

Supplement to ‘Deforestation in the Amazon: A Unified Framework for Estimation and Policy Analysis’

EDUARDO SOUZA-RODRIGUES*
Department of Economics, University of Toronto

FOR ONLINE PUBLICATION

1 Introduction

In this supplemental document, I provide a robustness analysis for the estimated regressions presented in the main paper (Section 2), and a detailed description of how the data set is constructed (Section 3).

Section 2 is divided into five subsections. In Subsection 2.1, I discuss the estimated coefficients on the covariates in the land use quantile regressions. In Subsection 2.2, I present the estimated results when I exclude the proxies for monitoring efforts: the number of fines issued and the distance to the nearest IBAMA office (the Brazilian Environmental Protection Agency). Consistent with the discussion in the main paper (that historically there had not been much enforcement in the Amazon rainforest), the inclusion and the exclusion of these variables do not affect substantially the coefficients on costs to ports, as well as on other regressors. Policy implications also are robust to the different specifications. In Subsection 2.3, I investigate how spatial dependence can affect the results. Coefficients on spatially lagged covariates are statistically significant, but results are not sensitive to the choice of weighting matrix nor to the cutoff distance. Subsection 2.4 drops the logit assumption in the paper and estimates a semiparametric model. The results establish the suitability of the logit model once there are no significant differences between the logit and the semiparametric models. Finally, in Subsection 2.5, I discuss the use of different productivity indices to calculate the demand for deforestation. I describe briefly the underlying economic model associated with the different indices, I present the corresponding estimated demands for deforestation, and I provide evidence that local yields do not respond to changes in transportation costs.

*University of Toronto, Department of Economics. E-mail: e.souzarodrigues@utoronto.ca. Financial support from Charles V. Hickox Fellowship, Arthur K. Watson Fund, Yale University and Kernan Brothers Environmental Fellowship is gratefully acknowledged. All errors are mine.

I describe the data set in Section 3. It includes a detailed explanation about how the share of deforestation (Subsection 3.1), the transportation costs (Subsection 3.2), the productivity indices (Subsection 3.3), and the covariates (Subsection 3.4) were constructed. In Subsection 3.1, in addition to explaining how deforestation is calculated using the Brazilian Agricultural Census, I also compare it with satellite-based deforestation measures. I discuss their differences and similarities, and I show that when they are most comparable (i.e., when the share of private land in the municipality is large and the share of clouds/unobserved areas in satellite images is small), their correlation in the data is significantly high. Furthermore, I find no evidence that the census deforestation varies systematically with the exogenous variables after conditioning on the satellite deforestation measure.

2 Robustness

2.1 Land Use Quantile Regression – Main Specification

In this section, I discuss briefly the estimated coefficients on the covariates for the main specification. Table 1 presents the results for the IVQR estimator at the median. For brevity, it omits the coefficients on the spatially lagged regressors, but I present and discuss them in Subsection 2.3.

First, the coefficients on altitude are small in magnitude and not significant. Second, higher levels of temperature and precipitation in the Amazon are worse for agriculture, as should be expected. The already high levels of temperature and rainfall in the region make agriculture unattractive because, among other reasons, cattle are more susceptible to parasites and insect pests, crops are more subject to rotting, yields are depressed by light-limiting cloud cover, mechanization is difficult, and forest burning is incomplete (Chomitz and Thomas (2003)). Third, locations with steeper slopes are worse for agriculture as well (although not significantly so). Fourth, soils with worse quality induce less deforestation when compared to good soils, as expected, although the effects are not monotonic in the rank of soil quality (the proportion of good soils is omitted in the regressions).

As discussed in the main text, protected areas (PAs) may have spillover effects that affect the value of nearby forestlands. Except for large farms, the coefficients on distance to protected areas are small in magnitude and are not statistically significant, and most of them have a positive sign: the smaller the distance, the lower the deforestation. This is consistent with the regularity that protected areas tend to be located on land facing lower deforestation pressure (as documented by Pfaff, Robalino, Herrera, and Sandoval (2015)), and the possibility that the aggregated nature

of the data may mask multiple heterogeneous spillover effects. (On this note, Robalino, Pfaff, and Villalobos-Fiatt (2017) present evidence consistent with that in Costa Rica, using detailed disaggregated data.)

Mining, the presence of power plants either in the municipality or within 75 km from the municipal seat, and the local population are meant to capture factors that influence local demand and supply for non-tradables. The direction of the impacts are not clear *ex-ante*, as they depend on how these factors affect the demand for local products, the demand and supply for labor, and which sector (agriculture or forestry production) is more labor-intensive. Most of the estimated coefficients are negatively correlated with the share of agricultural land, but they are not significant in most cases. I discuss below robustness results related to local population and why it may be difficult to estimate effects of factors that affect local labor market conditions in the Amazon rainforest.

I also discuss below the robustness of the results with respect to the proxy for property rights, and with respect to the proxies for monitoring efforts (the number of fines and the distance to the nearest IBAMA agency). Here, I emphasize only that their coefficients are not statistically significant.

Robustness: Local Population. I now investigate whether the exclusion of local population in the land use regressions affects the main results. Table 2 compares the estimated coefficients on the regressors for each farm size. I present results only for the median, but the pattern is similar for other quantiles. As mentioned previously, the local population does not have significant coefficients. Furthermore, omitting it in these regressions does not affect the coefficients on the other regressors significantly. The population in the Amazon is sparsely distributed and, so, labor is likely a scarce factor in the region. Medium-sized and large farms, therefore, are likely to have production functions that are intensive in land, but not in labor. Indeed, the average number of agricultural workers per hectare of agricultural land is: 4/ha for small farms; 0.24/ha for small–medium farms; 0.054/ha for medium–large farms; and 0.012/ha for large farms. Therefore, the share of wages on costs is probably small, which may help explain why factors that shift the demand and supply of labor are not significant in the regressions presented in Table 1. So, even though one may argue that the population is endogenous, instrumenting it likely does not change the primary results.¹

¹One may expect local labor markets to be imperfect, maybe as a result of the sparsely distributed population. A symptom of this imperfection is that farmers rely more on family labor than on hired labor. Indeed, the proportion of family workers in the total number of workers averaged 83% in the Amazon, ranging from a low of 10% in more developed regions in the South Amazon to a high of 99%–100% in the isolated areas in the Western Amazon.

Robustness: Proxy for Property Rights. As discussed in the main text, another problem that may invalidate the instruments is the potential lack of property rights. Table 3 presents the results for the median regressions with and without the proportion of private land with land title, which is a proxy for tenure security. The inclusion and the exclusion of the proxy in the regressions do not affect the estimates significantly. Although both estimates could be equally biased, the results are reassuring. Most farmers do have their land titles (85 percent) and it is possible that they may not need to deforest as much to guarantee their property rights. Or, at least, their land tenure status may not affect how their deforestation decisions relate to transportation costs.

2.2 Penalties and Monitoring Efforts

I consider two proxies for monitoring efforts: the total cumulative number of fines (over the period 2002–2005), and the distance to the nearest IBAMA agency. The number of fines may capture the recent increase on penalties, while the distance to IBAMA may capture more permanent monitoring efforts. Table 4 presents the results for the IVQR median regressions. It omits the coefficients on all the regressors except for the costs to ports and the proxies for penalties. The first column of the table presents the specification that excludes both proxies; the second column includes only the number of fines; the third column includes only the distance to IBAMA; and the last column includes both (the main specification).

The estimated coefficients on both proxies are not statistically significant, and the magnitudes are small. Specifically, when covariates are fixed at the sample mean and when I increase the number of fines by 10 (corresponding to approximately one-fifth of its sample standard deviation) the estimated marginal effects range from negligible to a 0.5 percentage point decrease in the share of deforestation, depending on the farm size and the quantile. Recall that the average number of fines in the data is 23; increasing that amount by 10 is therefore a substantial change in the number of fines.

Although the signs of the coefficients are in the expected direction, it is clear that the number of fines may suffer from simultaneity problems, which may cause upward biases on their estimated coefficients (i.e., the estimated magnitudes are closer to zero). As discussed in the main text, it is a nontrivial task to find good instruments for environmental transgressions in the Amazon; and the biases on the coefficients on costs to port caused by this simultaneity problem are less clear *ex-ante*. Against that, ignoring the number of fines in the regressions may lead to omitted variables bias. 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 the omission may

result in small biases.² The fact that the inclusion and the exclusion of the number of fines do not affect substantially the coefficients on costs to ports, as well as on other regressors, is reassuring. This is consistent with the interpretation that the increased monitoring efforts were too recent to substantially affect the total accumulated deforested land.

The coefficients on distance to IBAMA are on the expected direction for medium-sized and large farms: the greater the distance, the lower the monitoring, and so the larger the deforested area. Similarly to the number of fines, the coefficients are not statistically significant, their magnitudes are small, and the inclusion and exclusion of this variable do not change results significantly. Because distance to IBAMA captures more permanent monitoring effects, this is consistent with the interpretation that the legislation has not been fully enforced historically – at least up to 2006.

Policy Implications. Table 5 presents the policy implications for the different specifications. The first column presents the value of the uniform tax that would lead to 20 percent of agricultural land in private properties in the Amazon; the second column presents the associated farmers’ lost surplus from the uniform tax; the third column presents farmers’ lost surplus under the (perfectly enforced) ‘80 percent rule;’ and the fourth column, the amount of avoided emissions of carbon under a carbon tax of one dollar per ton of CO₂. The top panel presents the case in which the counterfactual policies are superimposed on the status quo scenario (I denote it by “status quo” counterfactual); the bottom panel considers the case in which the total number of fines during 2004–2005 is fixed at the same level as in the previous period 2002–2003. That is, it “shuts down” the recent increase in the monitoring efforts (see Section 7.2 in the main text).

In the counterfactual in which taxes are superimposed on the status quo scenario, the values of the uniform tax and their associated lost surpluses vary little across specifications: uniform taxes range from US\$ 37.50/ha to US\$ 42.50/ha; and lost surpluses, from US\$ 425 million to US\$ 482 million. The lost surpluses from the ‘80 percent rule’ range from US\$ 3.24 billion to US\$ 5.25 billion, which is 7 to 11 times more expensive than the uniform taxes for farmers. The avoided emissions are all close to each other: between 4.17 billion and 4.21 billion tons of carbon. The policy implications are therefore robust to the inclusion and exclusion of the proxies for monitoring in the model specification.

For the counterfactual in which the number of fines post-2004 is fixed at pre-2004 levels, I consider only the main specification as well as the specification with no distance to IBAMA in

²The raw correlation between number of fines and straight-line distance to ports is 0.1; and the raw correlation with distance to the nearest capital is 0.04. When regressing the total number of fines on the instruments for costs to ports and on all other covariates, I obtained small and insignificant coefficients for the instrumental variables.

the regressions. Both specifications deliver similar results. The value of the uniform tax increases to US\$ 47.50/ha, compared to US\$ 42.50 in the “status quo” counterfactual. The higher tax is not surprising given that the burden of reducing deforestation is much more relegated to the counterfactual policy. Similarly, the lost surpluses from the tax, from the ‘80 percent rule,’ and the amount of avoided emissions also increase compared to the “status quo” counterfactual, but they are all of comparable magnitudes. Note that the estimated lost surpluses consider only the losses from the counterfactual policy and ignore the benefits farmers enjoy from reducing the existing monitoring efforts.

2.3 Spatial Dependence

There are different potential sources of spatial dependence in the data. I considered several possibilities, which I briefly discuss now.

First, I have included spatially lagged regressors of the local demand and supply shifters because inter-related local markets may create spatial dependence in farmers’ decisions across municipalities. Spatial lags for proximity to protected areas are also included among regressors to allow for further spatial spillover effects of PAs on deforestation.

I have chosen to not include a spatially lagged dependent variable among the regressors in my final specification. When I estimated such a model using spatially lagged regressors as instrumental variables, I obtained unstable and unreasonable estimates. Such a model would capture direct interactions between farmers across municipalities. My interpretation is that this type of interaction across municipalities seems less likely or less important than interactions through equilibrium prices or local neighbor interactions.

