 4 billion per year of lost farmer surplus. The US$ 42.5/ha tax would also result in 80 percent of forest cover but be substantially less expensive: farmers' lost surplus would be approximately US$ 479 million per year, provided the tax revenues were redistributed to them.[8] This corresponds to a cost saving from the land use tax of approximately 90 percent of the cost of a perfectly enforced '80 percent rule,' which is substantially higher than the cost saving estimates from allowance trading in pollution markets, ranging from 20 to 47 percent (see Schmalensee and Stavins, 2017). The '80 percent rule' would be substantially more expensive than taxes because the more productive farms would use less land for agriculture and so forgo greater profits. Although land use taxes and payment programs differ

[6]The establishment of protected areas on public lands is one of the leading forest conservation policies worldwide, and in Brazil in particular (Pfaff et al., 2015). I do not estimate causal impacts of protected areas on deforestation because my focus is on policies that affect land use on private properties (though I do allow for spillover effects – see discussion in Section 4.2). The types of payment program that I consider involve payments to avoid deforestation. Although payments to replant forests are important, they are not studied here.

[7]Instead of a perfectly enforced tax, one may interpret US$ 42.5/ha as the expected tax that farmers would pay.

[8]Tax revenues would be approximately US$ 658 million per year (0.37 percent of the Brazilian federal budget for 2006).

in several respects (including the distribution of preservation costs, and practical implementation issues), they share the same predictions in terms of land use and lost surpluses in the present context. As such, payments of US$ 42.5/ha would yield the same forest cover as taxes, but would require US$ 2.61 billion per year of transfers to farmers (approximately 1.45 percent of the Brazilian federal budget for 2006).[9]

Third, by combining the estimated demand for deforestation with the geographic distribution of the carbon stock in Brazil, I obtain a 'supply of avoided emissions.' If a carbon tax of (or a program paying) US$ 1 per ton of CO_2 per year were implemented, farmers would be willing to avoid emissions of approximately 4.17 billion tons of carbon. That corresponds to about 4.5 years of worldwide emissions from land use change during 2002 to 2011 (IPCC, 2013). Given the low agricultural returns and the large stock of carbon on the ground, a carbon tax set at the social cost of carbon (estimated to be US$ 21/tCO_2 for 2010 according to Greenstone et al. (2013), and US$ 18.5/tCO_2 for 2015 according to Nordhaus (2014) could virtually eliminate all agricultural land in the Amazon.

The rest of the paper is organized as follows: Section 2 places the analysis in the context of the related literature. Section 3 provides relevant background to the Brazilian Amazon. Section 4 presents the empirical framework and a detailed discussion of the identification strategy. Section 5 describes the data. The estimated regressions are shown in Section 6, and Sections 7 and 8 discuss the estimated demand for deforestation and the policy implications, respectively. Section 9 concludes.[10]

2 Related Literature

This paper builds on important prior work examining land use, deforestation, and environmental policies and regulations. The IPCC (2007, 2013) and the Stern Review (2007) make extensive use of 'engineering' models, in which the values of alternative land uses are calculated from the revenues and costs of the different alternatives facing a representative farm.[11] Although that approach has proven fruitful, it does not incorporate unobserved heterogeneity across farms. As a result, all

[9]A perfectly targeted policy making payments *only* to those who would deforest their lands and *not* paying those who would *not* deforest would require less than half of the non-targeted program transfers: US$ 1.18 billion per year. The geographic pattern of deforestation under taxes or payments would also be different from the pattern under the '80 percent rule' being more concentrated in the South Amazon, which is arguably the most productive area. As a result, forests in the central regions of the rainforest would have been less fragmented, which may be advantageous from a biodiversity point of view.

[10]The Supplemental Material complements the main text with various robustness exercises and a detailed explanation of the construction of the variables used in the paper.

[11]See Kindermann et al. (2008), Nepstad et al. (2007), and, for a thorough review, Lubowski and Rose (2013).

farmers in a region would prefer to not deforest when taxes (or payments) reduce the *average* value of agricultural land sufficiently compared to the *average* value of forested land. When farms are heterogeneous, however, the marginal unit of land in a region differs from the average unit and so the estimated impacts on deforestation and the estimated costs of policies may be biased.

Revealed-preference methods incorporating unobserved heterogeneity were first developed in the seminal contributions of Stavins and Jaffee (1990) and Stavins (1999). Typically, existing studies estimate reduced-form parameters of farmers' land use choice models using short panel data, exploiting variables with a high degree of variation across time, such as prices or revenues. Different empirical approaches identify farmers' short-run or long-run responses, depending on the time frame covered in the data.[12]

The distinction between farmers' short-run versus long-run responses is important in order to evaluate the performance of alternative policies. Policy makers may prefer implementing policies that can be put into effect for a long period of time. As such, estimates obtained from short panel data that exploit say, year-to-year variation in prices may not provide a reliable indication as to how farmers would react to a counterfactual permanent policy change. In other words, counterfactual simulations may suffer from an external validity problem.

One possible solution is to estimate a forward-looking dynamic structural model of land use choice. In a recent innovative contribution, Scott (2013) implements a structural model in a US setting and finds that dynamic reduced-form models (i.e., with myopic agents) likely understate long-run land use responses significantly. Yet structural dynamic models require access to rich panel data (involving detailed information on land use transitions as well as the evolution of local prices and production costs of different agricultural activities) and such data are not always available, especially in developing countries, where most of the worldwide deforestation has been occurring (FAO, 2016). When access to data is somewhat limited, another approach is necessary (see Timmins and Schlenker, 2009; Brady and Irwin, 2011).

In this paper, I exploit regional variation in transportation costs as a means to recover farmers' responses to permanent policies. The first step of my strategy builds on a growing literature that estimates the determinants of land use.[13] The typical exercise in that literature involves regressing deforestation on covariates using ordinary least squares (OLS), flexible methods such as matching, or fixed-effect methods using panel data. A common focus has been on the impacts of

[12] See Lubowski et al. (2006), Busch et al. (2012), Mason and Platinga (2013), and, for a recent survey of the literature, Brady and Irwin (2011).

[13] See Chomitz and Gray (1996), Pfaff (1999), Andersen et al. (2002), Chomitz and Thomas (2003), Brady and Irwin (2011), and the literature cited therein.

roads on deforestation; existing studies do not attempt to estimate how farmers would respond to counterfactual incentive-based policies.

Compared to that literature, the first part of my strategy adds a plausible instrumental variables approach to address the potential endogeneity of roads and measurement error in transportation costs, in a similar spirit to Chomitz and Gray (1996) (see Subsection 4.2 for a detailed discussion of the identification strategy). In addition, I allow for heterogeneous impacts across municipalities through the use of quantile regressions (specifically, I make use of the instrumental variable quantile regression estimator developed by Chernozhukov and Hansen (2008)). Conditional on observables, a highly deforested location may be well-suited to agriculture in terms of unobservables so that transportation costs would have to increase considerably to reduce the amount of agricultural land; better preserved locations, in contrast, may be more sensitive to changes in transportation costs. An implication is that the relative value of the agricultural land also likely differs at the upper and lower tails of the conditional distribution of deforestation. These differences may lead to non-trivial impacts on the aggregate costs of the policy interventions. My results (shown in detail in Sections 6 and 7) confirm this: the estimated coefficients differ across quantiles, and have important impacts on the total policy costs.

A related literature focuses on the impacts of roads on deforestation and on local economies using treatment effects methods. Those studies construct treatment and control groups based on regions (say, census tracts) that received road investments versus regions that did not. Examples include Banerjee et al. (2012) for impacts on local GDP in China, and Pfaff and Robalino (2013) for impacts on deforestation in the Brazilian Amazon. Although those approaches are appealing, they are not well-suited for my current purposes. First, the estimated differences in land use between locations close to and distant from local roads can explain the *differential* impacts on the groups, but I need to estimate the *overall* impacts of roads on land use to back out farmers' private values. In addition, they do not take into account improvements in roads elsewhere in the transportation network (except by use of adjacency matrices). I instead consider roads as a network and estimate aggregate impacts of transportation networks in a spirit similar to Donaldson and Hornbeck (2016).

An important literature studies recently-implemented 'payment for ecological services' programs.[14] Those studies use treatment effect techniques and find mixed results. Possible explanations include variation in program design, and the fact that evaluations have taken place in countries – mainly Mexico and Costa Rica – in which deforestation rates were declining over the period of

[14]See Pattanyak et al. (2010), Jack (2013), Alix-Garcia and Wolff (2014), Jayachandran et al. (2017), Simonet et al. (2018), and the literature cited therein.

the program (Alix-Garcia and Wolff, 2014). The current paper complements that literature by providing a framework for estimating the potential effects of policies yet to be implemented.

Several studies examine environmental regulations and policies in various contexts, beyond land-use and deforestation issues. Schmalensee and Stavins (2017) compare the performance of incentive-based and command-and-control policies focusing on the most prominent emissions trading systems implemented over the past 30 years in developed countries; Greenstone and Jack (2015) provide a recent survey of a growing body of research investigating environmental regulations in developing countries. In the context of climate change issues, Burke et al. (2016) stress the importance of improving our understanding of the empirical consequences of several policies (including carbon pricing schemes, tradable obligations, and taxes), as well as their optimal design, with an emphasis on impacts on developing economies. My paper contributes to this literature by investigating the cost effectiveness of alternative relevant policies related to the deforestation process in a major developing country.

3 The Brazilian Amazon – Relevant Background

3.1 A Brief History of the Occupation of the Amazon

Before the 1960s, the Amazon was barely occupied. Access to the rainforest was typically open and local economic activity was based on subsistence and a few extraction activities, mainly involving rubber and Brazil nuts. Most municipal seats were established by late 1800s and early 1900s as a result of these local activities.[15] During the 1960s and 1970s, the military dictatorship promoted the occupation of the region with the explicit objective of securing national borders and integrating the region's economy. Hydroelectric facilities, mining, ports, and around 60,000 km of roads were constructed during this period. The first overland connection between Amazonia and the rest of the country was completed in 1964 – a highway linking Belem, an Amazonian state capital, and Brasilia, the country's capital city located in the central region. During the 1980s, an economic recession and hyperinflation led the government to cut investment. Then, after the 1990s, ecological concerns started to shape the policies in the Amazon. In 1996, the required share of forest cover on private land in the Amazon increased from 50 percent to 80 percent.

Deforestation rates had increased substantially since the 1960s (INPE, 2017). After the peak of deforestation in 2004, when an area almost the size of Belgium was deforested in a single year, the Brazilian government launched a new and ambitious conservation program that focused on

[15] Source: https://cidades.ibge.gov.br/

two main areas: improvements in remote sensing-based monitoring and the expansion of protected areas. Then, in 2008, Brazil issued a list of municipalities for the first time, updated annually, that were to be subject to more rigorous monitoring and stricter policy enforcement. The deforestation rate slowed down significantly after 2009.[16]

3.2 Occupied Area

Figure 1 shows the map of Brazil in the left panel, along with the location of the Amazon rainforest, the political divisions and the names of the Amazonian states. The right panel shows the deforested area in 2006 according to satellite images produced by the Brazilian National Institute of Space Research – INPE. Most of the deforested area is concentrated in the southern and eastern parts of the Amazon, normally referred to as the 'Arc of Deforestation.'

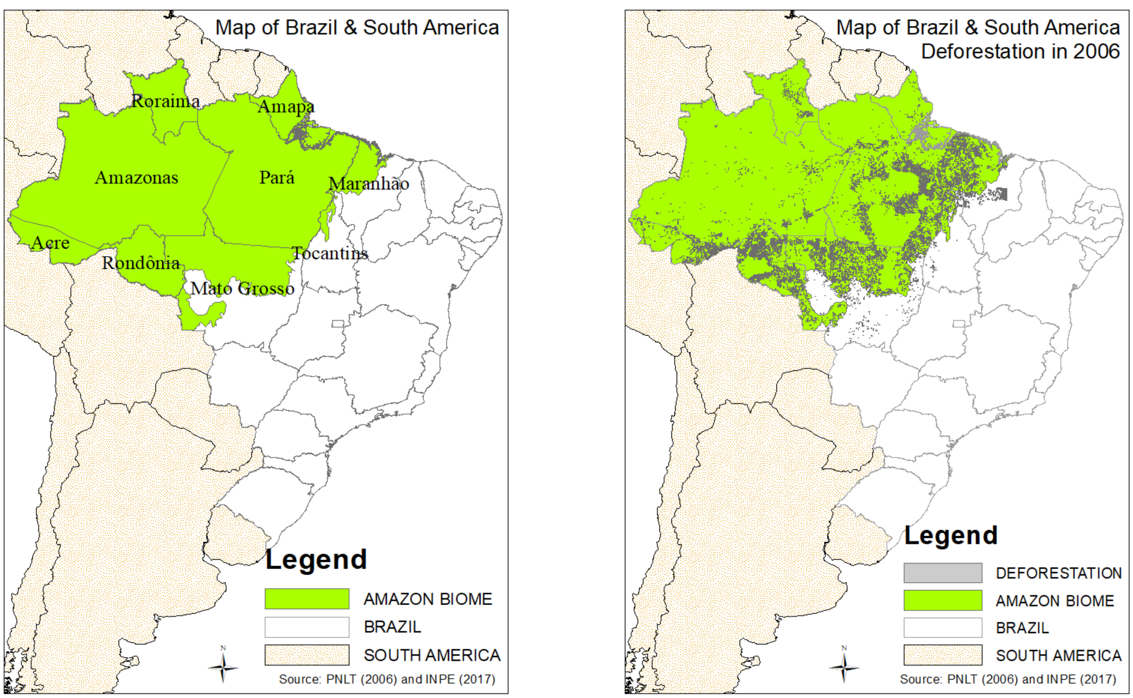

Figure 1: Map of Brazil and Cumulative Deforestation in the Amazon, 2006

Private land accounts for around 18 percent of the Amazon, but the proportion varies depending on the region: it occupies 45 percent of the *South Amazon*, 19 percent of the *Eastern Amazon*, and

[16] Assunção, McMillan, Murphy, and Souza-Rodrigues (2018) estimate the impact of the blacklist policy on deforestation (and investigate how such policy could be further refined). The average annual deforestation between 1988 and 2006 was 1.8 million hectares. It stabilized after the policy changes, the average annual rate falling to 0.6 million hectares during the period 2009–2017 (INPE, 2017).

