WORKING PAPER · NO. 2023-86

Carbon Prices and Forest Preservation Over Space and Time in the Brazilian Amazon

Juliano J. Assunção, Lars Peter Hansen, Todd Munson, and José A. Scheinkman

JUNE 2023

Macro Finance Research Program of the Becker Friedman Institute

EPIC
ENERGY POLICY INSTITUTE
AT THE UNIVERSITY OF CHICAGO

5757 S. University Ave.
Chicago, IL 60637
Main: 773.702.5599
bfi.uchicago.edu

Carbon prices and forest preservation over space and time in the Brazilian Amazon*†

Juliano Assunção (Climate Policy Initiative and PUC-Rio)
Lars Peter Hansen (University of Chicago)
Todd Munson (Argonne National Laboratories)
José A. Scheinkman (Columbia University)

April 10, 2023

Abstract

Some portions of land in Brazilian Amazon are forested, and other portions used in agriculture. Deforestation (reforestation) emits (captures) carbon, which has consequence for the global climate. The social and private productivities for these alternative land uses vary across locations within the Amazon region. In this research, we build and analyze a spatial/dynamic model of socially efficient land allocation to establish a benchmark for *ad-hoc* policies. We show how to incorporate the stochastic evolution of agricultural prices, and we explore the consequences of ambiguity in the location-specific productivities on the socially efficient policy. Finally, we assess the consequences of imposing alternative social costs of carbon emissions on the spatial/dynamic allocation of land use Our results indicate that with modest transfers per ton of net CO_2, Brazil would find it optimal to choose policies that produce substantial capture of greenhouse gasses in the next 30 years, suggesting that the management of tropical forests could play an important role on climate change mitigation in the near future.

*This draft is preliminary and incomplete.

†We thank João Pedro Vieira and Daniel (Samuel) Zhao for their expert research assistance and to Diana Petrova and Carmen Quinn for editorial assistance.

1 Introduction

The Amazon forest contains 123 ± 31 billion tons of captured carbon that can be released into the atmosphere, equivalent to the historical cumulative emissions of the United States (Malhi et al. (2006), Friedlingstein et al. (2022)). The Brazilian Amazon occupies 60% of the 2.7 million square miles that comprise the Amazon. From 1985 to 2021, the agricultural area in the Brazilian Amazon increased from 68.6 to 240.5 thousand square miles. The associated deforestation, comprising an area the size of Texas, has resulted in high emissions, setting the Brazilian Amazon as a substantial outlier in a plot of countries' emissions per-capita *vs.* GDP per-capita. (see Figure 1.)

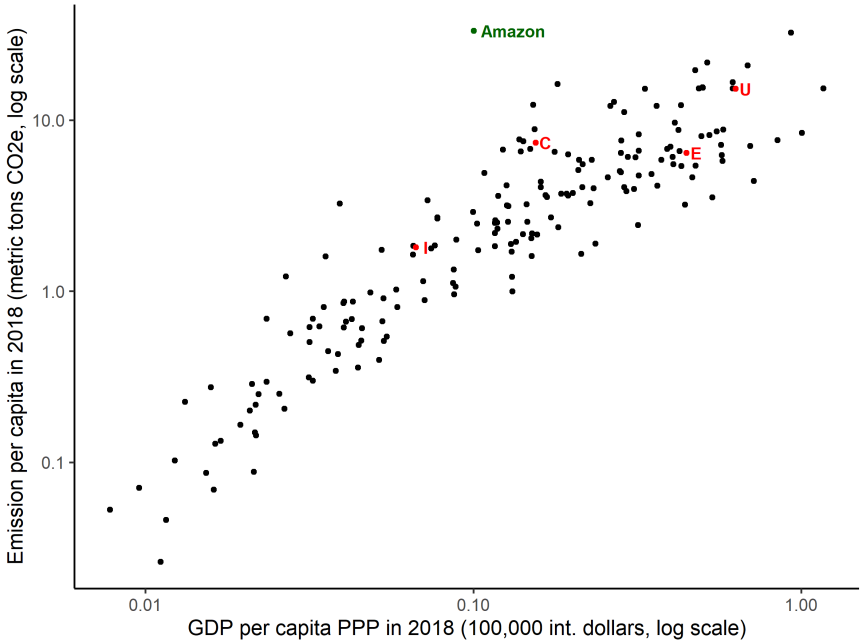

Figure 1: Each dot represents a country in 2018, except for the European Union and the Brazilian Amazon. Highlighted letters stand for (C)hina, (I)ndia, (E)uropean Union, and (U)nited States. Sources: World Bank Data, downloaded on March 2021; Fatos da Amazônia 2021 (www.amazonia2030.org).