Although local interaction among neighbor farmers is likely important (evidence of that can be found in Robalino and Pfaff (2012) and the references cited therein), the aggregated nature of the present data set limits my ability to incorporate them in the analysis. Another possible specification involves interactions of farms of different sizes in the same municipality. Doing so would require an extended model with a simultaneous system of equations in which deforestation in a given type of farm would directly affect the land use on other types of farms. To identify and estimate such a system I would have to find instrumental variables satisfying exclusion restrictions as in standard simultaneous equations problems. Finding such instruments is challenging; they are not available in the present data set (none of the regressors are farm-size specific). I therefore interpret my estimated results as the (reduced-form) equilibrium outcome of these local interactions.

Finally, I have not included spatial correlations in the unobservables because, as far as I am

aware, there is no currently available technique that incorporates such correlations into quantile regressions. Although spatial dependence in unobservables could affect the standard errors, presumably it does not bias the point estimates.

Table 6 presents the results for the IVQR median regressions when spatial dependence is taken into account. As before, the table omits the coefficients on all the regressors except for the costs to ports and the spatially lagged regressors. The first column presents the specification that ignores spatial dependence; the second column includes the spatially lagged regressors and considers a cutoff distance of 50 km; the third column considers a cutoff distance of 75 km (the main specification); and the fourth column, a cutoff distance of 100 km. The spatial weight matrices are of the power functional type with distance decay parameter that equals one. Most of the estimated coefficients on the spatially lagged regressors are not statistically significant; an exception is the presence of mining for medium-sized and large farms. As previously discussed in Subsection 2.1, it is difficult to predict the direction of the impacts of the local demand and supply shifters. The results are not highly sensitive to the choice of the cutoff distance.³

2.4 Semiparametric Model

As a robustness exercise, I relax functional form restrictions by dropping the logit assumption to check whether it may drive the results. Specifically, I estimate the semiparametric model:

$$G_s(Y_m(s), u) = X_m\beta_{su} - TC_m, \quad (1)$$

where the function $G_s(\cdot, u)$ is unknown. Because scale and location normalizations are necessary, the coefficient on transportation cost is normalized to be minus one and the constant, to be zero. The normalization is without loss as long as TC_m affects deforestation negatively.

2.4.1 Implementation of the Estimator

To estimate the model (1), I use the Penalized Sieves Minimum Distance (PSMD) estimator proposed by Chen and Pouzo (2009, 2012). I assume (Z_m, X_m) is independent of $U_m(s)$ for any s , where $Z_m = (Z_m^p, Z_m^c)$ denotes the straight-line distances to the nearest port and to the nearest capital. For each farm size s and quantile $u \in (0, 1)$, the moment restriction implied by the model is:

$$E[\rho_{su}(Y_m(s), X_m, TC_m; G_{su}, \beta_{su}) | Z_m, X_m] = 0, \quad (2)$$

³Consistent with the discussion in Le Sage (2014), results are robust to the choice of the weighting matrix. In light of the results presented in Assunção, Lipscomb, Mobarak, and Szerman (2016), I kept the presence of power plants in neighbor municipalities in all specifications.

where the residual function is

$$\rho_{su}(Y_m(s), X_m, TC_m; G_{su}, \beta_{su}) \equiv 1 \{G_s(Y_m(s), u) \leq X_m \beta_{su} - TC_m\} - u, \quad (3)$$

and the conditional moment function is

$$m(Z_m, X_m; G_{su}, \beta_{su}) \equiv E[\rho_{su}(Y_m(s), X_m, TC_m; G_{su}, \beta_{su}) | Z_m, X_m]. \quad (4)$$

I approximate G_{su} using an artificial neural network (ANN) sieve approximation, because this non-linear sieve is often, in practice, better able than alternatives to allow for nonlinearities in the unknown function (Chen (2007)). More specifically, I use a sigmoid ANN defined by:

$$sANN(k_m) = \left\{ \sum_{j=1}^{k_m} \alpha_j S(\gamma_j Y_m + \gamma_{0,j}) : \alpha_j, \gamma_j, \gamma_{0,j} \in \mathbb{R} \right\} \quad (5)$$

where $S: \mathbb{R} \rightarrow \mathbb{R}$ is a sigmoid activation function. Note that the dimension of $sANN(k_m)$ is $3k_m$. I opted for a Gaussian activation function, i.e., I take $S(\cdot)$ to be a normal distribution function.

For each s and u , the function G_{su} in $sANN(k_m)$ and the finite dimensional parameter β_{su} are chosen to minimize the criterion function:

$$Q(G_{su}, \beta_{su}) = \left\{ \frac{1}{M} \sum_{m=1}^M \widehat{m}(Z_m, X_m; G_{su}, \beta_{su})' \widehat{m}(Z_m, X_m; G_{su}, \beta_{su}) + \lambda_m \widehat{M}(G_{su}, \beta_{su}) \right\}$$

where M is the number of municipalities; $\widehat{m}(\cdot)$ is an estimator for $m(\cdot)$; the penalization parameter, $\lambda_m \geq 0$, converges to zero as $M \rightarrow \infty$; and $\widehat{M}(G_{su}, \beta_{su})$ is the penalization function.

I take $\widehat{m}(\cdot)$ to be a series least square estimator of $m(\cdot)$. Let $\{p_1(Z_m, X_m), p_2(Z_m, X_m), \dots\}$ be a sequence of known basis functions that approximate any square integrable real-valued function. Denote $p^{J_M}(Z_m, X_m) = (p_1(Z_m, X_m), \dots, p_{J_M}(Z_m, X_m))'$ a $(1 \times J_M)$ -vector, and define the $(M \times J_M)$ matrix $P = (p^{J_M}(z_1, x_1), \dots, p^{J_M}(z_M, x_M))'$. The series least square estimator is given by:

$$\widehat{m}(z, x; G_{su}, \beta_{su}) = p^{J_M}(z, x)' (P'P)^{-} \sum_{m=1}^M p^{J_M}(Z_m, X_m) \rho_{su}(Y_m(s), X_m, TC_m; G_{su}, \beta_{su})$$

where $(P'P)^{-}$ is the pseudo-inverse matrix of $P'P$. The penalization function used is:

$$\widehat{M}(G_{su}, \beta_{su}) = \|\nabla_y^2 G_{su}\|_{L^2(\widehat{f}_{Y(s)})}$$

where ∇_y^2 denotes the second derivative with respect to y and $\|\cdot\|_{L^2(\widehat{f}_{Y(s)})}$ denotes the L^2 -norm with the empirical measure $\widehat{f}_{Y(s)}$ of $Y_m(s)$.

Next, I discuss the specific approximations I adopted. Because the asymptotic theory provides guidance on the rate at which k_m must increase with the data, but not the specific value for

k_m , I choose k_m so that the number of coefficients estimated does not exceed the number of moment restrictions.⁴ Specifically, the total number of parameters to be estimated is $\dim(\beta) + 3k_m$. The number of covariates in the regressions is 20 (including the spatially lagged regressors), so $\dim(\beta) = 20$. By choosing $k_m = 3$, the total number of parameters becomes 29. To approximate the conditional moment $m(Z_m, X_m)$, I use a basis function $p^{J_M}(Z_m, X_m)$ with dimension $J_M = 31$, so that the number of moment restrictions ($J_M = 31$) is greater than the number of model parameters ($\dim(\beta) + 3k_m = 29$). With respect to the penalization parameter, I use $\lambda_m = 10^{-6}$.

Although the nonparametric function G_{su} depends only on the scalar $Y_m(s)$, approximating the function $m(Z_m, X_m)$ is difficult because of the curse of dimensionality. For this reason, I had to be very parsimonious in choosing the number of terms for $p^{J_M}(Z_m, X_m)$. I opted to use P-Splines(2,4) for both Z_m^p and Z_m^c , i.e., quadratic splines with four knots (chosen at the respective 0.01, 0.25, 0.75 and 0.99 quantiles of Z_m^p and Z_m^c), and an interaction term $Z_m^p \times Z_m^c$. Each spline has dimension 5; the two splines plus the interaction therefore result in 11 terms. The remaining exogenous terms, X_m , enter linearly in $p^{J_M}(Z_m, X_m)$, yielding the dimension $J_M = 31$. When I experimented by approximating G_{su} with higher dimensions, say, $k_m = 4$, the dimensions of the splines for $Z_m = (Z_m^p, Z_m^c)$ were also increased to cubic splines, i.e., P-Splines(3,4). Including many more higher order terms for any of the instruments in $p^{J_M}(Z_m, X_m)$ made the estimation routine more difficult, because the matrix $(P'P)$ becomes singular very quickly. Note that to restrict the function $m(Z_m, X_m)$, it would be necessary to specify the conditional distribution of the endogenous variables $(Y_m(s), TC_m)$ given the exogenous (Z_m, X_m) . Such a specification is not imposed by the model.

In the estimation routine, I used the estimates from the IVQR model as an initial guess for β_{su} and the best $sANN(k_m)$ approximation for the rescaled logistic function as an initial guess for G_{su} (where the rescaling used IVQR estimates). Therefore, as the minimization routine proceeds, it indicates whether or not the estimates of the IVQR are good estimates for the SPQIV. As shown below, the results indicate that they are indeed good estimates. Even when I started the minimization routine using different initial guesses, I could not obtain better results for the criterion function.

⁴Chen and Pouzo (2012) provide very general results for the rates of convergence for NPQIV models. However, to date, there exists no explicit rate of convergence in terms of the number of observations in the data for non-linear sieves such as the $sANN$ used here.

2.4.2 Results for the Semiparametric Model

In this section I present results only for medium-sized and large farms. I do not present results for small landholders because, as shown in the main text, they do not seem to respond significantly to transportation costs.

I focus on the specification with $k_m = 3$ because I obtained very similar results by setting $k_m = 4$. Figure 1 compares the estimates of the approximating function for G_{su} using $sANN$ (3) and the logistic function for the median. The left panel presents the results for small-to-medium farms (5–50 ha); the middle panel reports estimates for medium-to-large farms (50–500 ha); and the right panel, for large farms (more than 500 ha). To make the comparisons fair, I rescaled the logistic function from the IVQR estimates to make the coefficient on TC_m equal -1 and the constant term in the single-index equal zero. The $sANN(k_m)$ approximations for G_{su} are surprisingly similar to the rescaled logistic functions (and similar results hold for almost all the quantiles).⁵

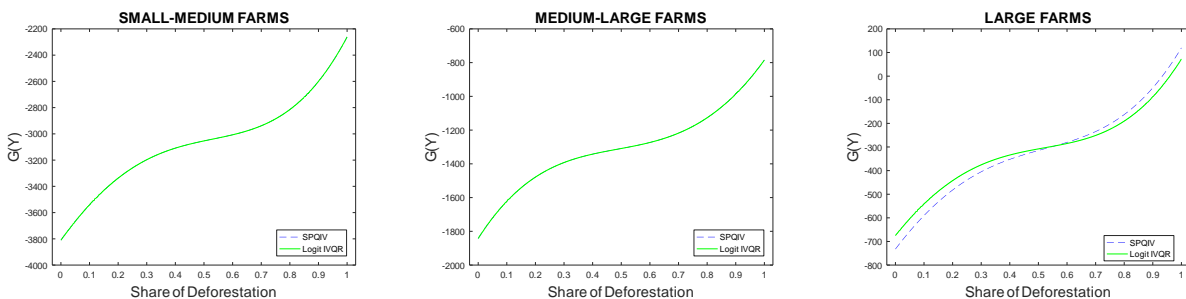


Figure 1: $G(Y,u)$ and $\text{Log}(Y/1-Y)$ at the Median

Not only is the link function, G_{su} , similar for the SPQIV and the IVQR estimates, but the estimated finite dimensional coefficients, β_{su} , are also similar. Table 7 compares the estimated coefficients using the PSMD and the normalized IVQR estimates.⁶ I illustrate only the results for the median regressions, but this pattern is similar for other quantiles. It is clear how (surprisingly) close the estimates are to each other. Although I do not run any formal tests between the IVQR and SPQIV estimates, the results illustrated in Figure 1 and Table 7 suggest that the logit assumption

⁵Note that, in Figure 1, neither the logistic nor the G_{su} functions seem to tend to minus and plus infinity at the extremes: that is because of the scale of the figure and because the share of deforestation in the data is truncated away from 0 and 1. Although not presented in the text, I have also approximated G_{su} using P-Splines and Hermite polynomials, instead of $sANN(k_m)$. The estimated functions based on P-Splines and Hermite polynomials are close to the rescaled logistic function, but only when Y is around 0.5. The global approximations in these cases are too flat and fail to capture any curvature of G_{su} when Y is close to either 0 or 1. As a result, P-Splines or Hermite approximations might predict the share of agricultural land outside the unit interval. Despite the fact that the instruments are strong in a linear sense and able to identify β , they may be weak in a “non-linear” sense, i.e., may fail to identify curvatures of G_{su} in a finite data set. Using the $sANN(k_m)$ approximation helps impose more restrictions on G_{su} in the data.