4.5 percent of the *Western Amazon*.[17] Conservation Units and Indigenous Reserves accounted for approximately 44 percent of the Amazon by 2010 (INPE, 2017). Approximately 38 percent of the region consists of unprotected public land.[18]

Most of the private land is used for pasture (about 49 percent) and the grazing cattle are used mainly to produce beef. Brazil has the second largest number of cattle in the world (232 million cattle in 2017, of which approximately 40 percent were reared in the Legal Amazon), and it is the third largest producer of beef. Ten percent of the private properties in the Amazon are occupied by crops. Soy is the most important crop (occupying about 22 percent of the crop area), followed by corn (11 percent), manioc (11 percent), rice (8.4 percent) and other types of bean (4 percent). Brazil also is the second largest exporter of soybeans in the world. Cropland has expanded lately in the South Amazon, where most of the deforestation has been occurring. Forests occupy about 37 percent of private land, including forests that are not being exploited and managed forests. Among the forestry products, the most important in terms of the value of production in 2006 were açaí, an Amazonian fruit (41 percent) and timber (39 percent).[19]

There are important differences between logging activities and agricultural production. First, logging activities cause *forest degradation* (i.e., forest cover that, while not intact, has not yet been totally removed), while agriculture causes *deforestation* as it requires complete land clearing. Second, although annual forest degradation and annual deforestation areas are comparable in magnitude (INPE, 2017), agriculture production is the activity that most affects total forest cover in the long term. Forest degradation appears to occur mostly as a result of a single selective logging event and is typically associated with low-intensity forest damage (Pinheiro et al., 2016). Signs of degradation may not be visible even one year after being detected in the satellite imagery. Deforestation, in contrast, accumulates over time: the total accumulated deforested area in 2006 was 14 percent of the Legal Amazon in 2006 (INPE, 2017).[20]

---

[17] The *South Amazon* comprises the states of Rondônia and Mato Grosso; the *Eastern Amazon*, the states of Pará, Amapá, part of Tocantins and part of Maranhão; and the *Western Amazon*, the states of Amazonas, Acre and Roraima. The *Legal Amazon* is an administrative area in the northern part of Brazil that includes the nine states indicated in Figure 1.

[18] The unprotected public land can still be occupied and claimed by squatters. Despite this fact, most farmers have land titles (85 percent), and the proportion of farms with no land titles is higher among small landholders (20 percent). Small landholders in the current paper are those who own farms less than 5 hectares in size.

[19] The production of soybeans and corn is located mostly in the South Amazon and is directed to international markets. Manioc, rice and beans are consumed domestically, with manioc being more concentrated in pristine areas, possibly for subsistence. The logging industry is located along the South and Eastern Amazon and it directed 36 percent of its production to international markets in 2009 (data sources: USDA (www.fas.usda.gov/psdonline/), the Brazilian Agricultural Census for 2006, and the Municipal Livestock Production Survey – both produced by the Brazilian Institute of Geography and Statistics).

[20] Based on remote-sensing data for 2008, Almeida et al. (2016) estimate that approximately 90 percent of the total deforested area in the Brazilian Amazon is used for agriculture (pasture and cropland), 3 percent corresponds

3.3 Transportation Network

Figure 2 presents key geographic information relevant to the analysis. The top left panel of Figure 2 shows the navigable rivers and main ports. Rivers have always been important in the Amazon, especially in the western region, where they are the only viable means of transportation for the local populace. The top right panel of Figure 2 presents the railway network and the Amazonian state capitals. Railroads are not very prevalent in Brazil, are concentrated in the southeast, and mainly connect into ports. The main ports in the country are also located in the southeast, the most important being the Port of Santos and the Port of Paranaguá. Not only is the infrastructure of these ports better than in the rest of the country; the roads linked to them are also of higher quality, making them a more attractive option for exporters than the ports in the north.

The bottom left panel shows the location of roads, distinguishing paved from unpaved. Most of the roads in the Amazon are unpaved (89 per cent, according to the Ministry of Transportation). The few paved roads in the region tend to connect the main state capitals. The bottom right panel combines the transportation network and the deforested area in 2006, making the spatial correlation between them apparent.

3.4 Legislation and Penalties for Illegal Deforestation

If a farmer wants to clear a fraction of his land, he needs to hold many licenses and authorizations, including a detailed plan of management that must be approved by the state and national environmental protection agencies. The requirements are costly and time-consuming to fulfil, and may take several months to be approved (Hirakuri, 2003). Sanctions for forest-related violations include fines ranging from US$ 2,300 to US$ 23,000 per hectare, the seizure of products and equipment, and the suspension of production activities. The fines are extremely high for the typical farmer, given the average gross revenue per hectare – only US$ 120/ha according to the Agricultural Census of 2006.

There is evidence that the legislation has not been fully enforced. For example, between 2005 and 2009, IBAMA (the Brazilian Environmental Protection Agency) applied 24,161 fines totaling about US$ 7.34 million, but the revenues collected from these fines were only 0.6 percent of their total value (TCU, 2009). Perhaps more importantly, the proportion of deforested area (according to satellite images) that received fines was small before 2006: approximately 0.15 percent in 2003,

to mining, urban areas, and others, and the remaining 7 percent is unobserved (i.e., areas whose land cover cannot be interpreted due to cloud shade or smoke from recently burned areas).

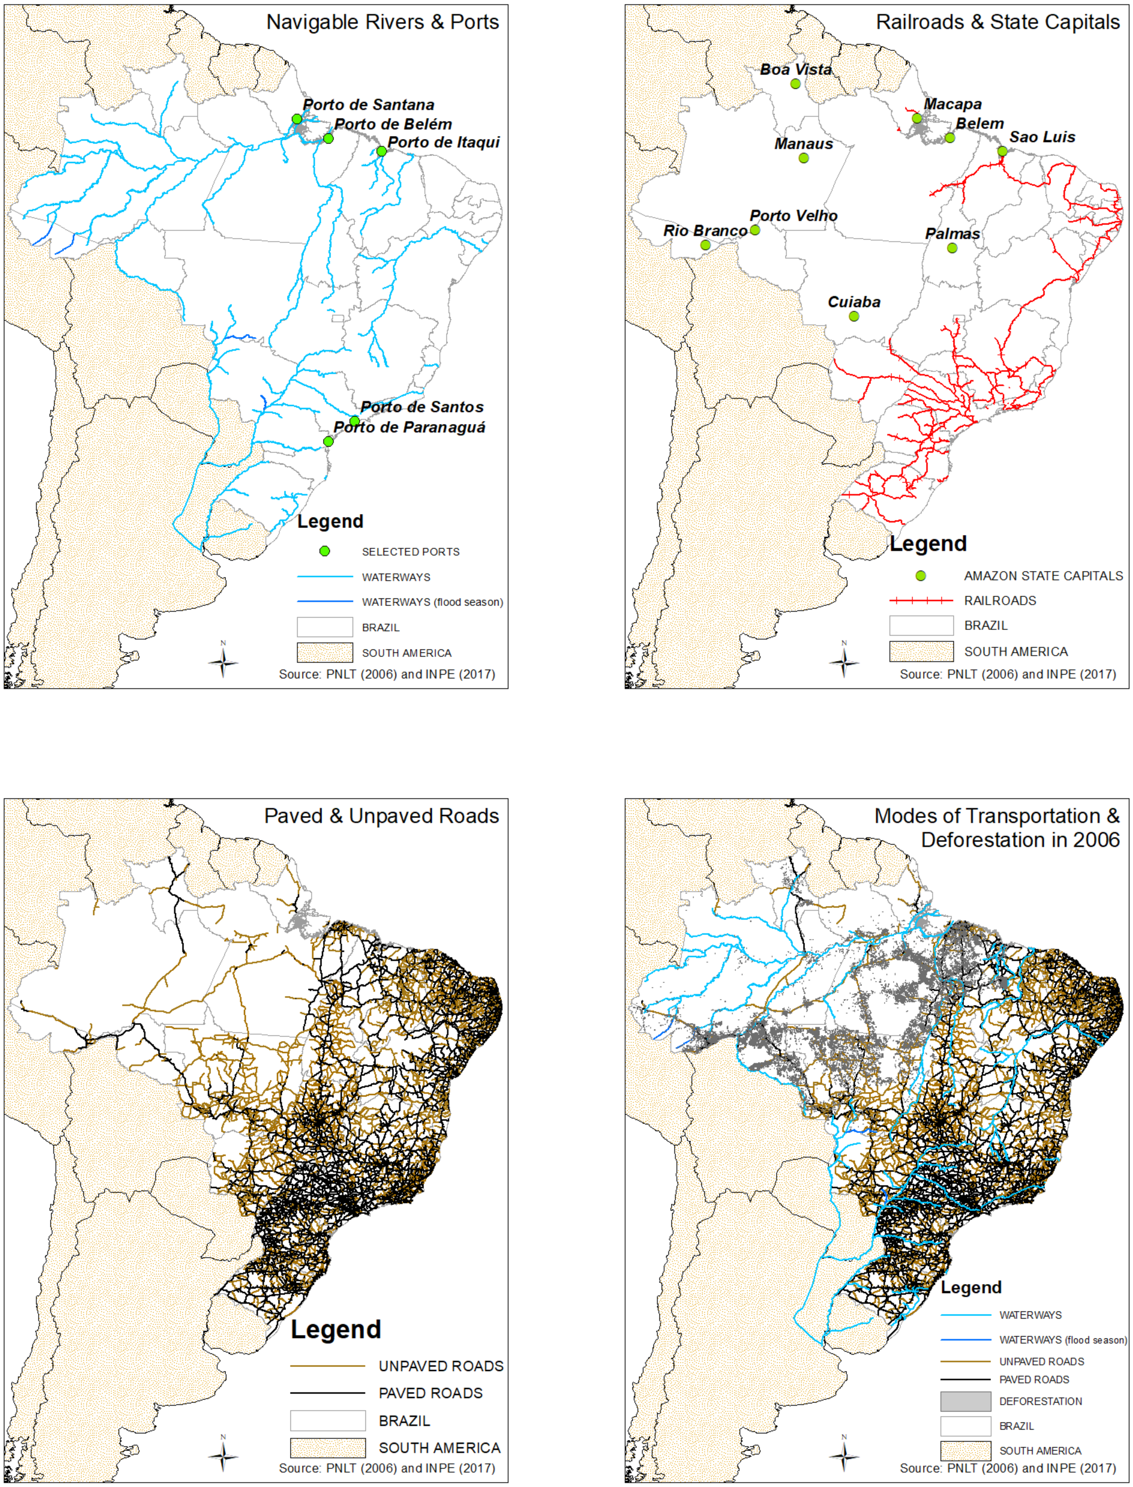

Figure 2: Transportation Network and Deforestation

0.1 percent in 2004, 1.2 percent in 2005, and 7.9 percent in 2006.[21] Brito and Barreto (2006) provide related evidence: only a small fraction of court cases involving environmental violations led to offenders being found liable (2 percent in a sample of 55 cases in Pará State during 2000–2003).