This paper investigates the potential social gains to designing prudent policies that combat deforestation through the lens of a spatial and dynamic model. We build the model to capture the trade-off between agricultural production and forest preservation or regeneration.

We use the model to provide insights into structuring policy improvements that realign economic incentives for allocating land use. We make the model quantitative through the use of detailed spatial information from multiple data sets. Our data document large cross-sectional variability in cattle farming productivity and in the potential absorption of carbon in the Brazilian Amazon. To account for this variability, we divide the Amazon region into various subregions or sites making the spatial dimension important in our analysis. While the model has considerable cross-sectional richness, it is nevertheless highly stylized for reasons of tractability and transparency.

We pose the model in continuous time. The cross-sectional heterogeneity in productivities and the natural state constraints on the land allocation preclude standard recursive methods for solving so-called Hamilton-Jacobi-Bellman (HJB) equations. Instead, we use and extend methods from Modified Predictive Control (MPC) that were originally developed in control theory and engineering to study multi-plant production in real time. The inequality constraints on the states are approximated using an interior point method. These MPC alogorithms allow for uncertainty specified as a Markov process and are implemented by incorporting a shorter uncertainty horizon than the overall control horizon as a means of approximation. We also explore parameter ambiguity from the standpoint of the social planner by making a model-determined robustness adjustment to the subjective probabilities for the unknown parameters. We are once again pushed to use numerical methods, in this case a Metropolis Hasting algorithm implemented as a Markov Chain Monte Carlo (MCMC) simulator. In this way, we confront what is sometimes referred as "deep uncertainty."

Our analysis of the model implications proceeds in four steps. First, we use the model to elicit an estimate of the shadow price for emissions revealed by the deforestation during 1995-2008. The year 1995 is the first date at which we have reliable price data on cattle prices.[1] The year 2008 marks the beginning of the Amazon Fund, financed primarily by the Norwegian and German governments. Their funding was a pay-for-performance scheme based on an emissions price of $5 per ton of CO2e. It generated to USD 1.2 billion in payments in 2008-2017 (Angelsen (2017), Correa et al. (2019)) and provides and example of how deforestation may be influenced by outside payments. We use this shadow price to produce simulations when we assume that "business as usual" prevails.

We then study the impact of adding outside payments for *net* capture of CO_2 in the Amazon. We produce results at three levels of aggregation in steps two, three and four. In step two we construct the finest grid by considering 1059 sites in the Amazon biome with each measuring 67.5km$\times$ 67.5km. For this level of detail, we produce results without accounting for uncertainty in the price of agricultural output, by assuming (a) the price that corresponds to the highest quartile of our observed price series; (b) the price that corresponds to the lowest quartile of our price series; and (c) the price corresponding to the (stationary) average of the two-state Markov process we fit to the observed agricultural prices. We use these deterministic solutions to produce bounds on the value functions for the stochastic case and ascertain whether and by how much Brazil would gain if the country signed an agreement for a set of hypothetical dollar transfers per net unit of CO_2. Our model output shows both the social gains to overall preserving or enhancing the Brazilian rain forest and to reallocating production in the cross section.

In a third step, we consider 81 sites, each of which is 270km$\times$270km. For this level of aggregation, we produce results that take into account stochastic changes in the price of agricultural output. The results are qualitatively similar to the case of 1059 sites.