⁶The t-statistics for the normalized IVQR estimates were computed using the delta method.

is well suited for this data set.

2.5 Local Yields

In this section, I discuss the use of local yields to calculate the demand for deforestation. As mentioned in the main text, 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 in the Amazon to construct two productivity indices. The details of how I calculated the indices are explained in Section 3.3. Here, I briefly (a) discuss the underlying economic model associated with them, (b) present the corresponding estimated demands for deforestation, and (c) provide evidence that local yields do not respond to changes in transportation costs.

I consider two productivity indices. The first one, denoted by $q_m^{cp}(s)$, is based on the production of beef and the yields of the most representative crops (those discussed in Section 3 of the main paper); it corresponds to the main specification presented in the paper. The second index, denoted by $q_m^c(s)$, includes only the main crops and ignores pasture land.

Multinomial Choice Models. Different productivity indices are associated with different underlying multinomial choice models. The first index, $q_m^{cp}(s)$, is associated with a simple “fixed-proportions” model. The underlying assumption in this case 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 for different municipalities, but they are fixed within the municipality. If the importance of substitution patterns within municipalities is not central to the present problem, the “fixed-proportions” model should provide a suitable estimate of the demand for deforestation.

The second index, $q_m^c(s)$, is associated with a “vertical” multinomial choice model: once the land is cleared for agriculture, it is used only for the representative crops. In this case, farmers do not substitute forest for pasture. The “vertical” model severely restricts the substitution patterns among land uses and provides a conservative upper bound on the demand for deforestation. To see why, note first that the counterfactual share of agricultural land $Y_m(s, t)$ is increasing in $q_m(s)$ (see Section 4.1 in the main paper). Because the production of beef is land-intensive, $q_m^{cp}(s)$ must be smaller than $q_m^c(s)$ (as it is in the data). Therefore, the demand curve based on the “vertical” model should be above the demand curve based on the “fixed-proportions” model, and also should be above any demand curve that allows for rich substitution patterns between forest and pasture.

In general, rich substitution patterns can be recovered by exploiting choice-specific variables

that shift the value of each land use independently of the value of the other options (Berry and Haile (2014)). However, there is no variable satisfying such a requirement in the present data set. Note that although it is possible to estimate a parametric multinomial choice model in the absence of the choice-specific variables, the estimated parametric model may be inconsistent with the underlying model associated with the productivity indices; and the indices are necessary to recover the demand for deforestation. For this reason, I choose to be agnostic in terms of the way the agricultural area is divided when estimating the impacts of TC_m on deforestation, and experimented with the two different indices associated with the different multinomial models.

Demand for Deforestation. Figure 2 compares the demand for deforestation based on the “fixed-proportions” model and the demand based on the “vertical” model. The demand based on the “vertical” model indeed lies above the curve based on the “fixed-proportions” model, implying smaller impacts from taxes. As mentioned previously, the “vertical” demand curve provides an upper-bound for any demand curve with flexible substitution patterns among the different agricultural land uses. It results in a very conservative upper bound because the assumption that all forests can only be converted to cropland (which in turn can then be converted to pasture) is too restrictive for the Brazilian Amazon.⁷

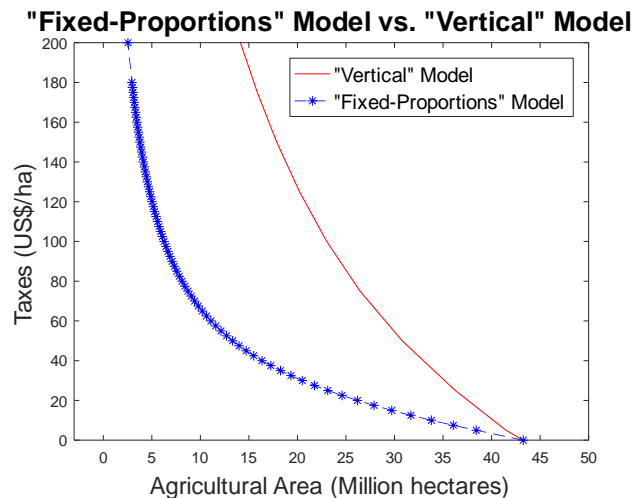


Figure 2: Demand for Deforestation – Different Specifications

⁷Based on remote-sensing data, Almeida *et al.* (2016) provide evidence that a significant fraction of deforested areas was converted directly to pasture. I do not present the supply curve of avoided emissions for the “vertical” model because it provides an extremely conservative bound that is not very informative for carbon emissions.

Yields Regression. I now investigate whether the local yields respond to transportation costs. Specifically, I estimate the model

$$q_m(s) = X_m \gamma_s + \delta_s TC_m + \varepsilon_{m,s}^q$$

for both $q_m^{cp}(s)$ and $q_m^c(s)$, using the same set of covariates and instrumental variables as in the land-use regressions. Recall that the crops-pasture index $q_m^{cp}(s)$ corresponds to the main specification in the paper (the “fixed-proportions” model).

Tables 8 and 9 report the estimated coefficients on the transportation costs and omit the coefficients on the other regressors. Table 8 presents results for the crops-pasture index, $q_m^{cp}(s)$. Except for the small farms, all estimated coefficients on costs to ports are insignificant and small in magnitude.

Table 9 presents the results for the crops-index, $q_m^c(s)$, as well as the results for each individual crops. The estimated effects of transportation costs on $q_m^c(s)$ are statistically significant (except for small to medium-sized farms). However, when separately regressing the yields of each individual product on covariates, I obtained impacts of transportation costs that are always insignificant. This conforms with Roberts and Schlenker’s (2013) results: farmers seem to adjust the extensive margin (land use), but not the intensive margin (yields), in response to changes in transportation costs. The significant coefficients observed in Table 9 are therefore explained by the impacts on the weights of the indices, i.e., on shares of land use. This suggests some substitution patterns among land uses. However, as previously discussed, identifying rich substitution patterns in the present context is difficult. I conclude that the response of local yields to transportation costs is not a first-order issue.⁸

3 Data

This section provides a detailed explanation about how the share of deforestation, the transportation costs, and the productivity indices were constructed. It also provides a brief explanation about the covariates and other useful variables that were not utilized in the regressions, but that were helpful in providing a sense of magnitude for the results.

⁸The number of observations in the yields regressions are slightly smaller than in the land use regressions because of the presence of some outliers. The yields outliers are due to manioc: few observations report an unreasonably large amount of manioc produced. I consider manioc yields above 10 ton/ha measurement error. Because manioc is typically produced for subsistence, it is conceivable that is more likely subject to measurement error problems.

The richest data set available for the agricultural sector is the Brazilian Agricultural Census, produced by the IBGE (Brazilian Institute of Geography and Statistics).⁹ The unit of analysis in the census is the “agricultural establishment,” be it a household or firm producing any animal or plant output. The data are aggregated up to the municipal level. There are 523 municipalities in the data set. All municipalities with a positive fraction of their area in the Amazon Biome and with complete information about all covariates are included in the sample.

3.1 Dependent Variable: Deforestation

The land use in the Agricultural Census is divided into the following categories: annual cropland, perennial cropland, pasture (planted and natural), forest (planted and natural), ponds and lakes, constructions, degraded land, and unusable land (for economic activities). I define agricultural land as the sum of annual and perennial cropland, pasture, constructions and degraded land. Forested area is the sum of the remaining land uses, which includes managed forests and forests that are not currently being exploited.

The online information available for the Agricultural Census provides land use data separately for 18 classes of farm size. For confidentiality reasons, the land use data are missing whenever the municipality has only one or two farmers in a given class. Yet, by aggregating different classes of farm sizes so that the final aggregated class has more than two farmers in the municipality, it is possible to recover the total land-use information. I therefore aggregated the 18 classes into four in trying to achieve a good balance among: (a) not losing too much information from the missing data, (b) obtaining more or less homogeneous classes, and (c) easily exposing the results. The final group classification is: (i) small farms (those with fewer than 5 hectares); (ii) small-to-medium farms (those with an area between 5 and 50 hectares); (iii) medium-to-large farms (those with an area between 50 and 500 hectares); and (iv) large farms (those with more than 500 hectares). For each group of farm size, the share of deforestation in a municipality divides the total agricultural area of the farms by the total area they occupy.

3.1.1 A Comparison of Census Data and Satellite Data

Because the existing policies restrict the amount of deforestation on private lands, one may wonder whether farmers underreport their share of agricultural land in the census. To investigate whether this is an important issue, I contacted the IBGE staff for further clarifications about the census

⁹ Available at <http://www.ibge.gov.br/home/estatistica/economia/agropecuaria/censoagro/default.shtm> and <http://www.sidra.ibge.gov.br/bda/agric>

methodology; I also contacted staff at the Brazilian National Institute of Space Research (INPE) who are responsible for the official measurement of deforestation based on satellite imagery; and I compared the census and the satellite deforestation data myself.

Agricultural Census Data. In personal communication with the IBGE staff, they told me that, although they do not select a subset of farmers to serve as validation data for the Agricultural Census, they do verify the information by comparing the census with several other sources of data that are collected regularly (the oldest survey started in 1935, and the most recent, in 1972). The other data sets are the Systematic Survey of Agricultural Production, the Municipal Agricultural Production Survey, the Extraction of Forest Products and Silviculture Survey, the Municipal Livestock Production Survey, and the Quarterly Survey of Livestock Slaughter.¹⁰ In all surveys, IBGE verifies the information by checking first for filling errors (such as fields left blank or invalid values) and calculation errors (such as planted area smaller than harvested area). Based on central tendency measures (mean, median, mode), IBGE constructs acceptable ranges to correct for outliers. It also checks variables' percentage variation over time to detect extremely discrepant differences. Information on weather conditions and any records of abnormalities that may have affected the normal crop evolution are taken into account (IBGE (2002)). The staff argue that the local IBGE agencies have interacted with farmers systematically and continuously and so accumulated substantial local knowledge (there exist more than 500 IBGE agencies in the Brazilian territory, and each agency is responsible for collecting data covering approximately 11 municipalities on average). They also argue that farmers have little incentive to misreport because (a) they are informed that the census is intended solely for statistical purposes and cannot legally be used against farmers; and (b) farmers are themselves interested in having access to accurate IBGE statistics to make well-informed planting decisions.¹¹

¹⁰The *Systematic Survey of Agricultural Production* is a monthly survey that forecasts and monitors the performance of agricultural crops for the current year with the objective of helping public policies. It provides planted and harvested area estimates, production, and average yield at the municipal level, and this information is available to the public only aggregated at the state level. The *Municipal Agricultural Production Survey* collects detailed and disaggregated annual information on crop area planted and harvested, quantity, and production value for several crops at the municipality level. The *Extraction of Forest Products and Silviculture Survey* provides annual municipal-level information on the quantity and value of production obtained by the exploitation of natural forest resources (timber, medicines, vegetal oil, fruits, etc.), as well as the exploitation of planted forest regions. The *Municipal Livestock Production Survey* collects annual municipal-level information on the number of cattle, pigs, chickens and other animals, as well as animal products (such as milk and wool). Finally, the *Quarterly Survey of Livestock Slaughter* provides quarterly data on the number of animals slaughtered and the total carcass weight per animal species investigated. In all cases, the information is obtained from local farmers, and complemented by cooperatives, agribusinesses, companies that provide technical assistance, slaughterhouses, etc. (IBGE (2002)).