Historically, the expected cost of punishment seems to have been low (at least up to 2006), in spite of the recent increase in monitoring efforts after 2004. One might therefore expect farmers to have slashed-and-burned to clear the land without having authorization. (IBAMA staff estimate that more than 90 percent of deforestation is illegal.)[22]

4 Model and Estimation

In this section, I present a stylized model to guide the empirical analysis. Before setting out the model details, several remarks about the general formulation are in order. First, I consider a continuum of farmers making binary choices (to deforest or not) aggregated up to the municipal level, given that the available data are already aggregated up to that level.[23]

Second, I focus on landowners' choices *within* private properties. From a policy perspective, it makes little sense to tax (or pay for) land that no one owns, and the '80 percent rule' does not apply to public land – deforestation of public land is an important problem in the Amazon, but one that I do not investigate here. Third, I split the sample into different farm sizes and conduct the analysis separately for each sub-group. Doing so allows for diminishing (or increasing) returns to agricultural land that may affect farmers' valuations. It may also be informative for policy makers: to the extent that policy makers view payment programs as a way to reduce poverty, they may want to adjust payments to small landholders.

Fourth, although land-use taxes and payment programs differ in several respects (including practical implementation issues and the distribution of conservation costs), they share the same predictions for land use decisions in the present context, as will be clear below. For this reason, I lump them into one policy and refer to them simply as "taxes," unless stated otherwise.

Next, I present the details of the model, and describe my identification strategy.

[21]Estimates provided by IBAMA. The fraction of deforested areas where fines were levied increased substantially after 2006: 49 percent in 2007; 44 percent in 2008; 51 percent in 2009, and 24 percent in 2010. It seems to be the result of redoubled government efforts to slow down deforestation (see Assunção et al., 2013).

[22]This is based on informal conversations with IBAMA staff. IBAMA does not have official numbers because the state agencies responsible to supply the information to the national system do not provide data on regularized deforested areas.

[23]The aggregated nature of the data prevents me from considering local neighbor interactions. See Robalino and Pfaff (2012) for an analysis of the importance of neighbor interactions based on detailed micro data for Costa Rica.

4.1 Model

Take a parcel of land i that belongs to a farm of size s located in municipality m. Assume there is a continuum of such parcels, and for each parcel, a farmer decides whether or not to clear it for agriculture. Let P_{ims} be a vector of input and output farmgate prices and X_{ims}, a vector of other determinants of land use (e.g., productivity factors). Define $\Pi^a(P_{ims}, X_{ims})$ as the expected discounted present value of current and future profits obtained by using the parcel for agriculture, and $\Pi^f(P_{ims}, X_{ims})$ as the corresponding value obtained from leaving the plot as managed forest. The agricultural value Π^a incorporates upfront conversion costs and expected penalties for illegal deforestation (when applicable). The forest value Π^f includes profits from forest products (e.g., wild fruits and timber) and the option value of deforesting in the future, as well as possible non-pecuniary benefits. Let Y_{ims} equal one if the plot i is cleared and zero otherwise. Then, plot i is deforested when its agricultural value exceeds its forest value:

$$Y_{ims} = 1\left\{\Pi^a(P_{ims}, X_{ims}) > \Pi^f(P_{ims}, X_{ims})\right\},$$

where $1\{.\}$ is the indicator function.[24]

The vector of determinants of land use can be decomposed as $X_{ims} = (X_m, U_m(s), \varepsilon^x_{ims})$, where X_m is a municipality-level vector of observed productivity shifters, such as soil quality and other agroclimatic conditions, as well as government monitoring efforts to deter illegal deforestation; $U_m(s)$ is a municipality-level unobserved productivity shock; and ε^x_{ims} captures farmers' unobserved idiosyncratic abilities, effort and the deviations from both X_m and $U_m(s)$ within m. Because the empirical analysis is conducted separately for each farm size, it is possible to allow the unobservable $U_m(s)$ to be indexed by s, which allows for a richer model than the usual municipality random effects model. That is, a municipality may be suitable for agriculture for large farms but less so for small landholders.

I assume farmers are price takers, that all production is sold in nearby markets or exported directly, and that a no-arbitrage condition holds. These assumptions imply that local prices are determined by the international price minus the transportation cost to the nearest port, i.e., $P_{ims} =$

[24] One may interpret Π^a and Π^f as choice-specific value functions of a fully dynamic model. A structural dynamic model would separately identify and estimate the different components of Π^a and Π^f (current payoff, conversion costs, continuation values, and so on); see, e.g. Scott (2013). Given that I do not estimate such a model, the estimated parameters of the choice-specific values are not invariant to certain policy changes. While the framework allows me to estimate the effects of permanent changes in these values, it cannot estimate, for instance, the impact of a change in the volatility of timber prices, as this affects continuation values (in particular, the option value to deforest) in a way that the 'reduced-form' payoff function cannot capture in cross-sectional data. I am grateful to the Editor and an anonymous referee for this observation.

$\overline{P} - TC_{ims}$. The transportation cost TC_{ims} can in turn be decomposed into two parts. First, the cost to transport a product from the municipal seat to the nearest port is denoted by TC_m; a proxy for this variable is observed in the data. Second, the deviation of the farm's transportation cost to TC_m is denoted by ε^t_{ims}, and is unobserved by the econometrician but observed by the farmer.[25]

Although the costs of transporting different products may not be equal across products, they are likely to be proportional: all products use the same transportation network and reach the same ports (under the no-arbitrage condition). Therefore, the transportation costs of different products should be highly collinear, which makes it difficult to separately identify their impacts on deforestation. I therefore proceed with a single measure for transportation costs to reflect differences in local prices. The exact proxy for TC_m is explained in Section 5.[26]

The existing literature typically projects the difference between Π^a and Π^f on the municipal-level variables $(TC_m, X_m, U_m(s))$ and collapses all individual heterogeneity into a single scalar, ε_{ims}, which is assumed to follow an extreme value distribution. Let $Y_m(s)$ be the aggregated share of agricultural land within farms of size s in municipality m. The resulting logit model can be estimated after taking the differences in log shares as:

$$\log\left(\frac{Y_m(s)}{1 - Y_m(s)}\right) = X_m\beta_s - \alpha_s TC_m + U_m(s), \tag{1}$$

where the coefficients can be different for different farm sizes.[27] The typical exercise in the literature estimates equation (1) using OLS (Pfaff, 1999).

My procedure builds on the standard approach, extending it in two ways. First, I use quantile, instead of mean, regression. Specifically, I estimate the quantile model:

$$\log\left(\frac{Y_m(s)}{1 - Y_m(s)}\right) = X_m\beta(U_m(s)) - \alpha(U_m(s)) \times TC_m, \tag{2}$$

[25] As discussed in Section 3.2, the most important products in the Amazon are exported (with the exception of manioc). Informal conversations with farmers in the industry (particularly, in the National Association of Soybeans Producers, as well as beef producers) suggest that the price farmers effectively receive is the price at the port minus the costs to transport the product, minus a margin for intermediaries. Given that information about intermediaries is difficult to obtain, I leave an investigation into their potential impact on deforestation for future research.

[26] Castro (2003) investigates freight values in Brazil for the period 1997–2003 and finds evidence that bulk products and sacks have similar transportation costs (because costs depend mostly on the weight of the product, and not on the type of packaging).

[27] Given the linear specification, the international prices $\overline{P}$ are incorporated into the constant term, and the idiosyncratic shocks ε_{ims} are composed of $\varepsilon^x_{ims} - \alpha_s\varepsilon^t_{ims}$. The variance of ε_{ims} may differ for different farm sizes s, which in turn may affect the scale of the estimated model parameters; this is a common feature of discrete choice models. Note that farm size, s, is not an explanatory variable in equation (1): I do not try to explain deforestation patterns by exogenously varying the size of the farms. Although the endogeneity of farm sizes has been discussed extensively in the literature, particularly in the literature that estimates impacts of cultivated agricultural area on rural productivity (see, e.g., Foster and Rosenzweig, 2011) and the references cited therein), potential endogeneity of farm sizes does not arise when estimating equation (1).

where $U_m(s)$ is now normalized to have a uniform distribution on $(0,1)$. Equation (2) is a random coefficients representation of the quantile function. More specifically, following Chernozhukov and Hansen's (2013) terminology, the function $u \mapsto X_m \beta_s(u) - \alpha_s(u) \times TC$ is the "quantile treatment response" (QTR) function. The QTR function is strictly increasing and continuous in $u \in (0,1)$, so that a location more suitable for agricultural activities is associated with a higher value of u. The coefficients on X and TC can depend arbitrarily on both the farm size s and the quantile u, which implies heterogeneous effects on deforestation. This flexibility relaxes both the single-index restriction and the logit assumption in determining the shape of the demand for deforestation. In light of this, I change notation slightly from now on and denote the coefficients by $(\beta_{su}, \alpha_{su})$.

Incorporating heterogeneous effects on deforestation is important because, as noted previously, two observationally equivalent locations (in term of X and TC) may respond differently to changes in transportations costs. A highly deforested location may be well-suited to agriculture in terms of unobservables so that increases in transportation costs may not reduce the cultivated area significantly; better preserved locations, in contrast, may be more sensitive to changes in TC. The impacts of transportation costs therefore likely differ at the upper and lower tails of the conditional distribution of deforestation. An implication of this is that the relative value of agricultural land also likely differs at the upper and lower tails of the conditional distribution: some locations may be more sensitive to taxes than others and so may face higher costs (lost surpluses) when taxes are imposed. In turn, these differences may lead to non-trivial impacts on the aggregated costs of policy interventions.

A second extension addresses the potential endogeneity of roads and measurement error in transportation costs. Specifically, costs to ports are instrumented for using straight-line distances to the main destinations. In the next subsection, I discuss reasons why transportation costs to the nearest port should be instrumented for and the conditions under which straight-line distances are expected to be valid instruments. Here, I simply note that, because the transportation costs are not exogenous, the conventional quantile regression estimator is inconsistent for estimating the QTR function. For this reason, I estimate equation (2) using the instrumental variable quantile regression (IVQR) estimator proposed by Chernozhukov and Hansen (2008).

Demand for Deforestation. Now I turn to the central component of the empirical framework: the demand for deforestation. Taking the logistic function $h(x) = \exp(x)/(1+\exp(x))$, the share of agricultural land for farms of size s in municipality m at a given quantile u is given by $h(X_m \beta_{su} - \alpha_{su} TC_m)$. The effect of raising the private value of forested land (relative to agricultural land) by US$ t per hectare on farmers' land-use decisions is given by

$$Y_m(s,t) = h\left(X_m\beta_{su} - \alpha_{su}\left(TC_m + \frac{t}{q_m(s)}\right)\right), \tag{3}$$

where $Y_m(s,t)$ is the counterfactual share of agricultural land, and $q_m(s)$ is the quantity (in tons) of agricultural output sold per hectare. I define the *demand for deforestation* for farms of size s in municipality m as the product of the total area they occupy and the counterfactual share $Y_m(s,t)$. The total demand then aggregates over s and m.[28]

Two aspects of the demand function warrant discussion. First, following Chernozhukov and Hansen's (2013) terminology, this demand satisfies the "rank invariance" assumption. Rank invariance is a common assumption in the applied literature and preserves the intuitive notion that, conditional on observables, a relatively highly deforested location in the data (i.e., a location associated with a high rank u) remains a relatively highly deforested location under alternative counterfactual policies (i.e., it preserves the rank u).[29]

Second, because the data are aggregated up to the municipality level, and because there are hundreds of products being produced in the Amazon, some care is needed in defining $q_m(s)$. I selected the most representative products (those discussed in the Subsection 3.2) and constructed a local productivity index (in which the weights are the proportions of the area utilized for each product). The underlying assumption here is that once the land is cleared for agriculture, it is used in fixed proportions for pasture and for the main crops: the proportions are allowed to differ across municipalities, but they are fixed within-municipality.[30]

[28]This structure can be extended to consider multinomial choice models (e.g., forest *vs.* crops *vs.* cattle ranching), and in principle, variables other than TC_m can be used to capture variation in private values (e.g., data on agricultural potential to identify crops *vs.* livestock). Yet, in order to be useful, any such variable must satisfy three requirements: (a) it must affect farmers' decisions significantly, i.e., it must have a coefficient that is different from zero; (b) it must be measured in dollars in a way that can be converted into the appropriate units; and (c) it must be able to capture persistent differences in the private returns to land use. None of the variables in X_m in the present data set satisfies all three requirements.

[29]Empirical applications that estimate QTR functions under the rank invariance assumption include Hausman and Sidak (2004), Chernozhukov and Hansen (2008), and Lamarche (2011).

[30]I also consider a second index that includes only the main crops and ignores pasture. This severely restricts the substitution patterns among land uses within municipalities and so provides a conservative upper bound on the demand for deforestation. More general substitution patterns could be recovered by exploiting choice-specific variables that shift the value of each type of land use independently of the value of the other options (Berry and Haile, 2014), but there is no variable satisfying such a requirement in the present data set. I therefore choose to be agnostic in terms of the way the agricultural area is divided when estimating the impacts of TC_m on deforestation, and report the results for both indices. I also examined whether the indices q_m respond to TC_m, and found no such evidence: farmers seem to adjust the extensive margin (land use), not the intensive margin (yields), which is consistent with Roberts and Schlenker (2013). (See the Supplemental Material.)