Estimates of some crucial parameters in the model are subject to non-trivial uncertainty. In the

[1] Until mid-94 Brazil went through a period of very high and volatile inflation.

third step, we explicitly consider parameter uncertainty from the perspective of the social planner. We start with a prior distribution over the full array of site-specific productivity parameters for carbon sequestration and agricultural productivity. Since the choice of prior distribution is *ad hoc* and uncertain, we engage in a prior sensitivity analysis subject to penalization. This sensitivity analysis uses the prior as a baseline distribution and the planner's objective function to ascertain the which departures from the baseline are of most concern to the planner. The magnitude of the penalty parameter limits how much sensitivity is entertained and serves as the inverse of an ambiguity aversion parameter. For this investigation to be tractable, we use a very coarse grid to construct a small number of heterogeneous sites.

While we assume that the planner can directly control deforestation, we also view the solution to the planner's problem as providing a benchmark for comparing the outcomes of alternative *ad hoc* policies, and suggesting improvements over current policies.[2]

The rest of the paper is organized as follows: In the next section, Section 2 we review some of the relevant literature. This is followed in Section 3 with an exposition of our theoretical model. Section 4 discusses the numerical methods used to compute solutions to the social planner's maximization problem. Section 5 describe how we confront ambiguity about the subjective distributions for unknown parameters and how this approach is related the smooth ambiguity model from decision theory. Section 6 explains how we use a large collection of relevant data sets to calibrate the model. Our results are presented in Section 7, which is followed by our conclusions and suggestions for further work.

2 Related Substantive Literature

Griscom et al. (2017) identify and quantify "natural climate solutions" (NCS), which include tropical forests. Heinrich et al. (2021) focuses on the potential of the Brazilian Amazon. As noted by Balboni et al. (2022), most studies on agricultural expansion and deforestation are static.

A recent branch of the literature uses discrete-choice models in order to study the link between agriculture and deforestation (Souza-Rodrigues (2019), Dominguez-Iino (2021), Araujo et al. (2022)). Souza-Rodrigues (2019) and Dominguez-Iino (2021) develop static approaches, emphasizing the role of the transportation network and trade to the design of policies. Araujo et al. (2022), on the other hand, presents a dynamic model along the lines of Scott (2014), allowing farmers to internalize the social value of carbon. However, the dynamics in Araujo et al. (2022) is restricted to the forward-looking behavior of farmers.

In contrast to the existing literature, our dynamic approach not only accounts for how expected future prices of agricultural goods influence optimal current land use, it also incorporates carbon

[2]Brazil has experience implementing effective policies to curb deforestation. The launch of satellite-based monitoring systems, the creation of protected areas, the enactment of conditioned-credit measures and the creation of a priority list of municipalities comprised a new approach to combat deforestation, resulted on a reduction of more than 80% in the deforestation rates between 2004 and 2012 (Gandour (2018), Assunção and Rocha (2019), Assunção et al. (2020), Assunção et al. (2022a) and Assunção et al. (2022b)).

emissions from deforestation and carbon capture from forest regeneration. In addition, we also take into account the uncertainty on the forest carbon measures. Consequently, we provide a framework that integrates the impact of carbon prices on deforestation, forest restoration, and agriculture.

Finally, our results contribute to the literature on climate policy design. Our simulation shows that deforestation in the Amazon will cross the tipping point of $20-25\%$ suggested by Lovejoy and Nobre (2018) in the scenario with carbon price equal to the shadow emission price elicited from the 1995-2008 period with no additional international payments.[3] On the other hand, additional payments of at least \$15/ton would not only safeguard the tipping point, but would also trigger forest restoration on large scale. In this sense, the carbon sink potential of secondary forests emphasized by Griscom et al. (2017) and Heinrich et al. (2021) can be realized with these additional carbon payments.

3 Model

We pose the problem of a fictitious social planner who considers the tradeoff between using land for agriculture and nurturing or preserving forests that function as carbon sinks. This planner internalizes the externalities resulting from deforestation. The planner's problem is dynamic with explicit heterogeneity across regions in the Amazon. Guided by empirical measurements, the regions have two important sources of heterogeneity: i) agricultural productivity and ii) ability to absorb atmospheric carbon.