¹¹IBGE staff also told me that the Agricultural Census follows the United Nations' recommendations established by the Food and Agriculture Organization (FAO) contained in the "Programa del Censo Agropecuario Mundial 2010"

Although IBGE’s answers are encouraging, I investigate further whether misreporting is a first-order problem by comparing the deforestation measurements based on satellite and census data. Before presenting the comparison, a brief explanation about how INPE estimates deforestation using remote-sensing data is in order.

Satellite Sensor Data. INPE uses the Landsat Thematic Mapper images (TM/Landsat), with spatial resolution of 60×60 meters, to calculate the following land covers in the Brazilian Amazon: forest, deforested, non-forest (mostly cerrado, which is similar to savanna), hydrography, clouds, and unobserved. Land cover classification is performed in several steps. First, INPE estimates a linear spectral mixture model for each pixel in the data to obtain the pixel’s fraction of different components that help predict land cover.¹² It then groups adjacent pixels in larger regions based on their spectral similarities. After the image segmentation, it implements a cluster unsupervised classification algorithm to generate the land cover classifications. Finally, it verifies the classifications and calculates the deforestation rate.¹³

There are several differences between the deforestation measured by the census and that measured by the interpretation of satellite imagery. First, while the census data capture land use within private properties, satellite data cannot distinguish between deforestation on private and public land. Farms are not geo-referenced in the census data, which makes it impossible to obtain deforestation on private land using satellite images. As explained in the main text, this is the reason why I opted for the census data in the econometric analysis.

Second, the census data include the short-term fallow (land not used for four years) in either the crop land or the pasture land, depending on the farmer’s activity; while the long-term fallow

(available at <http://www.fao.org/economic/ess/ess-wca/es/>).

¹²The pixel components they consider are soil, vegetation, and shade. The images-fraction shade and soil help in the process of identification of deforested areas; images-fraction shade is helpful for areas dominated by tropical forests due to the various strata in the forest structure and the irregularity of the canopy, which contrasts with a low amount of shade in deforested areas; and finally the images-fraction soil helps detect transition/contact areas between forest formations and those of cerrado. Formally, the linear spectral mixture model is the following: for each pixel, let $i = 1, \dots, I$ be the number of spectral bands (e.g., blue, red, near-infrared, etc.); and $j = 1, \dots, J$ be the number of components assumed for the problem (soil, shadow, and vegetation). Assume the linear model $R_i = a_1 X_{i1} + \dots + a_J X_{iJ} + e_i$, where R_i is the mean spectral reflectance for the i -th spectral band (e.g., how much the pixel reflects red in the sensor data); X_{ij} is the spectral reflectance of the j -th component in the pixel for the i -th spectral band (e.g., how much soil reflects red – which is assumed known, obtained from other studies); a_j is the proportion of the j -th component in the pixel; this is the parameter of interest. Assuming all a_j are positive and add up to 1, a constrained least squares is used to estimate the coefficients a_j (Camara, Valeriano, and Soares (2006)).

¹³Several photointerpreters have the task of analyzing the thematic polygons generated by the classifications (directly on the computer screen having as background, for comparability, the original image in color composition). The classified polygons are accepted or reclassified into other land use categories based on the experience of the photointerpreter, whose evaluation depends on specific contexts and, if necessary, relies on historical data (Camara, Valeriano, and Soares (2006)).

(land not used for more than four years) is included in the forested land. INPE, in contrast, considers deforestation irreversible, so that land converted back to forest is counted as deforested. In this sense, the satellite measure underestimates the amount of forest. To have a better sense of whether the underestimation is substantial, I make use of the TerraClass data, which consider a finer classification of land uses (including forest regrowth) and which have been released recently by INPE.¹⁴ According to this data, forest regrowth is responsible for about 22 percent of the total accumulated deforested area. Although highly informative, TerraClass cannot be used directly to compare deforested versus forested areas and the census data for two reasons: first, TerraClass data cover every two years from 2004 to 2014, except for 2006 (exactly the year of the agricultural census). Second, forest regrowth classification is not based on biologic methods, and so it includes various stages of development, from initial stages, when the canopy is homogeneous and has few species, to advanced stages, when canopy heterogeneity and species diversity are similar to original forest (Almeida *et al.* (2016)).¹⁵

In addition to forest regrowth, deforestation measured by INPE includes more categories than the agricultural land in the census. It includes urban area, mining, and others (such as rock outcrops, river beaches, and sandbars). However, these land covers are responsible for only about 2 percent of the total accumulated deforested land, according to the TerraClass data.

Finally, while the census data may suffer from misreporting, satellite data may suffer from prediction errors. There are three reasons for prediction errors. First, cloud cover prevents the classification of true land use; this is an important issue for the Amazon rainforest. Second, it is common to include a residual classification to allow for cases in which the classifier algorithm is not able to clearly identify the land cover. INPE includes such a class, denoted by “unobserved area,” that essentially results from poor radiometric quality of the sensor data. The average proportion of the sum of clouds and “unobserved area” per municipality in my data is 10 percent, which is a non-negligible share of municipality areas.¹⁶ Finally, there are the unavoidable statistical misclassification errors. Unfortunately, INPE does not report the land cover misclassification probabilities, so I cannot assess the accuracy of their deforestation measurement.

¹⁴TerraClass data are available at http://www.inpe.br/cra/ingles/project_research/terraclass.php

¹⁵The spectral similarities in the reflectance of forests at different stages of succession and regeneration lead to high classification errors when one tries to distinguish secondary forests from deforested areas, cerrado vegetation, and mature forests. Assessments of forest regrowth based on remote-sensing data are commonly viewed as a nontrivial difficult task (Caviglia-Harris *et al.* (2015)).

¹⁶Clouds result in a missing data problem. However, the data are not exactly missing-at-random because some areas are covered by clouds systematically more often than others. The data can be more convincingly considered missing-at-random only if *conditioned* on the latitude and longitude of a location.

Direct Comparison. Although the two deforestation measures are not entirely comparable, in order to make a comparison that is as clean as possible, I proceed as follows: first, I calculate the municipalities’ share of deforestation according to both census and satellite data sets. Then I compare the two measures in municipalities in which both the share of public land and the share of clouds/unobserved areas are small. Specifically, I consider municipalities with more than 90 percent of private land and with less than 3 percent of clouds/unobserved areas.¹⁷ The idea is that in municipalities that have negligible share of public land and that are not contaminated by clouds or by “unobserved areas,” the two measures should be close to each other.¹⁸

As expected, the correlation between the two deforestation measurements increases as we move from the entire sample to the restricted sample: the correlation increases from 0.74 to 0.84. The average shares of deforestation also increase and get closer to each other when I use the restricted sample, again as expected. The average share of deforestation in the full sample according to the census is 27 percent, while according to the satellite data it is 42 percent (which is consistent with the previous discussion). In the restricted sample, the census average share is 68 percent, and the satellite average share is 75 percent. Although the gap is reduced by half, there is still a non-negligible difference. Yet, when I subtract the 2008 TerraClass data’s forest regrowth from the INPE deforestation, I obtain an average deforestation of 65 percent, which is much closer to the census average. Although TerraClass data are not directly comparable to INPE or census data (as I mentioned earlier), these results are encouraging.

Perhaps more important than comparing raw correlations and simple averages is to investigate whether the census deforestation systematically varies with exogenous variables (particularly with the instrumental variables) *after* controlling for the satellite deforestation. The idea is that, if we consider the satellite data to provide a good measure of deforestation, then, after conditioning on it, the exogenous variables should contain no additional information about deforestation in the census data.

To investigate this possibility, I regress the census deforestation measure on satellite deforestation and other explanatory variables. In this case, I augment the restricted sample to include municipalities with more than 80 percent of private land (because the sample with more than 90 percent of private land does not have enough degrees of freedom). Table 10 presents the results. The first column presents the estimated coefficients for the full sample, and the second column, for

¹⁷To minimize potential problems with INPE misclassification probabilities, the restricted sample also ignores municipalities in which the sum of land areas classified in INPE is greater than or equal to the municipality area. Correlations are not very sensitive to this restriction.

¹⁸I am grateful to an anonymous referee who suggested such comparison.

the restricted sample.¹⁹

As expected, the estimated coefficients on the satellite deforestation measure are positive and statistically significant in both samples. The estimated coefficients on most of the other explanatory variables are not statistically significant in the full sample, including the coefficients on the instrumental variables (distance to port and distance to capital), which are very small in magnitude. In the restricted sample, when the census and satellite measures are most comparable, none of the regressors' coefficients are significant except for the coefficient on satellite deforestation.²⁰

The correlations I find are consistent with recent results obtained by Caviglia-Harris *et al.* (2015). To the best of my knowledge, theirs is the only published paper that combines and compares survey data (with geo-referenced household property boundaries) and remote-sensing data for the Brazilian Amazon. They perform a series of reliability tests, and find a high degree of consistency between households' responses and satellite-derived land cover.

3.2 Endogenous Regressor: Transportation Costs

The proxy for transportation costs is defined as the minimum unit cost (US\$/ton) to transport 1 ton of goods to the nearest port. This cost was 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 the freight rate data collected by SIFRECA (the Freight Information System). The implementation used ArcGIS cost distance tools – see Allen and Arkolakis (2014) for an excellent description of this type of algorithm. Calculating this proxy requires definitions about (a) which ports are included in the calculations, and (b) how to combine the freight rate data with the network of roads, railroads, and navigable rivers.

As discussed in the main paper, when directed to international markets, the main products in the Amazon normally use either the ports in the Amazon (Port of Santana, Port of Belém and Port of Itaquí) or the ports in the South (Port of Santos and Port of Paranaguá); see Figure 2 in the main paper. Therefore, I selected these five ports to be the main destinations.

Given the selected destinations, the least cost path to the nearest port is computed in ArcGIS.

¹⁹The regressions include the proportion of clouds/unobserved areas among the covariates to further mitigate measurement problems in the satellite deforestation. Note that the coefficient on clouds/unobserved areas increases substantially in the restricted sample; this is because the scale of that covariate changes from an average of 10 percent in the full sample to an average of 0.0004634 percent in the restricted sample. For both samples, I used the Eicker-White robust standard errors. I obtained qualitatively similar results when I used the spatially dependent HAC robust standard error proposed by Conley (1999). For the spatially dependent robust standard errors, I used a uniform spatial kernel, and a 75 km distance cutoff. The results do not change when cutoffs are 50 km or 100 km.

²⁰Results are qualitatively similar when I subtract the TerraClass 2008 forest regrowth data from satellite deforestation.

The calculation divides the entire country into cells corresponding to 1 km², 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. Given these travel costs, the optimization routine in ArcGIS determines the least cumulative cost path from each origin to the nearest destination.²¹ 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 section (see below). 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 1 ton of soybeans to the nearest port. Note that, although the costs to transport different products may differ, they should be highly collinear: all products use the same transportation network and reach the same ports (under the no-arbitrage condition). Furthermore, bulk products and sacks must have the same transportation costs as soybeans (Castro (2003)).²²

Table 11 summarizes the unit costs used to compute the least accumulative cost in ArcGIS. The first column discriminates between the possible modes of transportation considered in the calculations; the second column reports the unit travel costs used in Brazilian currency (R\$); and the third column converts these unit costs into US\$. Note that differences in unit costs are in the expected direction, as waterways with good infrastructure are the cheapest mode of transportation, followed by railroads and paved roads. Unpaved roads inside the Amazon are the worst mode of transportation followed by navigable rivers with poor infrastructure and unpaved roads outside the Amazon. Finally, because ArcGIS allows for traveling by land, I imposed high costs to transport goods by land with no mode of transportation so that ArcGIS would avoid computing traveling costs using these cells.