4.2 Identification Strategy

Endogeneity. There are several reasons why one needs to instrument for transportation costs in land-use regressions. First, these costs are likely measured with an error. The proxy for transportation costs is defined here as the minimum unit cost (US$/ton) to transport one ton of goods to the nearest port using the most cost-effective route. It is a common proxy used in the literature (Allen and Arkolakis, 2014; Donaldson and Hornbeck, 2016), but might not provide an accurate measure of the real costs that farmers incur and so is potentially mismeasured. If the measurement error is classical, it may induce an attenuation bias in the OLS estimates.

Second, previously deforested regions may have a higher demand for improvements in local infrastructure conditions, including more and better roads, which leads to reverse causality in cross-sectional data. Third, roads may have been built in response to profitable agricultural conditions. As a common example, unobservable (to the econometrician) soil quality for agriculture in a given location may induce both deforestation there and road construction to access that location. Both the simultaneity and the omitted variable problems may lead the OLS estimates to overstate the impact of transportation costs.

In the present case, the omitted variable problem does not necessarily lead to an upward bias. As mentioned in Subsection 3.1, early occupation of the Amazon was based on the extraction of rubber, and regions that are well-suited for the growth of rubber trees (*Hevea brasiliensis*) are not necessarily well-suited for agriculture. The soil quality in the Amazon is actually poor for agriculture in most regions (see Section 5). Because good navigable rivers might have been used and unofficial roads might have been built in the past to access valuable trees, but more recent roads (and any improvement in them) may have been directed to agricultural regions, the direction of the bias in the OLS estimator is not clear *ex-ante* (Pfaff et al., 2009).

Instruments. I use straight-line distances to the nearest port and to the nearest state capital as instruments for transportation costs. I now discuss (a) why one should expect straight-line distances to serve as strong instruments – this is testable, and (b) under what conditions one should expect the instruments to satisfy an exclusion restriction condition.

First, it is evident that distances to the nearest port should correlate with transport costs to ports. Furthermore, to the extent that state capitals are connected to better transportation infrastructure (see Subsection 3.3), a location close to a state capital should have lower costs (*ceteris paribus*) of reaching the ports. Therefore, the distance to the nearest capital should also be positively correlated with transportation costs.

The conditions under which the instruments satisfy the exclusion restriction are more involved. I start by following the reasoning presented in Chomitz and Gray (1996): because locations of major towns – in the present case, ports and state capitals – were determined by geography and historical factors long before the expansion of the roads in the 1970s, I can construct an exogenous network of roads by linking the major centers by straight lines.[31] The distances computed using the virtual network should be correlated with transportation costs to ports, because the location of towns creates links between the major centers, but not the precise routing. Similar to the ports and state capitals, most municipal seats in the Amazon were established long before the occupation of the Amazon. Specifically, they were established by the late 1800s and early 1900s and – as discussed earlier – were not necessarily located in areas where agricultural activity was more valuable. It is conceivable therefore that the virtual road network is exogenous to the agricultural activities that took place in the Amazon after the 1970s. Given that using the entire virtual network and computing straight-line distances directly to the main destinations provides the same information, I opted for the simpler solution.

Although the virtual road network can be viewed as exogenous to recent agricultural activities, it is still possible that the straight-line distances correlate with factors that affect farmers' decisions to deforest. It is therefore necessary to control for those factors. As discussed in the previous subsection, farmers' decisions depend on productivity shifters, government monitoring efforts, and on farmgate input and output prices. Once those factors are taken into account, straight-line distances should not influence farmers' choices. In the application, I control for differences in productivity using measurements of soil quality and various agroclimatic variables; I consider two proxies for monitoring effort (which I discuss in more detail below); and variation in local prices is explained by variation in transportation costs to the nearest port – at least for tradable goods.

The instruments may be invalid if there are inputs or outputs whose prices are not fixed in the international market. In such a case, local market conditions may affect local prices and correlate with straight-line distances to the main destinations. Consider local labor markets, for example: wages may have to increase in the municipalities that are found further away from the nearest capital, all else equal, to compensate workers for working far from more desired locations. Municipalities further away from the capital may deforest less than a location close to the capital because of wage differences. If the wage differences are not controlled for in the regression, and

[31] All state capitals and ports used in this paper were established before or during the 19th Century, except for Porto Velho (founded in 1907).

correlate with the instruments, then the proposed instruments are invalid. A similar problem may arise if there are other non-tradable inputs and outputs.

To address this possible issue, I include in the regressions factors that shift local demand and supply for non-tradable inputs and outputs that may correlate with straight-line distances. Specifically, I include local population, the presence of power plants (mainly hydroelectric facilities), and local mining. The local population shifts the supply of labor and increases the demand for non-tradables, and power plants and mining shift both the demand for labor and non-tradables. Because inter-related local markets may create spatial dependence in terms of farmers' decisions across municipalities, I also include spatially lagged local demand and supply shifters among the covariates.[32]

To capture government monitoring efforts, I consider the number of fines issued in a municipality for environmental infractions, and the distance of the municipal seat to the nearest IBAMA office. Data on fines are publicly available after 2002 (see Supplemental Material). I consider the cumulative number of fines over 2002–2005, which may capture the recent increase in monitoring. The distance to IBAMA is intended to capture the possibility that monitoring and punishing farmers for illegal deforestation is more difficult for farms located in more pristine areas, i.e., the farther away the farm is, the less monitoring there will be, and so the more incentive farmers will have to deforest. Ignoring this possibility may lead to underestimating the impacts of transportation costs.

There are two other factors that potentially affect farmers' decisions to deforest and that may correlate with distances to ports and to state capitals: proximity to protected areas (PAs), and the potential lack of property rights. PAs may have spillover effects that influence the value of nearby forestlands. Herrera (2015) presents evidence for the Brazilian Amazon showing that PAs may indeed reduce deforestation in surrounding areas. To allow for this potential mechanism, I include the distance to the closest PA as well as its spatially lagged value in the land-use regressions.

The potential lack of property rights is another possible determinant of land use. Historically, farmers had incentives to deforest as a way to secure land tenure (Andersen et al., 2002). It is conceivable then that the further farms are from a state capital, the less secure the land rights are, and so the more incentive farmers have to deforest – that is, the greater the distance, the larger the deforested area. To address this issue, I include a proxy for property rights in the regressions. In particular, in the Brazilian Agricultural Census, the best proxy for property rights

[32]Assunção et al. (2016) provide evidence that power plants affect land use decisions in the neighborhood of power plants in Brazil. In the Supplemental Material, I present a detailed discussion of other potential sources of spatial dependence, as well as the results of several robustness exercises.

is the proportion of private land with a land title – presumably, the higher the tenure security, the larger is the proportion of land with land titles.

While these control variables mitigate concerns related to the validity of the instruments, they are not by themselves without potential issues. The siting decisions of power plants, IBAMA offices, and PAs may depend on deforested areas. Similarly, the number of fines and the tenure security proxy may suffer from simultaneity problems in the land-use regressions. Despite these issues, I provide empirical evidence, as well as an extensive discussion, in the Supplemental Material showing that these are not likely to be first-order concerns in the present context.[33]

5 Data

In this section, I explain how deforestation and transportation costs are measured. Then I present relevant summary statistics. The set of covariates I use in the regressions includes: soil quality, temperature, rainfall, altitude, slope, local population, local mining, local power plants, total number of fines (up to 2005), distance to IBAMA offices, distance to protected areas, the share of private land with land titles, and spatially lagged variables for the local demand and supply shifters as well as for distance to protected areas. A detailed description of the variables construction is provided in the Supplemental Material.

5.1 Dependent Variable: Deforestation

The land use classification in the Brazilian Agricultural Census of 2006 is divided into several categories that I aggregated up to two: agricultural and forested land. Agricultural land includes pasture and crops, while forested land aggregates managed forests and forests that are not currently being exploited. The farm size categories considered here are: (a) small farms (those with less than 5 hectares); (b) small-to-medium farms (those with areas between 5 and 50 hectares); (c) medium-

[33]To summarize, the location of power plants is plausibly exogenous to farmers' land use decisions, given that their location depend mainly on costs of hydropower dam construction, which in turn depend on topographic factors, most of which are controlled for in the land-use regressions. The IBAMA offices are located throughout the entire Brazilian Amazon (and the correlation between distance to IBAMA and distance to ports is small and insignificant). Until 2006, the majority of PAs were located on land facing lower deforestation pressure (as documented by Pfaff et al. (2015)). While the inclusion of fines in the regressions leads to the usual simultaneity problems, ignoring them may lead to omitted variables biases. The direction of the omitted variables bias depends on the correlation between fines and the instruments for TC_m. In the data, this correlation is small, which suggests that omitting fines in the regressions may result in very small biases. Indeed, the estimated coefficients on fines have the expected sign, but they are not statistically significant, their magnitudes are small, and the coefficients on other regressors are not substantially affected by the inclusion or the exclusion of fines. (This is consistent with the interpretation that the increased monitoring efforts were too recent to affect the total accumulated deforested land substantially by 2006.) Similar results are observed related to the inclusion and exclusion of the tenure security proxy. See a detailed discussion of these issues in the Supplemental Material.

to-large farms (those with areas between 50 and 500 hectares); and (d) large farms (those with more than 500 hectares).[34]

5.2 Endogenous Regressor: Transportation Costs

As a proxy for transportation costs, I use the minimum unit cost (US$/ton) to transport one ton of goods to the nearest port. This cost is calculated by combining information from the Brazilian transportation network for 2006 produced by the Ministry of Transportation for the National Highway Plan (PNLT, 2006), and freight rate data collected by SIFRECA (the Freight Information System). The implementation makes use of ArcGIS cost distance tools (see Allen and Arkolakis (2014) for an excellent description of this type of algorithm). The calculation divides the entire country into cells, and the cost to travel over each cell in the grid depends on whether it contains a segment of road (paved or unpaved), railroad, navigable river or no transportation mode. To assign the travel costs for each transportation mode, I use the freight rate data collected by SIFRECA. I adjust for road quality in the calculations using the Vehicle Cost Module of the World Bank's Highway Design Model (HDM-VOC-4), which is designed to calculate unit road user costs for different types of road surface. The optimization routine in ArcGIS determines the least cumulative cost path from each origin to the nearest destination. Because almost all the information I obtained from SIFRECA for the Amazon corresponds to costs of transporting soybeans, the proxy TC_m measures the minimum cost of transporting one ton of soybeans.

5.3 Summary Statistics

Table 1 presents summary statistics. There are 523 municipalities in the data set. Farms occupy almost 40 percent of the municipal area on average; and the average fraction of private land used for agriculture is 65 percent. The cost to transport one ton of soybeans from the South Amazon (the region where the soybean is primarily grown) to the Port of Santos is 30 percent of the price of soybeans at the port – a significant cost for farmers. As expected, high levels of temperature and rainfall are prevalent, and most of the soil is of poor quality for agriculture. The average number of fines per municipality increased steadily over the period 2002–2006, consistent with the numbers

[34] There is an alternative to using census data – namely satellite sensor data. However, remote-sensing data do not distinguish between deforestation on private and public land. Because the distinction is important in my analysis, I opted to use the census data. In the Supplemental Material, I present evidence showing that the census and the satellite data are consistent with each other in a restricted sample of municipalities in which the share of private land in the municipality area is large and the share of clouds/unobserved areas is small in the satellite images. That is, the two deforestation measures are consistent when they are most comparable. (Thanks to an anonymous referee who suggested these comparisons.)

Table 1: Summary Statistics

Variables	Mean	Std Dev	Min	Max
Land Use				
Number of Farms per Municipality	1225	1169	37	11544
Share of Farm Area in Municipal Area	0.39	0.27	0	0.98
Share of Deforestation in Farm Area	0.65	0.2	0.12	1
Transportation				
Cost to Port (US$/ton)	41.4	33.5	0	163
Distance to Port (km)	866	718	0	2627
Distance to Capital (km)	314	218	0	900
Agro-Climatic Variables				
Temperature (^{o}C/year)	26.5	0.56	25	27.5
Rainfall (mm/year)	184	32	116	272
Altitude (meters)	118	126	0	920
Slope (degrees)	0.54	0.39	0.03	3.03
Share of Good Soil	0.05	0.17	0	0.99
Share of Good/Medium-Quality Soil	0.08	0.22	0	1
Share of Medium-Quality Soil	0.04	0.18	0	0.99
Share of Low-Quality Soil	0.51	0.38	0	1
Share of Unsuitable Soil	0.31	0.35	0	1
Local Demand-Supply Shifters				
Number of Local Mines	0.81	2	0	22
Number of Local Power Plants	0.04	0.2	0	1
Local Population (thousands)	31	96.4	1.92	1405
Monitoring				
Number of Fines in 2002	2.09	4.85	0	41
Number of Fines in 2003	5.58	15.8	0	236
Number of Fines in 2004	7.47	25.2	0	449
Number of Fines in 2005	7.88	18.1	0	192
Number of Fines in 2006	10.7	23.9	0	178
Cumulative Number of Fines 2002–2005	23	54.6	0	735
Distance to IBAMA (km)	74.8	56.4	0	303
Other Variables				
Distance to Protected Areas (km)	40.2	32.4	0	167.7
Share of Land with Land Title	0.93	0.12	0.01	1
Carbon Stock				
Carbon Stock in Forested Areas (tC/ha)	188.5	68.24	49	339.6
Carbon Stock in Deforested Areas (tC/ha)	107.9	41.14	9.48	246.3

Notes: The unit of observation is a municipality in the Amazon.
There are 523 municipalities in the data.
Source: author's calculations.

provided in Subsection 3.4 (i.e., the small proportion of deforested areas that received fines) given the high levels of deforestation before 2006. The average difference between the carbon stock in forested and deforested areas is approximately 80 tons of carbon per hectare.