Let i denote a site index for $i = 1, 2, ..., I$ where I is the total number of sites and $t \in [0, T]$ the point in time. We use superscripts to denote sites and subscripts to denote dates. We adopt the notational convention that uppercase letters depict the actual state and lower case letters the potential state realizations. At date t,

Z_t	$\stackrel{\text{def}}{=}$	$(Z_t^1, Z_t^2, ..., Z_t^I)$	vector of area used for agriculture expressed in hectares
X_t	$\stackrel{\text{def}}{=}$	$(X_t^1, X_t^2, ..., X_t^I)$	vector of carbon captured expressed in Mg CO2e (CO2 equivalent)
A_t	$\stackrel{\text{def}}{=}$	$(A_t^1, A_t^2, ..., A_t^I)$	vector of agricultural output expressed in constant USD

We use the notation $Z \stackrel{\text{def}}{=} \{Z_t : 0 \leqslant t \leqslant T\}$ to denote the corresponding process that evolves over time, and similarly for other states and controls. In our base model the single aggregate state variable is P_t^a, an index of cattle prices in Brazil expressed in 2017 US dollars.[4]

The state vector Z_t is subject to an instant-by-instant and coordinate-by-coordinate constraint:

$$0 \leqslant Z_t^i \leqslant \bar{z}_i$$

where $\bar{z}_i$ is the amount of land available or agriculture at site i.[5] Let $\dot{Z}_t$ be the time derivative of

[3] Notice that our shadow emission price is close to the \$7.26/ton estimated by Araujo et al. (2022).

[4] We choose cattle prices because, in recent years, more than 85% of deforested land is dedicated to cattle grazing - soybean, the largest crop in the region, accounts for about 8% of the farming land (Mapbiomas - www.mapbiomas.org).

[5] For calibration of this and the other parameters see Section 6

Z at date t.

The evolution of X^i introduces an important asymmetry into our problem. We may write a "linear" version of this problem by introducing two site specific, scalar, non-negative control variables for our fictitious planner, U_t^i and V_t^i, that distinguish positive from negative movements in the derivative of Z_t^i :

$$\dot{Z}_t^i = U_t^i - V_t^i. \tag{1}$$

The site specific state variable process X^i evolves as:

$$\dot{X}_t^i = -\gamma^i U_t^i - \alpha \left[X_t^i - \gamma^i \left(\bar{z}^i - Z_t^i \right) \right] \tag{2}$$

where the parameters satisfy: $\gamma^i \geqslant 0, \alpha > 0$, for $i = 1, 2, ..., I$. The first term in the right side of (2) connects deforestation to a loss in captured carbon. The site-specific parameter $\gamma^i \geqslant 0$ denotes the density of CO2e that is present in a primary forest in site i.[6] The next term expresses the growth in captured CO2e, when the size of the forest in site i is held constant. The mean-reversion coefficient α guarantees that if one lets the forest grow undisturbed in a deforested area, it would reach $[1 - \exp(-\alpha 100)]\%$ of the maximal captured CO2e in 100 years as in Heinrich et al. (2021). In our case, we choose α such that $[1 - \exp(-\alpha 100)] = 99\%$. Notice that, holding constant the deforested area, the amount of carbon in a site converges to γ^i per hectare of remaining forest. In Remark 3.3, we argue that at the optimum $U_t^i V_t^i = 0$. Thus one of the controls is always zero at the boundary, which introduces additional binding constraints into the analysis.

We model cattle output as proportional to the land allocated to cattle farming,

$$A_t^i = \theta^i Z_t^i \tag{3}$$

where θ^i is a site-specific productivity parameter.

All of the locations contribute to emissions via the capture of carbon and emissions that result because of agricultural activity with a net impact given by

$$\kappa \sum_{i=1}^{I} Z_t^i - \sum_{i=1}^{I} \dot{X}_t^i, \tag{4}$$

where the parameter κ captures the emissions that result because of cattle farming. We include a cost of adjustment to changes in the use of land with contributions from each site. It is measured by

$$\frac{\zeta}{2} \left[\sum_{i}^{i} \left(U_t^i + V_t^i \right) \right]^2.$$

The price process P^a for the agricultural output evolves exogenously as an n-state Markov chain in continuous time with time invariant transitions. This process has an infinitesimal generator

[6] For simplicity, equation (2) assumes that all deforestation occurs in primary forest, what is not far from what has been observed in the Amazon.

represented as an intensity matrix $\mathbb{M}$ with nonnegative entries off-diagonal entries $\mathsf{m}_{\ell\ell'} \geqslant 0$ for $\ell' \neq \ell$ and diagonal entries

$$\mathsf{m}_{\ell\ell} = -\sum_{\ell'=1,\ell'\neq\ell}^{n} \mathsf{m}_{\ell\ell'}.$$

The implied transition probability matrix over an interval of time τ is $\exp(\tau\mathbb{M})$ computed using a a matrix counterpart to a power series.