Next, I discuss the unit costs for each mode of transportation.

²¹The network of modes of transport are in a polyline format in ArcGIS while the country is in a polygon format. I first have to transform these polylines and polygons into a raster format (i.e., a grid of cells with cost values). Then I use the command “cost distance” in ArcGIS to compute the least cumulative cost path to the nearest destination and “cost allocation” to identify which destination is the nearest for each cell. Finally, the command “extract values to points” is used to assign the total costs to the corresponding municipal seats.

²²I am grateful to Prof. Newton de Castro, who suggested the use of the World Bank’s HDM-VOC-4 model and the data from SIFRECA to compute transportation costs.

Roads. I purchased the freight values of routes from cities in the Legal Amazon to state capitals and ports for the main products in the Amazon.²³ I obtained the costs and distances to transport soybeans by roads for 105 routes and five destinations for 2006. The average cost to transport 1 ton of soybeans per km in these data is R\$ 0.0767 (US\$ 0.0353). That is the value I inputted as the unit cost to travel one cell with a paved road.²⁴

To capture the higher costs of traveling on unpaved roads, I increased the unit costs for paved roads by using the World Bank’s HDM-VOC-4. This model is designed to calculate unit road user costs for a road section with 1 km length and requires several inputs for the characteristics of the road. I maintained all inputs at the default values, except for changing the road characteristics from paved to unpaved and increasing the roughness of the road to the recommended value for a poor tertiary road. For heavy truck vehicles, the increase in the roughness raised the road user costs by 29 percent. I adopted the unit cost to travel 1 km on unpaved road to be 29 percent higher than the cost to travel on a paved road. Therefore, a cost of R\$ 0.0992 (US\$ 0.0457) to transport 1 ton of soybeans per km was assigned to unpaved roads.²⁵

Transporting large quantities on unpaved roads within the rainforest is extremely difficult. The poor conditions of the roads, combined with the excess of rain, especially in the rainy season, can make these roads inaccessible. For this reason, I decided to differentiate unpaved roads within the rainforest from the unpaved roads elsewhere. Unfortunately, the freight rate data purchased from SIFRECA do not cover the dense rainforest. To overcome this limitation, I called local companies directly and asked them how much it costs to transport soybeans from Sorriso (in Mato Grosso state), one of the main producers of soybeans, to the Port of Santarém on the Amazon River. The only route available in this case is the “Cuiabá-Santarém” road, which cuts the Amazon almost in the middle from the South to the North and provides a good measure of the difficulties in traveling in the dense jungle. The average cost per km that I obtained is R\$ 0.15 (US\$ 0.069). I used this unit cost for all unpaved roads within the dense rainforest.²⁶

²³The Legal Amazon is an administrative area in the northern part of Brazil that includes nine states and around 5 million km² of land (about 60% of the Brazilian national territory). The states are indicated in Figure 1 in the main paper. It consisted of 771 municipalities in 2006, and includes the Amazon Biome as well as other types of vegetation, in particular a savannah type of vegetation called the *cerrado*.

²⁴In SIFRECA’s sample, almost all cities were located in the state of Mato Grosso, the main producer of soybeans.

²⁵The International Roughness Index (IRI) is an index that measures the deviations of a surface from a true planar surface with characteristic dimensions that affect vehicle dynamics, ride quality, dynamic loads, and drainage. It is measured in m/km units. The value recommended for a good primary paved road is 2 m/km. The value I used in order to increase the costs of unpaved roads is 8 m/km, which corresponds to a poor tertiary road. The HDM-VOC-4 model is available at <http://web.worldbank.org/WBSITE/EXTERNAL>

²⁶I define the dense rainforest as covering the states of Acre, Amazonas, Rondônia, Roraima, Pará, and Amapá. The remaining states in the Legal Amazon (Maranhão, Tocantins, and Mato Grosso) are not in the dense rainforest.

Railroads and Navigable Rivers. Freight values for navigable rivers and railroads are more difficult to obtain because of the reluctance of the firms to disclose their data. The information I was able to obtain from SIFRECA includes the freight values to transport 1 ton of soybeans in two routes for a railroad, with an average value of R\$ 0.0608 (US\$ 0.0279), and one route for one of the most important navigable rivers in the Amazon – the Madeira River waterway. Because of the difficulty in collecting these data, SIFRECA does not provide recent freight rate data for these modes of transportation anymore.²⁷

Similar to roads, there are differences in the quality of navigable rivers, depending on the depth of the river, investments in signaling, investments in communications, and in the quality of the local ports. Based on conversations with governmental agencies responsible for the administration of the waterways, as well as local companies, I classified the navigable rivers in two types: those with good and those with poor infrastructure. The good rivers include the Madeira River waterway and the Amazon River waterway (linking Manaus to Belém), among a few others. The rivers with poor infrastructure are the remaining navigable rivers.²⁸

From SIFRECA’s information of the freight cost of the Madeira River waterway combined with the information I obtained directly from the companies for the cost of the Amazon River waterway (from Manaus to Belém), I arrived at a unit cost of R\$ 0.0444 (US\$ 0.0204) to transport 1 ton of soybeans per km. For the navigable rivers with poor infrastructure, I obtained an average value of R\$ 0.1139 (US\$ 0.0525) per ton of soybeans per km.²⁹

Land. Finally, I imposed high costs to transport products by land with no mode of transportation so that ArcGIS avoids computing traveling costs using these cells. For the Amazon Biome, I imposed a cost of R\$ 3 per km (US\$ 1.38) and for land outside the Amazon Biome, a cost of R\$ 1.5 (US\$ 0.69). The rationale is that moving within the rainforest should be much more costly than for other types of vegetation. There is little guidance on which values should be adopted for transportation

²⁷The routes for railroads are from Cascavel to Ponta Grossa (both in the Paraná state in the South) and from Porto Franco to São Luís (both in the Maranhão state in the Legal Amazon). The Madeira River waterway connects Porto Velho (the capital of the Rondônia state) to Itacoatiara, which is close to Manaus (the capital of the Amazonas state).

²⁸The Ministry of Transportation classifies two types of waterway: *more-navigable* and *less-navigable* rivers. I excluded the *less-navigable* rivers, because they do not seem to be used to transport large quantities of products. Hence, the classification of good- and poor-infrastructure rivers is restricted only to those *more-navigable* rivers.

²⁹The unit cost for poor-infrastructure rivers is obtained from the routes Manaus–Tabatinga, Manaus–Barcelos and Manaus–Boca do Acre, all of them in pristine regions in the Western Amazon. The larger costs obtained reflect not only poor signaling and capacity constraints, but also difficulties with the excess of curves, the depth of rivers (some may have about 1 meter depth in the dry season, while the Amazon River has, on average, 16 meters’ depth), as well as the presence of stones and sandbars that shift around over time. To gain a sense of magnitude, it can take about 10 days to go from Tabatinga (located at the border with Venezuela) to Manaus and about 18 days to navigate that stretch of the river in the opposite direction.

in land, but, because the municipal seats are connected to some segment of the network and to the extent that these unit cost values induce ArcGIS to use cells with some mode of transportation, the costs of moving on land should not impact the results significantly.³⁰

Figure 3 presents the map of transportation costs calculated in ArcGIS. The figure also presents the network of modes of transportation and the municipal seats in the Legal Amazon in the right panel. The darker the color in the map, the higher the transportation cost.

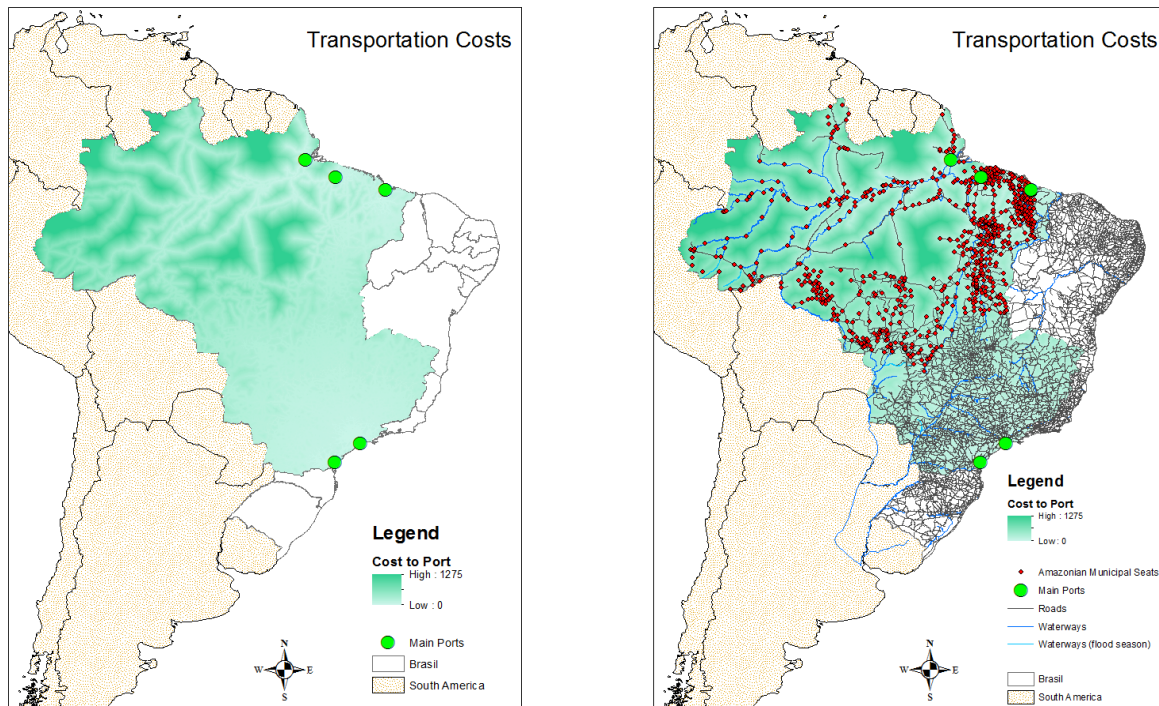


Figure 3: Transportation Costs and Transportation Network

3.3 Productivity Index

From the Agricultural Census, I obtained the quantity sold and the area occupied of major agricultural outputs for each municipality for farms of different sizes. I calculated two productivity indices: one considers only the production of crops, while the other considers the production of both crops and beef.

³⁰For those cities not connected to any mode of network in the map, I created straight lines joining them to the nearest road and assigned the cost of unpaved road (depending on whether they are inside or outside the rainforest) to travel over these straight lines. This procedure is reasonable, because the official map from the National Highway Plan misses the unofficial roads, and they are most likely unpaved.

I describe first the index based on crops only. I selected the main crops discussed in the paper: soy, corn, manioc, rice, and beans. For each product j , for each municipality m , and for each farm size s , I calculated the quantity of output sold per hectare for each product. The yields are denoted by q_{jm}^s , for $j = 1, \dots, 5$. The productivity index is the weighted average of the of q_{jm}^s across j , where the weights are the proportion of the area utilized for each crop, denoted by a_{jm}^s . The index for the crops is therefore:

$$q_m^c(s) = \sum_{j=1}^5 \left(\frac{a_{jm}^s}{a_m^s} \right) \times (q_{jm}^s)$$

where $a_m^s = \sum_{j=1}^5 a_{jm}^s$.

To add pasture in the productivity index, I first assumed that each ox weighs half a ton and that the entire ox can be used for beef consumption. I therefore multiplied the number of cattle sold by 0.5 to obtain the quantity of beef sold in tons.³¹ By dividing this result by the pasture area, I obtained the productivity for cattle, q_m^p . I do not have information on the number of oxen for different farm sizes, so q_m^p is the same for all farm sizes.