Table 2 provides information regarding the different farm sizes. The numbers in the cells are sample averages across municipalities. The concentration of land is clear from the table: despite the fact that large farms account for a small proportion of the total number of farms (6 percent), they occupy about 52 percent of all private farmland, while small farms account for 23 percent of the farms and occupy only 1 percent of the private land. The small landholders tend to deforest a large part of their land (90 percent), but the proportion of deforestation diminishes as farm size increases. (Recall that the existing legislation requires deforestation to be less than 20 percent of the property.)

Table 2: Summary Statistics by Farm Sizes (Sample Averages)

Statistics	Small	Small–Medium	Medium–Large	Large
Number of Farms per Municipality	302	414	353	46
Share of Total Number of Farms	0.23	0.34	0.31	0.06
Share of Total Farmland Area	0.01	0.11	0.37	0.52
Share of Deforestation in Farm Area	0.90	0.71	0.69	0.63
Local Yields (q_m)	0.70	0.42	0.41	0.27
Number of Municipalities	501	520	520	450

Notes: The unit of observation is a municipality in the Amazon. Small farms are those with less than 5 hectares; Small–Medium farms are those with areas between 5 and 50 hectares; Medium–Large farms are those with areas between 50 and 500 hectares; and Large farms are those consisting of more than 500 hectares. The numbers in the cells are sample averages across municipalities.
Source: author's calculations.

6 Effects of Transportation Cost on Deforestation

This section presents the estimated impact of transportation costs on deforestation. I first discuss the first-stage estimates for the IV approach, checking for the presence of weak instruments, then I turn to the land-use regressions.

6.1 First-Stage Regression

Table 3 presents the results from regressing transportation costs to the nearest port on straight-line distances; for brevity, the estimated coefficients on the other covariates are omitted. It is clear that both straight-line distances to ports and to the nearest capital are strong predictors of costs to ports and that there is no issue with weak instruments in this data set. Increasing the distance to the nearest port by 100 km raises the cost to transport one ton of soybeans by US$ 4 on average (which corresponds to 10 percent of the average transportation cost – see Table 1).

Table 3: First-Stage Regression

	Costs to Port
Distance to Port	0.041***
	(0.0014)
Distance to Capital	0.012***
	(0.0024)
Observations	523
F-statistic	862.39
R^2	0.92

Standard errors in parentheses, *** $p < 0.001$.

6.2 Land-Use Regressions

Next, I present results for the logit models, relating the share of deforestation to transport costs.[35] Table 4 reports the estimated coefficients for costs to the nearest port and the associated standard errors in parentheses for the OLS, 2SLS, quantile regression (QR), and the instrumental variable quantile regression (IVQR) for each farm size. The coefficients on the other regressors are omitted from the table; they are reported in the Supplemental Material, where I discuss several robustness exercises.[36]

I begin by comparing the OLS and 2SLS estimates. Recall that the typical exercise in the literature uses OLS to estimate land-use regressions. As discussed in Subsection 4.2, it is not clear *ex-ante* what the direction and magnitude of any bias from the OLS estimates would be. The OLS coefficients in Table 4 are small in magnitude and are not significantly different from zero. In addition, they predict implausible positive impacts of costs to ports on the share of agricultural land for small and medium-sized farms. When transportation costs are instrumented for using

[35] As a robustness exercise, I also drop the logit assumption and estimate a semiparametric quantile IV model using the penalized sieve minimum distance estimator proposed by Chen and Pouzo (2012). I find no significant differences between the logit and the semiparametric models. The quantile logit model appears to be sufficiently flexible for the current data set. The results of the semiparametric model are presented in the Supplemental Material.

[36] For all farm sizes, the overidentification tests do not reject the null of valid instruments. The corresponding p–values are: 0.33 for small farms; 0.31 for small–medium; 0.18 for medium–large; and 0.39 for large farms. Inference for the quantile regressions is based on the results of Chernozhukov and Hansen (2008). Currently, no estimation method that incorporates spatial correlation in the unobservables into instrumental variables quantile regressions exists. Despite this, in the Supplemental Material, I also present standard errors that are calculated based on a geographic cluster-bootstrap. The results are qualitatively similar, although the parameters are somewhat less precisely estimated.

Table 4: Land-use Regression Model by Farm Size – Coefficients on Costs to Ports

			Quantiles				
	OLS	2SLS	10	25	50	75	90
Small							
No IV	0.0116	-	0.0098**	0.0080*	0.0110**	0.0080	0.0185
	(0.0080)	-	(0.0044)	(0.0046)	(0.0052)	(0.0068)	(0.0154)
IV	-	-0.0024	0.0080*	0.0027	0.0045	-0.0003	-0.0121
	-	(0.0089)	(0.0045)	(0.0049)	(0.0060)	(0.0083)	(0.0249)
Small–Medium							
No IV	0.0022	-	-0.0034	0.0008	-0.0025	-0.0018	0.0078
	(0.0049)	-	(0.0031)	(0.0032)	(0.0031)	(0.0038)	(0.0072)
IV	-	-0.0063	-0.0049	-0.0023	-0.0053	-0.0063*	-0.0008
	-	(0.0052)	(0.0035)	(0.0035)	(0.0034)	(0.0038)	(0.0069)
Medium–Large							
No IV	0.0040	-	-0.0101***	-0.0085**	-0.0025	0.0034	0.0110**
	(0.0039)	-	(0.0033)	(0.0034)	(0.0034)	(0.0036)	(0.0050)
IV	-	-0.0058	-0.0184***	-0.0112***	-0.0079**	-0.0042	-0.0010
	-	(0.0039)	(0.0038)	(0.0037)	(0.0038)	(0.0040)	(0.0049)
Large							
No IV	-0.0023	-	-0.0113***	-0.0102***	-0.0077**	-0.0063	-0.0084
	(0.0047)	-	(0.0040)	(0.0039)	(0.0038)	(0.0045)	(0.0054)
IV	-	-0.0077	-0.0110***	-0.0109***	-0.0110***	-0.0107**	-0.0156***
	-	(0.0058)	(0.0042)	(0.0040)	(0.0038)	(0.0042)	(0.0050)

Notes: This table reports the estimated coefficients on transportation costs based on OLS, 2SLS, QR, and IVQR estimators. For each farm size, the dependent variable is the log odds ratio of the share of deforestation. The unit of observation is a municipality in the Amazon. The number of observations for small farms is 501, for small-medium farms is 520, for medium–large farms is 520, and for large farms is 450.
Standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

straight-line distances, the coefficients increase in magnitude (except for smallholders) and their signs become negative for all farm sizes. Similar to OLS estimates, however, the 2SLS coefficients are also imprecisely estimated.

Next I focus on the quantile regressions. The IVQR coefficients for medium-sized and large farms are negative, and almost all are significant and greater in absolute value than the QR coefficients. The coefficients differ across quantiles, so even after controlling for observable municipality-level variables, farms with different levels of deforestation appear to respond differently to changes in transportation costs.[37]

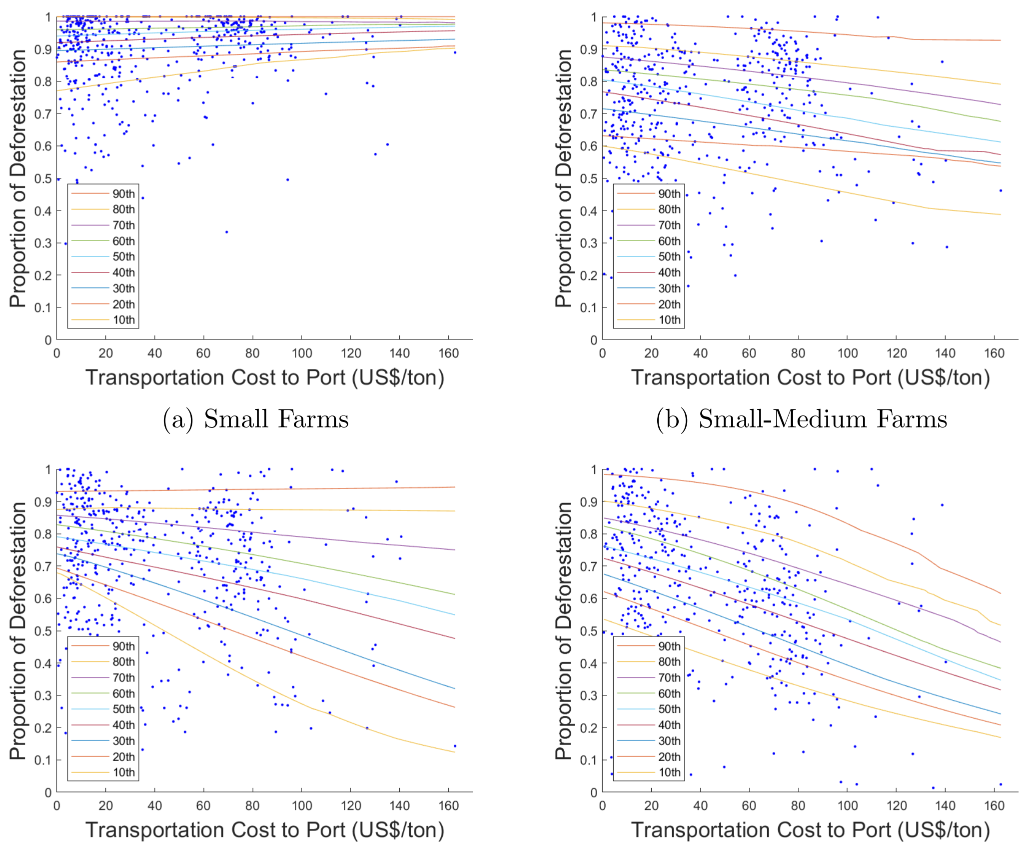

(a) Small Farms (b) Small-Medium Farms

(c) Medium-Large Farms (d) Large Farms

Figure 3: Quantile Functions – Share of Deforestation vs. Costs to Ports

The heterogeneity in responses across quantiles can be illustrated graphically. Figure 3 presents the data and the curves for the deciles $u = 0.1, 0.2, ..., 0.9$ obtained from estimating the IVQR.[38]

[37]The estimated results are similar when I use distance to port as the *only* instrumental variable for the costs to port. (See the Supplemental Material.)

[38]To compute the curves in the figure, I arrange the quantiles for each point in the data following the procedure proposed by Chernozhukov et al. (2010) to avoid quantile crossing. The model is estimated for quantiles ranging over $\{0.05, 0.1, ..., 0.95\}$ to avoid problems with extreme quantiles.

The regressors are fixed at the sample average and costs to ports vary over the observed range in the data. It is clear that small landholders are insensitive to transportation costs – almost all of their IVQR coefficients are not significantly different from zero and the lines in Panel (a) are flat for all but the lowest decile. This seems reasonable because small farms tend to be concentrated in isolated regions in the Western Amazon; they are most likely producing for subsistence and are not engaged in the broader market. As such, their decisions to deforest will be driven by the shadow value of food, and not by the costs to the nearest port. Even though they likely respond to payment programs or land use taxes, the econometric model does not seem to be well-suited for them and so my strategy most likely fails to identify their demand for deforestation. Despite these issues, the behavior of small landholders will not play a major role in environmental policies given that they occupy only one percent of private land overall.

Table 5: Marginal Effects of Transport Costs on the Share of Deforestation, by Farm Size

		Quantiles (IVQR)				
	2SLS	10	25	50	75	90
Small						
Share of Deforestation (%)	97.75	76.89	87.57	94.24	98.66	99.99
Marginal Effect	-0.053	1.421	0.298	0.245	-0.004	-0.001
Small–Medium						
Share of Deforestation (%)	78.58	49.99	61.54	74.17	85.02	93.29
Marginal Effect	-1.054	-1.230	-0.543	-1.020	-0.808	-0.047
Medium–Large						
Share of Deforestation (%)	75.72	46.64	58.87	72.13	83.41	90.80
Marginal Effect	-1.061	-4.575	-2.719	-1.581	-0.585	-0.086
Large						
Share of Deforestation (%)	69.74	38.80	50.64	65.18	77.99	89.53
Marginal Effect	-1.632	-2.605	-2.713	-2.505	-1.837	-1.466

Notes: This table reports the estimated shares of deforestation based on the 2SLS and the IVQR estimates presented in Table 4, holding covariates at the sample mean. It also shows the estimated marginal effects of transportation costs on the share of deforestation. Marginal effects are measured in percentage points and correspond to an increase of $10 per ton in transportation costs.