Since many carbon trading schemes are based on emissions, we assume that the planner takes as given a price for carbon emissions P^e, the initial price for agriculture and the Markov process that describes the future evolution of the price P_t^a for cattle and maximizes

$$\mathbb{E}\left\{\int_0^\infty \exp(-\delta t)\left[-P^e\left(\kappa\sum_{i=1}^{I} Z_t^i - \sum_{i=1}^{I}\dot{X}_t^i\right) + P_t^a\sum_i \theta^i Z_t^i - \frac{\zeta}{2}\left(\sum_i U_t^i + V_t^i\right)^2\right]dt\right\} \tag{5}$$

subject to equations (1)-(2), and the control restrictions:

$$U_t^i \geqslant 0, \qquad V_t^i \geqslant 0 \qquad t \geqslant 0.$$

where δ is the subjective discount rate. The emission price P^e considered by the planner would be the sum of the (constant) marginal value attributed by the planner to emission and any monetary transfers obtained from others, such as sales in carbon emission markets. In future work we intend to make P^e a state variable.

Remark 3.1. *The objective function* (5) *values agricultural output by the value of sales, thus assuming that inputs to production have no alternative use. This choice is dictated by lack of data on the cost of attracting or redeploying agricultural inputs, but it biases the results in favor of agricultural use.*

Remark 3.2. *The only interaction across sites in objective function* (5) *occurs through the adjustment costs. These interactions are intended to be the result of a less than perfectly elastic supply of resources needed for changing land use at the level of the whole Amazon. We could have introduced instead interactions across sites via nonlinearities in the valuation of agricultural output and/or emissions.*

Remark 3.3. *To argue that the controls U_t^i and V_t^i satisfy the complementary slackness condition $U_t^i V_t^i = 0$ for each pair (i,t) it is easier to consider a discrete time model. The proof for the analogous result for the continuous time case goes through by taking limits. Suppose you take take a point where the optimal trajectory involves $\min\{U_t^i, V_t^i\} > \Delta > 0$. If the planner lowers both controls by Δ, then at time t, one obtains an increase of Δ in X_t^i and lower emissions $\gamma^i\Delta$. Equation* (2) *implies that X_t^i would have a lower drift and converge over time to the stationary solution. This in turn implies that the sum of future emissions would increase by $\gamma^i\Delta$. However since the rate of discount is positive, the value of the problem would increase. Thus an optimal solution cannot involve simultaneously positive values for U_t^i and V_t^i.*

Remark 3.4. *Optimization problem 5 does not involve the stocks of (extended) carbon in the atmosphere generated by activities in the Amazon biome. However given the optimal solution one could derive the implied impact on the evolution of carbon stocks, by appending a model that maps emissions into carbon in the atmosphere.*

4 Solving the maximization problem

To achieve the needed degree of economic and environmental richness, we use numerical methods to obtain model solutions. To confront this locational heterogeneity, we necessarily have a large number state variables with state-constraints that bind at the optimal solution. To confront the inequality restrictions that are central to our problem, we use a so-called interior point method. This method is implemented by imposing penalties on logarithms on variables constrained to be nonnegative. While the interior point approximation pushes solutions away from their zero boundaries, in practice the solutions will be close enough to zero to identify the binding constraints.

In the absence of price uncertainty, we are able to solve our optimization problem by simply solving a static problem over all possible trajectories over the next 200 years, using interior point methods. Given obvious uniform bounds on possible utility flows and discount rates that exceed 2% we can argue that the resulting trajectories achieve close to the infinite horizon optima.