According to the freight data, it is 30 percent more expensive to transport frozen meat than soybeans. For this reason, I increased the weight for the beef in the productivity index by 30 percent so that the transportation costs of all products are measured in terms of the costs to transport 1 ton of soybeans. Let a_{mp}^s be the area occupied by pasture for farms of size s in location m . The second productivity index is a weighted average between q_m^c and q_m^p , where the weights are the proportion of crops and pasture areas:

$$q_m^{cp}(s) = \left(\frac{a_m^s}{a_m^s + a_{mp}^s} \right) q_m^c(s) + \left(\frac{a_{mp}^s}{a_m^s + a_{mp}^s} \right) (1.3) q_m^p$$

I averaged the indices over the micro-regions to reduce measurement error problems. Micro-regions are administrative areas larger than the municipalities; there are 85 micro-regions in the Amazon.

3.4 Covariates

Next, I briefly describe the set of covariates and the instrumental variables.³²

³¹Informal conversations with farmers suggest that a common rule of thumb is to assume a cow has about 150 kg of “available meat” and a bull has about 270 kg. Because a bull weighs around 500 kg, I am overestimating the productivity for pastures.

³²Some covariates were available at the municipal level for years other than 2006. Because municipal boundaries changed over time, I had to match the previous boundaries to convert the values of these covariates to the boundaries observed in 2006. I summed or averaged the values over the municipalities in 2006, taking a weighted average where appropriate.

Temperature and Precipitation. The Climate Research Unit (CRU) computed the annual average temperature and precipitation based on the average climate from 1961 to 1990. Given the high levels of temperature and rain in the region, one might expect these variables to be negatively correlated with deforestation.³³

Altitude. This variable is prepared by the IBGE. Although high altitudes may be good for the production of rice, I do not have *a priori* information about whether this variable should be positively or negatively correlated with deforestation.³⁴

Slope. To calculate the average slope in a municipality, I used data from the Shuttle Radar Topographic Mission (SRTM). Slope is the inclination of a surface in relation to the horizontal surface, and it is measured in degrees. One should expect places with steeper slopes to be worse for agriculture in the Amazon because of water runoff and because mechanization is more difficult.³⁵

Soil quality. These data are produced by IBGE and EMBRAPA (the Brazilian Agricultural Research Corporation) and were kindly made available by Professor Eustáquio Reis. They consist of the proportion of the municipal area on each of five aptitude classes of soil. The classes of soil aptitude for agriculture are: high quality, medium-to-high, low-to-medium, low quality, and unsuitable. The five aptitude classes were aggregated from the 13 soil types that exist in the Brazilian territory. Using the 13 soil types in conjunction with the data on local topography, data from ground surveys, and general familiarity with the land, EMBRAPA soil scientists created the digital map of soil aptitude for agriculture that is used here. Factors that could lead to a low ranking include high metal content, poor drainage, high flood risk, uneven ground, low nutrients, and steep slope (Anderson and Reis (2007)).

Local Population. These data come from the Demographic Census of 2000 produced by IBGE. In principle, the local population may affect local demand for non-tradables and local wages. Land use on private land data, however, is affected by the local population depending upon whether agriculture is more or less labor-intensive than extractive activities.³⁶

³³The data are available at <http://www.ipeadata.gov.br/>

³⁴The data are available at <http://www.ipeadata.gov.br/>

³⁵The data are available at <http://www.dpi.inpe.br/Ambdata/English/index.php>

³⁶Available at http://www.ibge.gov.br/home/estatistica/populacao/default_censo_2000.shtm

Presence of Local Mining and Power Plants. These data are available from the National Highway Plan. The presence of mining and power plant should shift the demand for non-tradables. To capture the effect that electrification may have on the neighborhood of power plants, I also created a dummy variable that equals one if a power plant exists within a 75 km radius of the municipal seat (but is not within the municipality itself), and zero otherwise.

Distance to Protected Areas. These data are available from the National Highway Plan. I computed straight-line distances from the municipal seats to the nearest protected area polygon using ArcGIS. Protected areas in the Amazon rainforest include indigenous reserves and conservation units (which in turn include national and state forests, wildlife reserves, extractive reserves, ecological stations, areas of relevant ecological interest, and national parks, among others).

Distance to IBAMA. I computed straight-line distances from the municipal seats to the nearest IBAMA agency using ArcGIS.

Number of Fines. The total number of fines includes all environmental infractions in a municipality for each year since 2002. The data are publicly available from IBAMA. For each municipality, I added all fines starting in the first year the data are available up to the year before the Agricultural Census.

Proportion of Private Land with Land Title. These data come from the Agricultural Census and measure the share of the private land that has a land title. The land with no title is either occupied by squatters or covered under one of the official settlement programs, which have not yet given the definitive titles to the farmers.

Instrumental Variables. I computed straight-line distances from the municipal seats to the nearest port and to the nearest capital using ArcGIS. I excluded Palmas, the capital of the Tocantins state, from the destinations for the straight-line distances because it is a planned city that was built in 1989 with the objective of helping to develop the region.

Carbon Stock. The amount of aboveground carbon stock is calculated by Baccini *et al.* (2012). I combined their raster data of carbon stock with the 2008 TerraClass shapefile of land uses in the Amazon to calculate the average carbon stock in forested and deforested areas in each municipality.

There are 18 municipalities with missing data on the carbon stock, almost all of them in the state of Maranhão, in the Eastern Amazon.

Other Useful Variables. Useful variables that are not utilized in the estimation procedure include:

Satellite Data. Satellite-based measures of deforestation in the Brazilian Legal Amazon are calculated by INPE (the Brazilian National Institute of Space Research). They include both the PRODES data (Monitoring Program of the Brazilian Amazon by Satellite), and the TerraClass data.³⁷

Labor. The Agricultural Census provides the number of workers in the farms (separating family member workers from hired workers), as well as rural wages.

Revenue and Expenditures. The Agricultural Census also contains information on both the revenues and the expenditures of the agricultural establishments. The revenue is the value from the sale of production and the expenditures include maintenance costs, salaries, rentals of machinery, and other expenses. Revenues and costs are not discriminated for each land use.

References

- [1] Allen, T. and C. Arkolakis (2014). “Trade and the Topography of the Spatial Economy,” *Quarterly Journal of Economics*, 129(3), 1085–1140.
- [2] 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.
- [3] Anderson, K. and E. Reis (2007). “The Effects of Climate Change on Brazilian Agricultural Profitability and Land Use: Cross-Sectional Model with Census Data,” Final report to WHRC/IPAM for LBA project Global Warming, Land Use, and Land Cover Changes in Brazil.
- [4] Assunção, J., M. Lipscomb, A. M. Mobarak, and D. Szerman. (2016). “Agricultural Productivity and Deforestation in Brazil,” Mimeo, Yale University.

³⁷The methodology is explained by Câmara, Valeriano, and Soares (2006). The data are available at <http://www.obt.inpe.br/prodes/index.html>

- [5] 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.
- [6] Berry, S. T. and P. A. Haile (2014). “Identification in Differentiated Products Markets Using Market Level Data,” *Econometrica*, 82(5), 1749–1797.
- [7] Camara, G., D. M. Valeriano, and J. V. Soares (2006). “Metodologia para o Cálculo da Taxa Anual de Desmatamento na Amazônia Legal,” São José dos Campos, INPE.
- [8] Castro, N. (2003). “Formação de Preços no Transporte de Carga,” *Pesquisa e Planejamento Econômico*, 33(1), 167–189.
- [9] Caviglia-Harris, J. L., M. Toomey, D. W. Harris, K. Mullan, A. R. Bell, E. O. Sills, and D. A. Roberts (2015). “Detecting and Interpreting Secondary Forest on an Old Amazonian Frontier,” *Journal of Land Use Science*, 10(4), 442–465.
- [10] Chen, X. (2007). “Large Sample Sieve Estimation of Semi-nonparametric Models,” Chapter 76 in *Handbook of Econometrics*, Vol. 6B, eds. James J. Heckman and Edward E. Leamer, North-Holland.
- [11] Chen, X. and D. Pouzo (2009). “Efficient Estimation of Semiparametric Conditional Moment Models With Possibly Nonsmooth Residuals,” *Journal of Econometrics*, 152, 46–60.
- [12] Chen, X. and D. Pouzo (2012). “Estimation of Nonparametric Conditional Moment Models with Possibly Nonsmooth Generalized Residuals,” *Econometrica*, 80(1), 277–321.
- [13] Chomitz, K. M. and T. S. Thomas (2003). “Determinants of Land Use in Amazônia: A Fine-scale Spatial Analysis,” *American Journal of Agricultural Economics*, 85(4), 1016–1028.
- [14] Conley, T. (1999). “GMM Estimation with Cross-sectional Dependence,” *Journal of Econometrics*. 92, 1–45.
- [15] IBGE (2002). “Pesquisas Agropecuárias,” *Série Relatórios Metodológicos*, Volume 6.
- [16] INPE, Instituto Nacional de Pesquisas Espaciais (2017), *Monitoramento da Floresta Amazônica Brasileira por Satélite – Projeto Prodes*. Available at <http://www.obt.inpe.br/prodes/>

- [17] Le Sage, J. P. (2014). “What Regional Scientists Need to Know About Spatial Econometrics,” *The Review of Regional Studies*, 44(1), 13–32.
- [18] 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. <https://doi.org/10.1371/journal.pone.0129460>
- [19] PNLT, National Highway Plan (2006). Available at <http://www.transportes.gov.br/index/conteudo/id/36604> and accessed on 11/24/2010.
- [20] Robalino, J. A., A. Pfaff (2012). “Contagious Development: Neighbor Interactions in Deforestation,” *Journal of Development Economics*, 97(2), 427–436.
- [21] Robalino, J. A., A. Pfaff, and L. Villalobos-Fiatt (2017). “Heterogeneous Local Spillovers from Protected Areas in Costa Rica,” *Journal of the Association of Environmental and Resource Economists*, 4(3), 795–820.
- [22] 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), 2265–2295.

Table 1: Land-use Median Regression by Farm Size

	Small	Small–Medium	Medium–Large	Large
Cost to Port	0.0045 (0.76)	-0.0053 (-1.58)	-0.0079** (-2.07)	-0.0110*** (-2.91)
Altitude	0.0010 (0.67)	0.0000 (0.74)	-0.0002 (-0.27)	-0.0011 (-1.48)
Temp	-0.0383 (-0.12)	-0.4680*** (-2.79)	-0.2649 (-1.54)	0.0229 (0.11)
Rain	-0.0038 (-0.97)	-0.0094*** (-3.53)	-0.0068** (-2.55)	-0.0051* (-1.65)
Slope	-0.0846 (-0.23)	-0.1015 (-0.50)	-0.2076 (-1.04)	-0.1965 (-0.85)
Soil 2	-2.1386*** (-2.83)	-1.8197*** (-4.25)	-1.5805*** (-3.06)	-1.6984*** (-2.62)
Soil 3	-0.9227 (-0.84)	-0.8729* (-1.93)	-0.9222* (-2.27)	-0.9942** (-2.28)
Soil 4	-1.4279** (-2.41)	-1.4894*** (-4.78)	-1.5151*** (-4.91)	-1.6160*** (-4.50)
Soil 5	-0.0730 (-0.11)	-0.8029** (-2.38)	-0.5831* (-1.84)	-0.6423* (-1.70)
Dist. to PA	-0.0059 (-0.90)	0.0005 (0.16)	0.0028 (1.13)	0.0055* (1.78)
Mining	-0.0136 (-0.31)	0.0061 (0.15)	-0.0432** (-1.99)	0.0075 (0.28)
Power Plant	-0.5107 (-1.29)	-0.2998 (-1.07)	-0.1086 (-0.33)	-0.0445 (-0.14)
Power Plant Neighbor	-0.4703 (-1.33)	-0.3035 (-1.25)	-0.3089 (-1.37)	-0.1575 (-0.63)
Population	-0.0008 (-0.89)	-0.0005 (-0.54)	0.0002 (0.14)	0.0003 (0.42)
Prop. Land Title	-0.3951 (-0.44)	0.4597 (1.02)	0.5894 (0.83)	0.6563 (0.74)
Fines	-0.0009 (-0.48)	-0.0012 (-1.07)	-0.0015 (-0.81)	-0.0015 (-0.94)
Dist. to Ibama	-0.0005 (-0.24)	-0.0009 (-0.74)	0.0013 (1.14)	0.0005 (0.37)
Constant	5.6863 (0.65)	16.317*** (3.41)	10.267** (2.08)	3.4441 (0.60)