For other farm sizes, the upper-tail quantile curves tend to be concave, while the lower-tail quantile curves tend to be convex, which conforms to the discussion in Subsection 4.1. These results reveal that the mean-effect estimates based on 2SLS mask considerable heterogeneity. The heterogeneity in responses, according to the IVQR estimates, can vary substantially across (a) farm sizes s, (b) quantiles u (for any given s and TC), and (c) transportation costs (for any s and u).

To have a sense of the magnitudes involved, Table 5 presents the estimated marginal effects of transportation costs on the share of deforestation for each farm size and for different quantiles based on the IVQR model, as well as the mean effects estimated using the 2SLS estimator. The marginal effects are measured in percentage points and correspond to an increase in transportation costs of \$10/ton, which is approximately one third of a standard deviation in the data – see Table 1. The covariates are fixed at the sample average. Consistent with Figure 3, the marginal effects tend to be greater at lower quantiles and for larger farms. Taking the (conditional) median, note that while increasing transportation costs by \$10/ton increases the estimated fraction of deforestation on small farms by only 0.24 percentage point, it reduces the corresponding fraction of deforestation on medium-sized farms by 1–1.5 percentage points, and on large farms by 2.5 percentage points.

7 Demand for Deforestation

In this section, I present the estimated demand for deforestation on private properties. I also assess the geographic distribution of demand. I then investigate how the demand changes when holding the total number of fines after 2004 at pre-2004 levels, to separate the effects of taxes from the increased government monitoring efforts observed after 2004.

To compute the demand function, I use the IVQR estimates of the logit model together with equation (3) in Section 4 to predict the fraction of agricultural land on private properties for each hypothetical tax, for each farm size, and for each municipality in the data set.[39] Then, for each tax level, I compute the total deforested area based on the predicted share of agricultural land. By summing over municipalities, I obtain the corresponding demand for each farm size. The total demand is then obtained by summing over farm sizes.[40]

Figure 4 presents the resulting demand for deforestation for each farm size and the total demand

[39]In practice, I impose the rank invariance condition in the following way: for each farm size s and each municipality m, I compute the estimated QTR function $X_m' \beta_s(u) - \alpha_s(u) \times TC_m$ for all quantiles. I then rearrange the quantiles according to the procedure proposed by Chernozhukov et al. (2010) to avoid quantile crossing. The rank u associated with observation (s, m) is the one that minimizes the difference between the observed $\log(Y_m(s) / (1 - Y_m(s)))$ and the estimated QTR function. I hold the rank fixed when calculating the demand for deforestation.

[40]An implicit assumption in this calculation is that taxes would not change the distribution of farm sizes. Possible impacts of taxes on the sizes of farms are beyond the scope of the paper (see Subsection 8.3 for discussion).

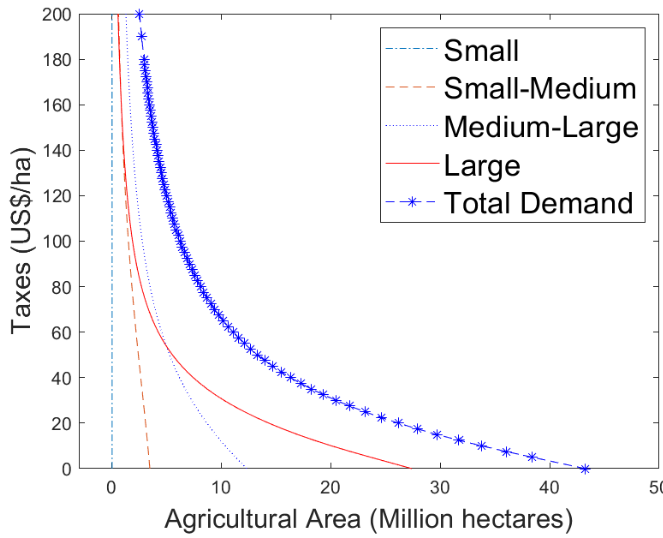

Figure 4: Demand for Deforestation, by Farm Size and Total

function. One may interpret the total demand in the following way: if the government had increased the relative value of forested land by imposing a perfectly-enforced tax that charged, say, US$ 42.5/ha of agricultural land per year, farmers would be willing to use approximately 15.5 million hectares for agriculture (20 percent of the private properties), reading off from the rightmost curve, instead of the actual 43.2 million hectares (56 percent) – the horizontal intercept of that curve. Because farmers' average gross revenue per hectare in the Amazon in 2006 was US$ 120/ha, such a tax would drive many farmers out of production.[41]

It is clear from Figure 4 that the shape of the total demand mainly comes from the demand curve for large farms. The extremely unequal distribution of land in the Amazon coupled with the response of large farms to taxes suggest that policies targeting *only* small landholders cannot promote significant conservation (recall that they occupy only a small fraction of the total land – see Section 5.3). Although taxes or payments can be effective in changing small farmers' behavior (see Simonet et al. (2018) for a recent small-scale PES program in Pará State), it is evident that payments are unlikely to reduce local poverty and deforestation simultaneously in a substantial way. This result is consistent with the analysis by Godar et al. (2014), who document that most of the deforestation in the Amazon between 2004 and 2011 occurred in regions dominated by large farms (those with more than 500 hectares).

[41]One may interpret the results in terms of expected taxes that farmers would pay instead of payments under a perfectly enforced tax.

7.1 Geographic Distribution

In Figure 5, I present the geographic distribution of deforestation under different taxes. The top left panel presents the total agricultural area computed from Census data – the darker the region, the larger the area occupied by agricultural land. The top right panel presents the counterfactual agricultural land for taxes of US$ 20/ha; the bottom left, the corresponding map for taxes of US$ 40/ha; and the bottom right panel gives the map for the US$ 100/ha tax. By comparing the top-left map (no tax) with, say, the top-right map (US$ 20/ha tax), it is possible to see which regions would be mostly affected by the US$ 20/ha land-use tax.[42]

The figure makes clear that farmers in the 'Arc of Deforestation' respond less to taxes. Even under a tax of US$ 100/ha, farmers in the South Amazon, the region where the soybean is produced, would be willing to use their land intensively – the opportunity costs of the agricultural area in this region are probably too high to *not* be used for agriculture. In contrast, forests in the central and western regions can be better preserved using taxes and so be less fragmented, which is advantageous from a biodiversity point of view.

7.2 Monitoring Efforts Post-2004

As discussed in Section 3, government monitoring increased after 2004, mostly through the use of satellite-based monitoring systems. In order to try to separate the effects of taxes from these recent monitoring innovations, I compute the demand for deforestation holding the total number of fines issued during the 2004–2005 period at the same level as the number of fines issued during the previous period available in the data, 2002–2003. This reduces the cumulative number of fines from an average of 23 fines per municipality to 15 (and reduces the standard deviation as well, from 55 to 39). In doing so, I 'shut down' the recent increase in monitoring effort in the estimated demand for deforestation.

Accordingly, Figure 6 presents the two demand functions. Eliminating the recent increase in monitoring amounts to an upward shift in deforestation levels, as expected. But the differences are economically small. In the absence of land-use taxes, the share of deforested area increases by 3 percentage points, from 56.2 percent to 59.3 percent on the private properties. The two demand curves are very close to each other throughout. Even though the use of satellite-based monitoring is an important innovation that likely has large effects on deforestation in the long run

[42] The numbers in the legend correspond to the quantiles of the agricultural area in the data, and are the same for all maps. The quantiles are $\{0.01, 0.05, 0.1, 0.2, 0.3, ..., 0.8, 0.9\}$.

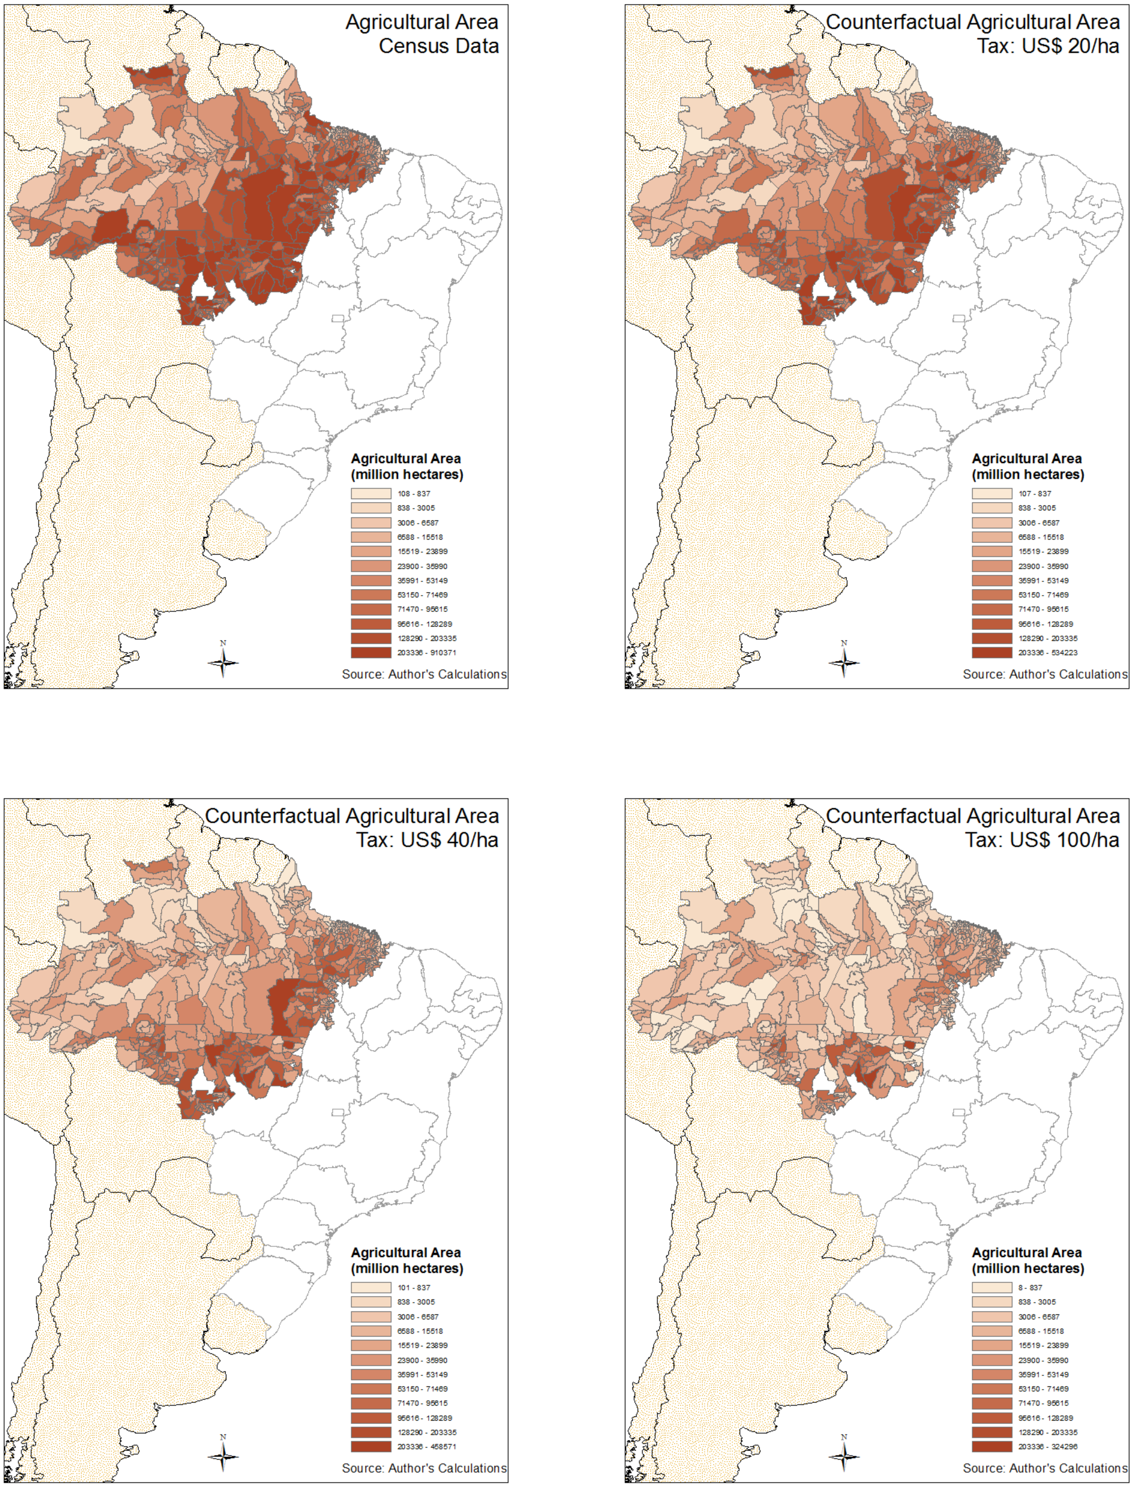

Figure 5: Geographic Distribution of the Demand for Deforestation

(see, e.g., Assunção et al., 2013), it appears to be too recent to affect the demand function in 2006 substantially.