When considering price uncertainty, we find "Modified Predictive Control" (MPC) methods (*e.g.* Scokaert and Rawlings (1998), Bemporad et al. (2002), Thangavel et al. (2018)) to be particularly suitable for solving our planner's problem. Our MPC approximation is implemented as follows.[7] Given the current period, say date zero, looking forward, we break the future into two segments: a) an uncertainty horizon of say τ time periods and b) the remaining $T - \tau$ time periods beyond this uncertainty horizon for which we abstract from uncertainty. While the cattle price distribution follows a Markov chain, to simplify our computations, we impose that the prices in periods $\tau + 1, \ldots T$ are set equal to the value that prevails at τ. We then apply the interior method to find the optimal trajectory at zero given P_0^a. We keep the optimal states for time $t = 1$ and repeat - that is we consider the problem stating at $t = 1$ with the new state vector and divide the future on two segments: an uncertainty segment of length τ and a remaining period of $T - \tau$. This step will produce an optimal state at period 2, we then repeat the procedure to produce the optimal state at period 3 and continue going forward. Thus we confront randomness in this problem by imposing appropriate "measurability" restrictions on the controls.

In practice, the dimensionality of the stochastic problem increases geometrically as a function of the uncertainty horizon, τ. This MPC method becomes tractable when the uncertainty horizon can be relatively short and still obtain good approximations. We determine an "adequate" uncertainty horizon τ^* by checking the difference in the value of the problem $V(\tau) - V(\tau - 1)$ for $\tau = 0, 1, \ldots, \tau^*$. In our example with 81 sites and price uncertainty, we chose $\tau^* = 6$.

Many of the results we will show entail projections into the future. We report results based a

[7] Related computational approaches have been proposed by Cai et al. (2017) and Cai and Judd (2023).

common randomly drawn sequence of cattle prices, P_t^a $t = 1, \ldots T$, using the observed P_0^a and the calibrated Markov chain.

5 Parameter uncertainty

We investigate a static formulation of robustness to parameter uncertainty. For each site, we consider the parameter pair $\beta^i = (\gamma^i, \theta^i)$ for $i = 1, 2, ..., I$. Let $\boldsymbol{\beta}$ denote full parameter vector including all sites and hence of dimension $2 \times I$.

Our planner is uncertain about these parameter values and instead has baseline probability distribution π. In addition, this planner is uncertain about what distribution to use and instead thinks of π as a rough approximation. We address this uncertainty by introducing ambiguity about the parameter distribution.

Let $\boldsymbol{d}$ be the vector of decisions and $f(\boldsymbol{d}, \boldsymbol{\beta})$ for the resulting value given the unknown parameter $\boldsymbol{\beta}$. We use a divergence measure to capture ambiguity about the parameter distribution. For $\int g(\boldsymbol{\beta}) d\pi(\boldsymbol{\beta}) = 1$, the relative entropy (or Kullback-Leibler) divergence

$$\int_{\mathcal{B}} \log g(\boldsymbol{\beta}) g(\boldsymbol{\beta}) \pi(\boldsymbol{\beta}) \geqslant 0,$$

is a commonly used measure of divergence between a probability $g(\boldsymbol{\beta})d\pi(\boldsymbol{\beta})$ and the baseline $d\pi(\boldsymbol{\beta})$. To produce optimal controls that are robust to the parameter uncertainty, solve

Problem 5.1.

$$\max_{\boldsymbol{d}} \min_{g, \int g d\pi = 1} \int_{\mathcal{B}} f(\boldsymbol{d}, \boldsymbol{\beta}) g(\boldsymbol{\beta}) d\pi(\boldsymbol{\beta}) + \xi \int \log g(\boldsymbol{\beta}) g(\boldsymbol{\beta}) d\pi(\boldsymbol{\beta})$$

where $\xi > 0$ is penalty parameter.[8]

We implement a full commitment to the baseline distribution by making ξ arbitrarily large. More modest settings capture a concern for robustness.