Notes: Land-use IVQR regression results for the median. Spatially lagged regressors are omitted from the table. The number of municipalities for small farms is 501; for small–medium is 520; for medium–large, 520; and for large farms, 450.

t-statistics in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Land-use Median Regression by Farm Size – Local Population

	Small		Small–Medium		Medium–Large		Large	
Cost to Port	0.0045 (0.76)	0.0059 (1.00)	-0.0053 (-1.58)	-0.0056* (-1.71)	-0.0079** (-2.07)	-0.0081** (-2.19)	-0.0110*** (-2.91)	-0.0109*** (-2.88)
Altitude	0.0010 (0.67)	0.0011 (0.75)	0.0000 (0.74)	0.0002 (0.29)	-0.0002 (-0.27)	-0.0003 (-0.49)	-0.0011 (-1.48)	-0.0011 (-1.39)
Temp	-0.0383 (-0.12)	-0.0414 (-0.14)	-0.4680*** (-2.79)	-0.4943*** (-2.92)	-0.2649 (-1.54)	-0.2883* (-1.68)	0.0229 (0.11)	-0.0013 (-0.01)
Rain	-0.0038 (-0.97)	-0.0034 (-0.86)	-0.0094*** (-3.53)	-0.0095*** (-3.58)	-0.0068** (-2.55)	-0.0067** (-2.48)	-0.0051* (-1.65)	-0.0054* (-1.73)
Slope	-0.0846 (-0.23)	-0.0188 (-0.05)	-0.1015 (-0.50)	-0.1023 (-0.50)	-0.2076 (-1.04)	-0.2026 (-1.02)	-0.1965 (-0.85)	-0.1949 (-0.85)
Soil 2	-2.1386*** (-2.83)	-2.1002*** (-2.79)	-1.8197*** (-4.25)	-1.7818*** (-4.10)	-1.5805*** (-3.06)	-1.5816*** (-3.08)	-1.6984*** (-2.62)	-1.7085*** (-2.65)
Soil 3	-0.9227 (-0.84)	-0.8473 (-0.78)	-0.8729* (-1.93)	-0.9393** (-2.10)	-0.9222* (-2.27)	-0.9181** (-2.39)	-0.9942** (-2.28)	-0.9793** (-2.25)
Soil 4	-1.4279** (-2.41)	-1.281** (-2.15)	-1.4894*** (-4.78)	-1.4712*** (-4.75)	-1.5151*** (-4.91)	-1.5439*** (-5.23)	-1.6160*** (-4.50)	-1.6292*** (-4.53)
Soil 5	-0.0730 (-0.11)	0.1017 (0.15)	-0.8029** (-2.38)	-0.7753** (-2.32)	-0.5831* (-1.84)	-0.5919** (-1.97)	-0.6423* (-1.70)	-0.6309* (-1.68)
Dist. to PA	-0.0059 (-0.90)	-0.0063 (-0.97)	0.0005 (0.16)	0.0006 (0.20)	0.0028 (1.13)	0.0030 (1.22)	0.0055* (1.78)	0.0055* (1.75)
Mining	-0.0136 (-0.31)	-0.0133 (-0.31)	0.0061 (0.15)	0.0125 (0.29)	-0.0432** (-1.99)	-0.0477** (-2.19)	0.0075 (0.28)	0.0053 (0.20)
Power Plant	-0.5107 (-1.29)	-0.4962 (-1.24)	-0.2998 (-1.07)	-0.3123 (-1.09)	-0.1086 (-0.33)	-0.172 (-0.56)	-0.0445 (-0.14)	0.0201 (0.06)
Power Plant Neighbor	-0.4703 (-1.33)	-0.5181 (-1.47)	-0.3035 (-1.25)	-0.3184 (-1.34)	-0.3089 (-1.37)	-0.2960 (-1.29)	-0.1575 (-0.63)	-0.1136 (-0.45)
Population	-0.0008 (-0.89)	- -	-0.0005 (-0.54)	- -	0.0002 (0.14)	- -	0.0003 (0.42)	- -
Prop. Land Title	-0.3951 (-0.44)	-0.3485 (-0.39)	0.4597 (1.02)	0.4670 (1.02)	0.5894 (0.83)	0.6071 (0.88)	0.6563 (0.74)	0.6556 (0.78)
Fines	-0.0009 (-0.48)	-0.0021 (-1.40)	-0.0012 (-1.07)	-0.0014 (-1.61)	-0.0015 (-0.81)	-0.0012 (-1.52)	-0.0015 (-0.94)	-0.0012 (-0.81)
Dist. to Ibama	-0.0005 (-0.24)	-0.0013 (-0.58)	-0.0009 (-0.74)	-0.0009 (-0.76)	0.0013 (1.14)	0.0012 (1.05)	0.0005 (0.37)	0.0006 (0.43)
Constant	5.6863 (0.65)	5.5114 (0.64)	16.317*** (3.41)	17.005*** (3.53)	10.267** (2.08)	10.888** (2.21)	3.4441 (0.60)	2.9096 (0.50)

Notes: Land-use IVQR regression results for the median. For each farm size, the column in the left presents the main specification, and the column in the right excludes “Population” among the regressors. Spatially lagged regressors are omitted from the table. t-statistics in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Land-use Median Regression by Farm Size – Prop. of Land Title

	Small		Small-Medium		Medium-Large		Large	
Cost to Port	0.0045 (0.76)	0.0055 (0.92)	-0.0053 (-1.58)	-0.0049 (-1.45)	-0.0079** (-2.07)	-0.0088** (-2.27)	-0.0110*** (-2.91)	-0.0119*** (-3.13)
Altitude	0.0010 (0.67)	0.0009 (0.62)	0.0000 (0.74)	0.0001 (0.10)	-0.0002 (-0.27)	-0.0003 (-0.38)	-0.0011 (-1.48)	-0.0009 (-1.27)
Temp	-0.0383 (-0.12)	-0.0284 (-0.09)	-0.4680*** (-2.79)	-0.4294** (-2.54)	-0.2649 (-1.54)	-0.3019* (-1.76)	0.0229 (0.11)	-0.0265 (-0.13)
Rain	-0.0038 (-0.97)	-0.0037 (-0.93)	-0.0094*** (-3.53)	-0.009*** (-3.38)	-0.0068** (-2.55)	-0.0072*** (-2.73)	-0.0051* (-1.65)	-0.0062** (-2.00)
Slope	-0.0846 (-0.23)	-0.0877 (-0.24)	-0.1015 (-0.50)	-0.1799 (-0.85)	-0.2076 (-1.04)	-0.1874 (-0.95)	-0.1965 (-0.85)	-0.2875 (-1.26)
Soil 2	-2.1386*** (-2.83)	-2.1276*** (-2.92)	-1.8197*** (-4.25)	-1.9908*** (-4.89)	-1.5805*** (-3.06)	-1.7028*** (-3.31)	-1.6984*** (-2.62)	-1.8049*** (-2.84)
Soil 3	-0.9227 (-0.84)	-0.8689 (-0.78)	-0.8729* (-1.93)	-.0.8805** (-1.96)	-0.9222* (-2.27)	-1.0041** (-2.52)	-0.9942** (-2.28)	-1.0573** (-2.50)
Soil 4	-1.4279** (-2.41)	-1.4054** (-2.37)	-1.4894*** (-4.78)	-1.5509*** (-4.95)	-1.5151*** (-4.91)	-1.5449 (-4.98)	-1.6160*** (-4.50)	-1.6795*** (-4.68)
Soil 5	-0.0730 (-0.11)	-0.0478 (-0.07)	-0.8029** (-2.38)	-0.9223*** (-2.79)	-0.5831* (-1.84)	-0.6658 (-2.11)	-0.6423* (-1.70)	-0.6728* (-1.80)
Dist. to PA	-0.0059 (-0.90)	-0.006 (-0.92)	0.0005 (0.16)	-0.0001 (-0.03)	0.0028 (1.13)	0.0027 (1.07)	0.0055* (1.78)	0.0056* (1.81)
Mining	-0.0136 (-0.31)	-0.0166 (-0.39)	0.0061 (0.15)	0.0054 (0.12)	-0.0432** (-1.99)	-0.0442** (-2.00)	0.0075 (0.28)	0.0097 (0.36)
Power Plant	-0.5107 (-1.29)	-0.4563 (-1.15)	-0.2998 (-1.07)	-0.3146 (-1.11)	-0.1086 (-0.33)	-0.1318 (-0.38)	-0.0445 (-0.14)	-0.0315 (-0.10)
Power Plant Neighbor	-0.4703 (-1.33)	-0.5514 (-1.57)	-0.3035 (-1.25)	-0.356 (-1.49)	-0.3089 (-1.37)	-0.3057 (-1.34)	-0.1575 (-0.63)	-0.007 (-0.03)
Population	-0.0008 (-0.89)	-0.0007 (-0.88)	-0.0005 (-0.54)	-0.0006 (-0.62)	0.0002 (0.14)	0.0001 (0.08)	0.0003 (0.42)	0.0001 (0.14)
Prop. Land Title	-0.3951 (-0.44)	- -	0.4597 (1.02)	- -	0.5894 (0.83)	- -	0.6563 (0.74)	- -
Fines	-0.0009 (-0.48)	-0.0009 (-0.46)	-0.0012 (-1.07)	-0.0012 (-0.97)	-0.0015 (-0.81)	-0.0015 (-0.81)	-0.0015 (-0.94)	-0.001 (-0.68)
Dist. to Ibama	-0.0005 (-0.24)	-0.0008 (-0.37)	-0.0009 (-0.74)	-0.001 (-0.85)	0.0013 (1.14)	0.0008 (0.75)	0.0005 (0.37)	0.0005 (0.41)
Constant	5.6863 (0.65)	4.9876 (0.57)	16.317*** (3.41)	15.783*** (3.28)	10.267** (2.08)	11.999** (2.48)	3.4441 (0.60)	4.432 (0.77)

Notes: Land-use IVQR regression results for the median. For each farm size, the column in the left presents the main specification, and the column in the right excludes “Proportion of Land Title” among the regressors.