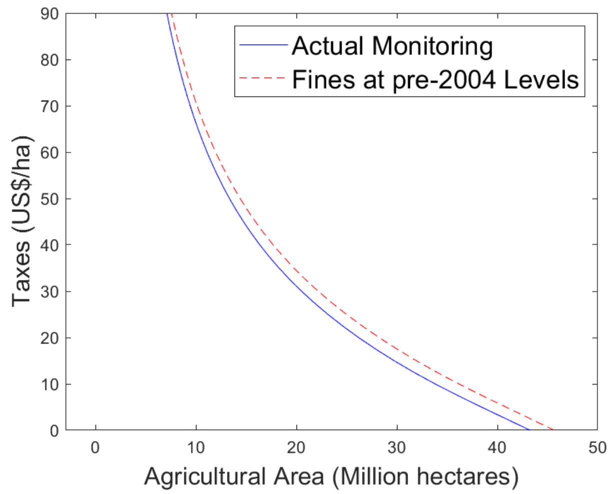

Figure 6: Demand for Deforestation – Eliminating Increases in Monitoring after 2004

8 Policy Analysis

In this section, I apply the framework to examine important policy questions. First, I consider the implications for carbon emissions and for the optimal tax. I then examine the resulting economic costs for the three policy interventions (taxes, payments and the '80 percent rule'). I close the section by discussing some limitations of the analysis.

8.1 Emissions of CO_2 and Optimal Tax

Avoided Emissions. Estimates of the carbon stock in forested and deforested areas are required in order to calculate potential emissions of carbon from land use change. Baccini et al. (2012) recently measured the geographic distribution of the above-ground carbon stock in Brazil. I combine their map of the carbon stock with the maps of deforestation from satellite images (INPE, 2017) and compute, for each municipality, the difference in the carbon stock in forested and deforested areas. These may vary across municipalities because forests are heterogeneous and the alternative land uses in agriculture may conserve more or less carbon on the ground. The average difference is almost 80 tons of carbon per hectare (tC/ha) – see Table 1.

Combining the differences in carbon stocks with the demand for deforestation results in a 'supply of avoided emissions.' Figure 7 presents this supply function. One may interpret the curve in the

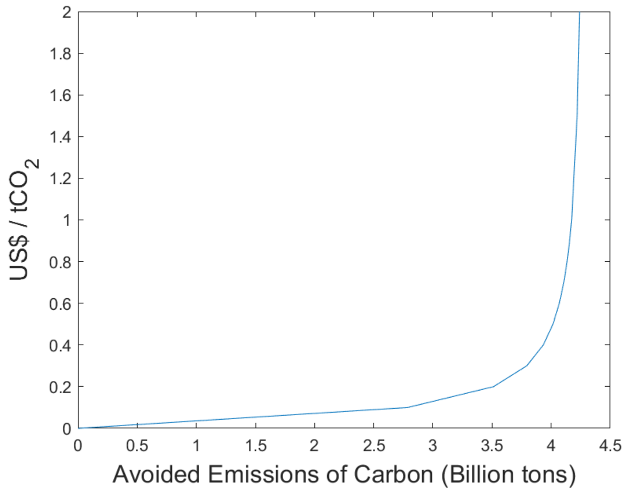

Figure 7: Supply of Avoided Emissions of Carbon

following way: if a carbon tax of (or a payment program paying) US$ 1 per ton of CO_2 per year were implemented, farmers would be willing to avoid the emissions of approximately 4.17 billion tons of carbon. The avoided emissions correspond to approximately 4.5 years of worldwide emissions from land use change during 2002 to 2011 (IPCC, 2013). The flat part of the curve results from the large stocks of carbon in forested areas combined with the fact that farmers would be responsive to taxes. The steeper portion, to the right, is the result of capacity constraints: with higher taxes, farmers would be willing to keep almost all private land forested.[43]

The shape of the supply curve I estimate is similar to that in other studies; see, e.g., Nepstad et al. (2007). Kindermann et al. (2008) and Lubowski and Rose (2013) also investigate the costs of avoided emissions from deforestation based on a number of different large-scale models (most of which are 'engineering' models). Although comparing my results to this literature is not trivial (because the approaches differ in several dimensions – the modelling assumptions, the underlying data, the spatial resolution, the region and time period covered, and the baseline deforestation rate in the absence of carbon taxes/payments, among others), the results underline the significant potential for emissions reduction through reducing deforestation. This has the potential to play an important role in global climate change policies.

A carbon tax of US$ $1/tCO_2$ per year is significantly smaller than the average price of carbon

[43]The calculation assumes that (a) the difference in the carbon stock comparing forested and deforested areas would be released into the atmosphere once the forest was removed (i.e., it ignores gradual decay); and (b) the carbon taxes would not affect the amount of carbon stock on agricultural land, although in principle, farmers could respond to carbon taxes by using new techniques that conserved more carbon on the ground. (To better understand the calculation, note that for a parcel in which the carbon difference between forested and deforested land is 80 tC, a land-use tax of \$10/ha is equivalent to a carbon tax of \$0.125/tC (and to $\$0.034/tCO_2$, because $1tC = 44/12\ tCO_2$).)

under the European Union Emissions Trading System. The carbon price was US$ 8.5/tCO_2 in December of 2006, and US$ 5.8/tCO_2 (2006$) ten years later.[44] The difference between the potential carbon tax and the market price of carbon suggests substantial opportunities for trade, but such opportunities have not yet been realized. One possible reason lies in the potentially large transactions costs involved. Measuring and monitoring the amount of carbon stocks may be expensive, and perhaps more importantly, measuring avoided emissions depends on the counterfactual emissions that would occur in the absence of payments, which is not trivial to compute.

The Optimal Tax. In general, a Pigouvian tax should be set equal to the marginal damage of the externalities caused by deforestation, such as emissions of carbon and biodiversity loss, at the social optimum. Because it is difficult to measure the production of all the relevant externalities and to value the corresponding damages, I focus on a lower bound for the optimal tax, which can be obtained from the estimated damages associated with an incremental change in CO_2 emissions. According to Greenstone et al. (2013), the central estimate of the social cost of carbon (SCC) for 2010 was US$ 21/tCO_2. Nordhaus (2014) estimates the central value of the SCC for 2015 to be US$ 18.5/tCO_2. As Figure 7 shows, imposing taxes of this magnitude would virtually eliminate all agricultural land in the Amazon.

8.2 The Costs of the Policy Interventions

I now examine the economic costs of the three policy interventions (taxes, payments and the '80 percent rule'). Starting with the '80 percent rule,' I calculate a lower bound for its hidden cost by means of perfectly-enforced land-use taxes that restrict farmers to use only 20 percent of their properties for agriculture. I compute these taxes separately for *each* municipality and for *each* farm size because they correspond to the most disaggregated information available in the data set. (Ideally, individual taxes would be calculated for each farm.) The corresponding farmers' lost surplus from these taxes is the sum of the trapezoid areas below each demand curve. That sum amounts to US$ 4 billion. Farmers are therefore estimated to be willing to pay around US$ 4 billion per year to avoid the enforcement of the 80 percent rule. Not surprisingly, farmers have tried, systematically, to alter the legislation since it was implemented.[45]

[44] The calculation converts dollars to 2006 using US Consumer Price Index data.

[45] The amount of money farmers would be willing to pay is likely even larger, for three reasons: (a) the command-and-control policy imposes the same limit on farmers' land use regardless of differences in opportunity costs of agricultural land, while the local land-use taxes I used to approximate the '80 percent rule' are (locally) a cost-effective price instrument; (b) the legislation does not allow for managed forests in preservation areas, except under very stringent conditions, but the forested area in the data potentially includes managed forests; and (c) the costs to

A uniform tax charging US$ 42.5/ha per year of agricultural land would also lead to 20 percent of land used for agriculture, but would be almost ten times cheaper than the '80 percent rule': farmers' lost surplus would be approximately US$ 479 million per year, provided the tax revenues were redistributed to them – tax revenues would be approximately US$ 658 million per year, which is 0.37 percent of the Brazilian federal budget for 2006. The difference corresponds to a cost saving from the land use tax of approximately 90 percent of the cost of a perfectly enforced '80 percent rule.' This is substantially higher than the cost saving estimates from allowance trading in pollution markets when compared to standard caps on emissions, which range from 20 to 47 percent (see Schmalensee and Stavins, 2017). The '80 percent rule' is not a cost-effective policy: the more productive farms would have to use less land for agriculture and so would have foregone greater profits, compared to taxes. Although, in theory, command-and-control approaches could also be cost-effective, they would require policy makers to obtain detailed information about each farmer's opportunity costs of land use, and set different limits on land use for each farm. Such information is not readily available to policy makers (Stavins, 2000).[46]

In terms of other possible policy responses, payment programs making payments of US$ 42.5/ha would also result in 80 percent of forest cover remaining, and lead to the same economic costs: US$ 479 million per year. Total transfers to farmers (potentially from the federal government) would have been US$ 2.61 billion per year (1.45 percent of the Brazilian federal budget for 2006). A perfectly targeted policy paying farmers who were going to deforest but *not* paying those who were *not* going to deforest would reduce total transfers by more than a half (US$ 1.18 billion per year). Perfect targeting is unlikely, however, because of asymmetric information problems (see Pattanyak et al., 2010; Busch et al., 2012; Mason and Platinga, 2013).

As noted previously, the 2SLS estimator cannot capture heterogeneous responses to changes in private values, and this may affect the estimated costs of policy interventions. Indeed, the 2SLS parameter estimates presented in Table 4 are relatively small in magnitude and do not reflect the high sensitivity of deforestation to changes in transportation costs at various quantiles. As a consequence, the land-use tax that results in 80 percent forest cover would be substantially higher based on the 2SLS estimates than the land-use tax based on the IVQR estimates: US$ 65/ha (versus US$ 42.5/ha). The corresponding economic cost would be approximately US$ 786 million, which is 64 percent higher than the estimated cost based on the quantile approach. The greater

replant vegetation add to the farmers' total costs since, by law, they must restore the forest at their own expense.

[46] Quantitative limits could be combined with trading, which would improve efficiency; see May et al. (2015) for an excellent discussion of environmental quotas to conserve or restore forests.

lost surplus from the 2SLS estimator thus points to the importance of allowing for heterogeneity in the demand for deforestation.

8.3 Limitations

The above analysis has several limitations. First, the estimated economic costs of the policies do not take into account monitoring and transactions costs. Given that Brazil does not yet have an established system to monitor the public expenditures directed at environmental issues (OECD, 2016), a useful proxy for existing monitoring costs is the sum of INPE's and IBAMA's combined annual budgets.[47] This is likely to provide an upper bound for the monitoring costs, since not all resources are directed to monitoring. INPE's budget for 2010 was US$ 125 million, and IBAMA's budget for 2011 was US$ 560 million. These numbers are comparable to the US$ 658 million of government tax revenues I obtained under a land-use tax of US$ 42.5/ha. This suggests that tax revenues could be sufficiently large to cover monitoring and enforcement costs.

A second limitation is that general equilibrium effects are not considered here. Given that the Amazon is responsible for 20 percent of Brazil's agricultural area, policies that affect land use there substantially may lead to non-trivial impacts on the country's trade balance, on the international prices of beef and soybeans, and possibly on the national economy. In addition, they may affect the total private land area (and the land distribution). For example, the '80 percent rule' – if perfectly enforced – may provide incentives to establish large farms. Also, large areas of unprotected public forested land might be occupied in response to payment programs. Such occupations could reduce the potential programs' effectiveness, increase their total costs, and increase disputes over land (along with the potential violence associated with these disputes). Although important, the full extent of such policy implications cannot be addressed with the present data set. The evidence I present should thus be viewed as one of the inputs – if a fundamental one – necessary for a complete evaluation of conservation policies.

9 Conclusion

In this paper, I have set out a unified econometric framework for estimating the impacts and cost effectiveness of alternative conservation policies. Based on data from the Brazilian Amazon, the main policy implications of this exercise are: (a) land-use taxes and payment programs can be

[47] INPE is responsible for the detection of 'hot spots' using satellite data and for providing the information to IBAMA, which then sends its inspectors to those areas, issues the appropriate fines, and follows up with the relevant administrative and judicial processes.

effective in avoiding deforestation and the emissions of carbon; (b) policy interventions that *only* target small landholders are not sufficient to promote substantial conservation; and (c) command-and-control policies that limit deforestation allowed on private properties are considerably more costly than incentive-based policies.

There are several directions for future research. First, micro-data on farmers' decisions are likely to provide a fuller picture of their opportunity costs and may help uncover the entire distribution of farmer's private valuations within each municipality, which in turn may help assess the use of auctions to allocate PES contracts. Second, a panel data set based on satellite images would allow a dynamic model of land use decisions to be estimated. Such a model could be used to study impacts of commodity prices on the rate of deforestation and how much these prices could alter the effectiveness of different policy interventions. Third, the framework presented here can be used to study impacts of improvements of roads on deforestation. This is an important topic, looking ahead, because the Brazilian government is paving existing unpaved roads in the Amazon to reduce the costs of exporting commodities.

References

Alix-Garcia, J. and H. Wolff (2014). Payment for ecosystem services from forests. *Annual Review of Resource Economics 6*, 361–380.

Allen, T. and C. Arkolakis (2014). Trade and the topography of the spatial economy. *Quarterly Journal of Economics 129(3)*, 1085–1140.

Almeida, C. A., A. C. Coutinho, J. C. D. M. Esquerdo, M. Adami, A. Venturieri, C. G. Diniz, N. Dessay, L. Durieux, and A. R. Gomes (2016). High spatial resolution land use and land cover mapping of the Brazilian Legal Amazon in 2008 using Landsat-5/TM and MODIS data. *Acta Amazonica 46(3)*, 291–302.

Andersen, L. E., C. W. Granger, E. Reis, D. Weinhold, and S. Wunder (2002). *The Dynamics of Deforestation and Economic Growth in the Brazilian Amazon.* Cambridge University Press, U.K.

Assunção, J., C. Gandour, and R. Rocha. (2013). DETERing deforestation in the Brazilian Amazon: Environmental monitoring and law enforcement. *Climate Policy Initiative Working Paper Series.*.

Assunção, J., M. Lipscomb, A. M. Mobarak, and D. Szerman. (2016). Agricultural productivity and deforestation in Brazil. *Mimeo, Yale University*.

Assunção, J., R. McMillan, J. Murphy, and E. Souza-Rodrigues (2018). Optimal environmental targeting in the Amazon rainforest. Working Paper, University of Toronto.

Baccini, A., S. J. Goetz, W. S. Walker, N. T. Laporte, M. Sun, D. Sulla-Menashe, J. Hackler, P. S. A. Beck, R. Dubayah, M. A. Friedl, S. Samanta, and R. A. Houghton (2012). Estimated carbon dioxide emissions from tropical deforestation improved by carbon-density maps. *Nature Climate Change 2*, 182–185.

Banerjee, A., E. Duflo, and N. Qian (2012). On the road: Access to transportation infrastructure and economic growth in China. *NBER Working Paper No. 17897*.

Berry, S. (2011). Biofuels policy and the empirical inputs to GTAP models. Technical report.

Berry, S. T. and P. A. Haile (2014). Identification in differentiated products markets using market level data. *Econometrica 82(5)*, 1749–1797.

Brady, M. and E. Irwin (2011). Environmental effects of land use change revealed through spatially explicit empirical models: Recent developments and challenges ahead. *Environmental and Resource Economics 48(3)*, 487–509.

Brito, B. and P. Barreto (2006). Sugestões para aumentar a eficácia da lei de crimes ambientais no IBAMA e nos tribunais de justiça no setor florestal do Pará. *11 º Congresso Brasileiro de Direito Ambiental: Biodiversidade e Direito em São Paulo. São Paulo.*.

Burke, M., M. Craxton, C. D. Kolstad, C. Onda, H. Allcott, E. Baker, L. Barrage, R. Carson, K. Gillingham, J. Graff-Zivin, M. Greenstone, G. Heal, S. Hsiang, B. Jones, D. L. Kelly, K. Kopp, M. Kotchen, R. Mendelsohn, K. Meng, G. Metcalf, J. Moreno-Cruz, R. Pindyck, S. Rose, I. Rudik, J. Stock, and R. S. J. Tol (2016). Opportunities for advances in climate change economics. *Science 352*, 292–293.

Busch, J., R. N. Lubowski, F. Godoy, M. Steininger, A. A. Yusuf, K. Austin, J. Hewson, D. Juhn, M. Farid, and F. Boltz (2012). Structuring economic incentives to reduce emissions from deforestation within Indonesia. *Proceedings of the National Academy of Sciences 9(4)*, 1062–67.

Castro, N. (2003). Formação de preços no transporte de carga. *Pesquisa e Planejamento Econômico 33(1)*, 167–189.

Chen, X. and D. Pouzo (2012). Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals. *Econometrica 80(1)*, 277–321.

Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010). Quantile and probability curves without crossing. *Econometrica 78(3)*, 1093–1125.

Chernozhukov, V. and C. Hansen (2008). Instrumental variable quantile regression: A robust inference approach. *Journal of Econometrics 142(1)*, 379–398.

Chernozhukov, V. and C. Hansen (2013). Quantile models with endogeneity. *Annual Review of Economics 5*, 57–81.

Chomitz, K. M. and D. A. Gray (1996). Roads, land use, and deforestation: A spatial model applied to Belize. *The World Bank Economic Review 10*(3), 487–512.

Chomitz, K. M. and T. S. Thomas (2003). Determinants of land use in Amazonia: A fine-scale spatial analysis. *American Journal of Agricultural Economics 85*(4), 1016–1028.

Donaldson, D. and R. Hornbeck (2016). Railroads and American economic growth: A market access approach. *Quarterly Journal of Economics 131(2)*, 799–858.

FAO (2016). The global forest resources assessment 2015: How are the world's forests changing? *Food and Agriculture Organization of the United Nations, Second Edition, Rome.*.

Foster, A. D. and M. R. Rosenzweig (2011). Are Indian farms too small? Mechanization, agency costs, and farm efficiency. *Mimeo, Yale University.*.

Godar, J., T. A. Gardner, E. J. Tizado, and P. Pacheco (2014). Actor-specific contributions to the deforestation slowdown in the Brazilian Amazon. *Proceedings of the National Academy of Sciences 111(43)*, 15591–15596.

Greenstone, M. and B. K. Jack (2015). Envirodevonomics: A research agenda for an emerging field. *Journal of Economic Literature 53*(1), 5–42.

Greenstone, M., E. Kopits, and A. Wolverton (2013). Developing a social cost of carbon for us regulatory analysis: A methodology and interpretation. *Review of Environmental Economics and Policy 7(1)*, 23–46.

Hausman, J. A. and J. G. Sidak (2004). Why do the poor and the less-educated pay higher prices for long-distance calls? *Contributions to Economic Analysis & Policy 3(1)*, 1–26.

Herrera, L. D. (2015). *Protected Areas' Deforestation Spillovers and Two Critical Underlying Mechanisms: An Empirical Exploration for the Brazilian Amazon.* Ph. D. thesis.

Hirakuri, S. R. (2003). Can law save the forest? Lessons from Finland and Brazil. *Centre for International Forestry Research, Jakarta, Indonesia*.

INPE (2017). Monitoramento da Floresta Amazônica Brasileira por satélite – Projeto Prodes. *Available at http://www.obt.inpe.br/prodes/*.

IPCC (2007). *Climate Change 2007: The Physical Science Basis.* Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

IPCC (2013). *Climate Change 2013: The Physical Science Basis.* Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Jack, B. K. (2013). Private information and the allocation of land use subsidies in Malawi. *American Economic Journal: Applied Economics 5(3)*, 113–135.

Jayachandran, S., J. de Laat, E. F. Lambin, C. Y. Stanton, R. Audy, and N. E. Thomas (2017). Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation. *Science 357(6348)*, 267–273.

Kindermann, G., M. Obersteiner, B. Sohngen, J. Sathaye, K. Andrasko, E. Rametsteiner, B. Schlamadinger, S. Wunder, and R. Beach (2008). Global cost estimates of reducing carbon emissions through avoided deforestation. *Proceedings of the National Academy of Sciences 105(30)*, 10302–307.

Lamarche, C. (2011). Measuring the incentives to learn in Colombia using new quantile regression approaches. *Journal of Development Economics 96*, 278–288.

Lubowski, R. N., A. J. Plantinga, and R. N. Stavins (2006, March). Land-use change and carbon sinks: Econometric estimation of the carbon sequestration supply function. *Journal of Environmental Economics and Management 51*(2), 135–152.

Lubowski, R. N. and S. Rose (2013). The potential for REDD+: Key economic modeling insights and issues. *Review of Environmental Economics and Policy 7(1)*, 67–90.

Mason, C. F. and A. J. Platinga (2013). The additionality problem with offsets: Optimal contracts for carbon sequestration in forests. *Journal of Environmental Economics and Management 66*, 1–14.

May, P. H., P. Bernasconi, S. Wunder, and R. Lubowski (2015). Environmental reserve quotas in Brazil's new forest legislation: an ex-ante appraisal. *Center for International Forestry Research (CIFOR), Occasional Paper 131.*

Nepstad, D., B. Soares-Filho, F. Merry, P. Moutinho, H. O. Rodrigues, M. Bowman, S. Schwartzman, O. Almeida, and S. Rivero (2007). The costs and benefits of reducing carbon emissions from deforestation and forest degradation in the Brazilian Amazon. *The Woods Hole Research Center. United Nations Framework Convention on Climate Change (UNFCCC) Conference of the Parties (COP), Thirteenth session. 3–14 December 2007.*, 3–14.

Nordhaus, W. (2014). Estimates of the social cost of carbon: Concepts and results from the DICE–2013R model and alternative approaches. *Journal of the Association of Environmental and Resource Economists 1(1/2)*, 273–312.

OECD (2016). OCDE – avaliações de desempenho ambiental Brasil 2015. *ECLAC*.

Pattanyak, S., S. Wunder, and P. Ferraro (2010). Show me the money: Do payments supply environmental services in developing countries? *Review of Environmental Economics and Policy 4*, 254–274.

Pfaff, A. (1999). What drives deforestation in the Brazilian Amazon? Evidence from satellite and socioeconomic data. *Journal of Environmental Economics and Management 37*, 26.

Pfaff, A., A. Barbieri, T. Ludewigs, F. Merry, S. Perz, and E. J. Reis (2009). Road impacts in Brazilian Amazonia. *Amazonia and Global Change, eds. Keller, M., M. Bustamente, J. Gash and P. Dias, American Geophysical Union*.

Pfaff, A. and J. A. Robalino (2013). REDD roads? Spatial dynamics and spatial variation in causal impacts. *Mimeo, Duke University*.

Pfaff, A., J. A. Robalino, D. Herrera, and C. Sandoval (2015). Protected areas' impacts on Brazilian Amazon deforestation: Examining conservation – development interactions to inform planning. *PLOS One 10(7)*, e0129460.

Pinheiro, T. F., M. I. S. Escada, D. M. Valeriano, P. Hostert, F. Gollnow, and H. Müller (2016). Forest degradation associated with logging frontier expansion in the Amazon: the BR-163 region in Southwestern Pará, Brazil. *Earth Interactions 20(17)*, 1–26.

Pizer, W. A. and R. Kopp (2005). Calculating the costs of environmental regulation. *Handbook of Environmental Economics, Elsevier 3*.

PNLT (2006). National highway plan. *Available at http://www.transportes.gov.br/index/conteudo/id/36604 and accessed on 11/24/2010.*.

Robalino, J. A. and A. Pfaff (2012). Contagious development: Neighbor interactions in deforestation. *Journal of Development Economics 97(2)*, 427–436.

Roberts, M. J. and W. Schlenker (2013). Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *American Economic Review 103*(6), 2265–95.

Schmalensee, R. and R. N. Stavins (2017). Lessons learned from three decades of experience with cap and trade. *Review of Environmental Economics and Policy 11(1)*, 59–79.

Scott, P. T. (2013). Dynamic discrete choice estimation of agricultural land use. *Working Paper*.

Simonet, G., J. Subervie, D. Ezzine-de Blas, M. Cromberg, and A. E. Duchelle (2018). Effectiveness of a REDD+ project in reducing deforestation in the Brazilian Amazon. *American Journal of Agricultural Economics*, aay028.

Stavins, R. N. (1999). The costs of carbon sequestration: A revealed-preference approach. *American Economic Review 89(4)*, 994–1009.

Stavins, R. N. (2000). Market-based environmental policies. *Public Policies for Environmental Protection. eds. P. R. Portney and R. N. Stavins. Resources for the Future, Washington, D.C.*.

Stavins, R. N. and A. B. Jaffee (1990). Unintended impacts of public investments on private decisions: The depletion of forested wetlands. *American Economic Review 80(3)*, 337–352.

Stern, N. (2007). The economics of climate change: The Stern Review. *New York: Cambridge University Press*.

TCU, T. d. C. d. U. a. (2009). Arrecadação de multas administrativas – exercício de 2009. *Available at http://portal2.tcu.gov.br/portal/page/portal/TCU/comunidades/contas/*.

Timmins, C. and W. Schlenker (2009). Reduced-form versus structural modeling in environmental and resource economics. *Annual Review of Resource Economics 1*, 351–380.