Spatially lagged regressors are omitted from the table.

t-statistics in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Land-use Median Regression by Farm Size – Fines and Distance to Ibama

	(1)	(2)	(3)	(4)
Small				
Costs to Port	0.0038 (0.64)	0.0042 (0.70)	0.0042 (0.70)	0.0045 (0.76)
Fines	-	-0.0010 (-0.52)	-	-0.0009 (-0.48)
Dist. Ibama	-	-	-0.0004 (-0.17)	-0.0005 (-0.24)
Small–Medium				
Costs to Port	-0.0055* (-1.66)	-0.0055 (-1.62)	-0.0048 (-1.42)	-0.0053 (-1.58)
Fines	-	-0.0012 (-1.02)	-	-0.0012 (-1.07)
Dist. Ibama	-	-	-0.0008 (-0.64)	-0.0009 (-0.74)
Medium–Large				
Costs to Port	-0.0101*** (-2.71)	-0.0093** (-2.43)	-0.0094** (-2.50)	-0.0079** (-2.07)
Fines	-	-0.0012 (-0.70)	-	-0.0015 (-0.81)
Dist. Ibama	-	-	0.0011 (0.98)	0.0013 (1.14)
Large				
Costs to Port	-0.0126*** (-3.48)	-0.0116*** (-3.10)	-0.0109*** (-2.89)	-0.0110*** (-2.91)
Fines	-	-0.0016 (-1.04)	-	-0.0015 (-0.94)
Dist. Ibama	-	-	0.0008 (0.64)	0.0005 (0.37)

Notes: Land-use IVQR regression results for the median. The first column excludes “Fines” and “Distance to Ibama;” the second column excludes only “Distance to Ibama;” the third column excludes only “Fines;” and the last column presents the main specification. t-statistics in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Monitoring Efforts and Policy Implications

	Tax (US\$/ha)	LS Tax (US\$)	LS 80% Rule (US\$)	Emissions at US\$1/tCO ₂
Status Quo				
Specification				
Main	42.50	479 million	4.41 billion	-4.17 billion tC
No Dist. Ibama	42.50	482 million	5.25 billion	-4.19 billion tC
No Fines	40.00	453 million	3.24 billion	-4.20 billion tC
No Fines/No Dist. Ibama	37.50	425 million	3.75 billion	-4.21 billion tC
Fines at pre-2004 Level				
Specification				
Main	47.50	582 million	4.58 billion	-4.23 billion tC
No Dist. Ibama	47.50	585 million	5.41 billion	-4.24 billion tC

Notes: This table reports the counterfactual policy implications for different model specifications. The first column presents the value of the uniform tax (US\$/ha) that leads to 80% of the forested area. The second column presents the lost surplus of a uniform tax. The third column presents the lost surplus (LS) of the 80% rule. And the fourth column shows the avoided carbon emissions (in tons of carbon) associated with a US\$1/CO₂ carbon tax. The top panel refers to the “status quo” counterfactual; the bottom panel refers to the counterfactual in which I hold the number of fines after 2004 at pre-2004 levels. The different rows correspond to different estimated models: (i) the main specification; (ii) the model that excludes “Distance to Ibama” among regressors; (iii) the model that excludes “Fines” among regressors; and (iv) the model that excludes both “Distance to Ibama” and “Fines” among regressors. All monetary values are in dollars.

Table 6: Land-use Median Regression by Farm Size – Spatial Dependence

	(1)	(2)	(3)	(4)
Small				
Costs to Port	0.0040 (0.70)	0.0054 (0.93)	0.0045 (0.76)	0.0049 (0.84)
Spatial Lag Population	-	0.0003 (0.20)	-0.0022 (-1.25)	-0.0019 (-0.74)
Spatial Lag Mining	-	-0.0066 (-0.13)	0.0147 (0.27)	0.0420 (0.56)
Spatial Lag Power Plant	-	0.1972 (0.29)	-0.2428 (-0.30)	-1.2371 (-0.88)
Spatial Lag Dist. PA	-	0.0084 (1.56)	0.0101 (1.33)	0.0042 (0.66)
Small–Medium				
Costs to Port	-0.0051 (-1.60)	-0.0049 (-1.51)	-0.0053 (-1.58)	-0.0059* (-1.75)
Spatial Lag Population	-	0.0016** (2.22)	0.0015 (1.17)	0.0012 (0.75)
Spatial Lag Mining	-	-0.0142 (-0.47)	-0.0193 (-0.57)	-0.0214 (-0.59)
Spatial Lag Power Plant	-	-0.1016 (-0.26)	-0.0574 (-0.09)	-0.1236 (-0.18)
Spatial Lag Dist. PA	-	0.0048 (1.52)	0.0010 (0.30)	-0.0011 (-0.22)
Medium–Large				
Costs to Port	-0.0073** (-2.01)	-0.0062* (-1.67)	-0.0079** (-2.07)	-0.0056 (-1.52)
Spatial Lag Population	-	0.0011 (1.61)	0.0011 (0.85)	0.0017 (1.11)
Spatial Lag Mining	-	-0.0369 (-1.35)	-0.0536* (-1.67)	-0.0285 (-0.58)
Spatial Lag Power Plant	-	-0.1519 (-0.41)	-0.1882 (-0.32)	-0.6725 (-0.69)
Spatial Lag Dist. PA	-	0.0029 (1.16)	-0.0022 (-0.76)	-0.0020 (-0.68)
Large				
Costs to Port	-0.0100*** (-2.63)	-0.0098*** (-2.57)	-0.0110*** (-2.91)	-0.0115*** (-3.08)
Spatial Lag Population	-	0.0022 (0.55)	0.0002 (0.13)	0.0010 (0.46)
Spatial Lag Mining	-	-0.0660* (-1.67)	-0.0768** (-2.08)	-0.0848* (-1.84)
Spatial Lag Power Plant	-	-0.4322 (-0.87)	0.0995 (0.19)	0.0149 (0.02)
Spatial Lag Dist. PA	-	0.0015 (0.48)	-0.0034 (-1.02)	-0.0032 (-1.01)
Average # of Neighbors	-	3.5	7.3	12

Notes: IVQR regression results for the median. Column (1) excludes spatially lagged regressors; in column (2) the cutoff distance is 50 km; in column (3) the cutoff distance is 75km; in column (4) the cutoff is 100 km. The spatial weight matrices are of the power functional type with a distance decay parameter that equals 1.

t-statistics in parentheses, * p < 0.1, ** p < 0.05, *** p < 0.01.

Table 7: Comparison between the IVQR and the SPQIV at the Median

	Small-Medium		Medium-Large		Large	
	IVQR	SPQIV	IVQR	SPQIV	IVQR	SPQIV
Cost to Port	-1	-1	-1	-1	-1	-1
Altitude	0.0083 (0.08)	0.0083	-0.0233 (-0.26)	-0.0233	-0.1012 (-1.24)	-0.1012
Temp	-87.887** (-2.01)	-87.887	-33.687* (-1.75)	-33.687	-2.0766 (-0.11)	-2.0766
Rain	-1.7606* (-1.66)	-1.8266	-0.8658** (-2.07)	-0.8658	-0.4625 (-1.63)	-0.4855
Slope	-19.06 (-0.49)	-19.06	-26.398 (-0.91)	-26.398	-17.802 (-0.82)	-17.802
Soil 2	-341.71 (-1.38)	-341.71	-201.02* (-1.66)	-201.02	-153.86* (-1.81)	-153.86
Soil 3	-163.93 (-1.48)	-163.93	-117.28* (-1.87)	-117.28	-90.068** (-2.1)	-90.068
Soil 4	-279.69 (-1.60)	-283.19	-192.7** (-1.99)	-199.93	-146.39** (-2.52)	-146.39
Soil 5	-150.77 (-1.49)	-150.77	-74.165 (-1.60)	-74.165	-58.186 (-1.57)	-58.186
Dist. to PA	0.0867 (0.16)	0.0867	0.3558 (0.96)	0.3558	0.5017 (1.45)	0.5017
Mining	1.1417 (0.15)	1.1417	-5.4907 (-1.56)	-5.4907	0.6757 (0.27)	0.6757
Power Plant	-56.307 (-0.87)	-56.307	-13.811 (-0.33)	-13.811	-4.0286 (-0.14)	-4.0286
Power Plant Neighbor	-56.998 (-0.95)	-56.998	-39.287 (-1.08)	-39.287	-14.272 (-0.60)	-14.272
Population	-0.0884 (-0.55)	-0.088	0.0274 (0.14)	0.0274	0.025 (0.41)	0.025
Prop. Land Title	86.32 (0.84)	86.32	74.961 (0.76)	74.961	59.457 (0.72)	59.466
Fines	-0.2332 (-0.82)	-0.2332	-0.1907 (-0.70)	-0.1907	-0.1316 (-0.84)	-0.1316
Dist. to Ibama	-0.1617 (-0.7)	-0.1617	0.1613 (0.92)	0.1633	0.0429 (0.36)	0.0429

Notes: Land-use regression results for the median. For each farm size, the column in the left presents the logit model, where the constant is normalized to 0 and the coefficient to transport cost is normalized to -1; the column in the right presents the semiparametric model.

t-statistics in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Yields Regression (2SLS) by Farm Size – Crop-Pasture Index

	Small	Small–Medium	Medium–Large	Large
Cost to Port	0.0038*	-0.0009	0.0010	0.0018
	(2.49)	(-1.17)	(-0.95)	(1.26)
Observations	496	517	513	446
Average Local Yields (q_m^{cp})	0.70	0.42	0.41	0.27

t-statistics in parentheses, *p<0.05.

Table 9: Yields Regression (2SLS) for Crops by Farm Size – Coefficients for Costs to Ports

	Small	Small–Medium	Medium–Large	Large
Index – Only Crops (q_m^c)	0.0073*	-0.0021	-0.0052*	0.0098*
	(3.33)	(-1.61)	(-1.98)	(2.46)
Soy	-0.0001	-0.0007	-0.0002	-0.0022
	(-1.01)	(-0.79)	(-0.19)	(-1.23)
Corn	0.0033	-0.0012	-0.0019	-0.0012
	(1.22)	(-0.63)	(-1.09)	(-0.45)
Rice	-0.0001	-0.0007	-0.0019	-0.0027
	(-0.09)	(-0.60)	(-1.28)	(-1.31)
Beans	0.0025	0.00003	-0.0003	-0.0012
	(1.48)	(0.03)	(-0.28)	(-0.85)
Manioc	0.0074	-0.0029	-0.0064	0.0052
	(1.57)	(-0.77)	(-1.79)	(1.64)
Observations	496	517	513	446
Average Local Yields (q_m^c)	1.03	1.04	1.16	0.87

t-statistics in parentheses, *p<0.05.

Table 10: Regress Census Deforestation on Satellite Deforestation and Covariates

	Full Sample	Restricted Sample
Satellite Prop. Deforestation	0.400*** (15.39)	0.330*** (3.74)
Dist. to Port	-0.00000694 (-0.55)	-0.0000595 (-0.71)
Dist. to Capital	0.0000664 (1.89)	-0.0000432 (-0.23)
Altitude	0.0000316 (0.57)	-0.0000456 (-0.30)
Temp	-0.0191 (-1.29)	-0.0339 (-0.69)
Rain	-0.00113*** (-4.03)	-0.00119 (-1.36)
Slope	-0.0337 (-1.91)	0.000476 (0.02)
Soil 2	-0.227*** (-5.22)	1.351 (0.67)
Soil 3	-0.188*** (-3.85)	-0.0343 (-0.44)
Soil 4	-0.250*** (-7.02)	-0.0606 (-0.78)
Soil 5	-0.231*** (-6.66)	0.0125 (0.16)
Dist. to PA	0.000994*** (4.82)	0.000169 (0.36)
Mining	-0.000816 (-0.44)	-0.00847 (-0.32)
Power Plant	-0.000879 (-0.05)	-0.00779 (-0.07)
Power Plant Neighbor	0.00828 (0.56)	0.0318 (0.99)
Population	-0.0000426 (-0.71)	0.000634 (0.57)
Prop. Land Title	0.101* (2.44)	0.0651 (0.28)
Fines	-0.000101 (-1.02)	-0.000305 (-0.92)
Dist. to Ibama	0.000118 (1.07)	0.000166 (0.29)
Prop. Clouds	-0.0381 (-1.19)	-29.61 (-1.21)
Constant	0.909* (2.07)	1.537 (1.08)
Observations	523	45
Adjusted-R ²	0.67	0.64

Notes: The second column presents the estimated coefficients for the full sample; and the third column, for the restricted sample. Both use the Eicker-Huber-White robust standard errors. The restricted sample includes municipalities in which the share of private land is greater than 80%, and the share of clouds/unobserved areas is less than 3%. t-statistics in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Unit Costs to Travel

Unit Cost	R\$/ton.km	US\$/ton.km
Paved Road	0.07678	0.0353
Unpaved Road – Outside Rainforest	0.0992	0.0457
Unpaved Road – Within Rainforest	0.15	0.069
Navigable River – Good Infrastructure	0.0444	0.0204
Navigable River – Poor Infrastructure	0.1139	0.0525
Railroad	0.0608	0.028
Land – Outside Rainforest	1.5	0.6912
Land – Inside Rainforest	3	1.3824