One nice feature of using relative entropy divergence to explore distributional sensitivity is that the minimization problem has a quasi-analytical solution:

$$-\xi \log \int_{\mathcal{B}} \exp\left[-\frac{1}{\xi} f(\boldsymbol{d}, \boldsymbol{\beta})\right] d\pi(\boldsymbol{\beta}) = \min_{g, \int g d\pi = 1} \int_{\mathcal{B}} f(\boldsymbol{d}, \boldsymbol{\beta}) g(\boldsymbol{\beta}) d\pi(\boldsymbol{\beta}) + \xi \int \log g(\boldsymbol{\beta}) g(\boldsymbol{\beta}) d\pi(\boldsymbol{\beta}) \tag{6}$$

where the minimizing g given by:

$$g^* = \frac{\exp\left[-\frac{1}{\xi} f(\boldsymbol{d}, \boldsymbol{\beta})\right]}{\int_{\mathcal{B}} \exp\left[-\frac{1}{\xi} f(\boldsymbol{d}, \tilde{\boldsymbol{\beta}})\right] d\pi(\tilde{\boldsymbol{\beta}})} \tag{7}$$

tilts the distribution towards smaller objectives for each given decision $\boldsymbol{d}$. The candidate solution

[8] Alternatively, we can think of ξ a as Lagrange multiplier on a relative entropy divergence constraint.

presumes that:

$$\int_{\mathcal{B}} \exp\left[-\frac{1}{\xi} f(\boldsymbol{d}, \boldsymbol{\beta})\right] d\pi(\boldsymbol{\beta}) < \infty$$

which implicitly limits how fat the tail can be for the distribution π.

Remark 5.2. *The formula on the left-side of* (6) *is a special case of a smooth ambiguity objective, first suggested by Klibanoff et al. (2005). They deduced a rationale for an ambiguity adjustment represented using a concave function distinct from the one used for expressing risk aversion. Thus the negative exponential in* (6) *is such a concave function. The logarithmic adjustment converts this to a certainty equivalent. Their axiomatic motivation is quite different from the distributional robustness that is of interest to us, however.*

For conceptual reasons, we also switch the order of the maximization and minimization. Under quite general conditions, we can invoke the min-max theorem:

Problem 5.3.

$$\min_{g, \int g d\pi = 1} \max_{\boldsymbol{d}} \int_{\mathcal{B}} f(\boldsymbol{d}, \boldsymbol{\beta}) g(\boldsymbol{\beta}) d\pi(\boldsymbol{\beta}) + \xi \int \log g(\boldsymbol{\beta}) g(\boldsymbol{\beta}) d\pi(\boldsymbol{\beta})$$

where $\xi > 0$ *is penalty parameter.*

Consider the inner maximization problem:

$$\max_{\boldsymbol{d}} \int_{\mathcal{B}} f(\boldsymbol{d}, \boldsymbol{\beta}) g(\boldsymbol{\beta}) d\pi(\boldsymbol{\beta})$$

where we are free to drop the relative entropy penalty as it does not depend on the decision $\boldsymbol{d}$. Provided that this problem has a solution for the outer g minimization, then the planner is maximizing against this particular (penalized) "worst case probability." This computation is of interest as a way to interpret the consequences of any given choice of the penalty parameter ξ. In practice, we find it revealing to explore alternative choices of ξ and deduce their implications for the implied worst case probabilities.[9]

For solving the robust problem numerically, we take an iterative approach, also supported by the min-max theorem. Specifically, we proceed as follows:

i) Given a g, we solve the maximization problem for a candidate $\boldsymbol{d}$. We ignore the relative entropy penalty term in this solution.

ii) For a given $\boldsymbol{d}$, we solve the minimization problem with the relative entropy term to obtain a new candidate for g.

iii) We repeat the steps until we achieve convergence.

For step ii), we use a Metropolis-Hastings simulation approach along with the quasi-analytical solution in computing (7). A numerical method is necessary because of the denominator term, and

[9]This follows a common practice for robust Bayesian methods.

Markov simulation gives us one way to explore a parameter space of potentially large dimension. Very similar to how a Metropolis-Hastings algorithm can be used to compute a Bayesian posterior in terms of a prior and a likelihood, we calculate the exponentially tilted solution using $d\pi(\boldsymbol{\beta})$ and $\exp\left[-\frac{1}{\xi}f(\boldsymbol{d},\boldsymbol{\beta})\right]$.

6 Calibration

In order to calibrate our model for the Brazilian Amazon, we combine data from multiple sources. Initially, it is important to notice two aspects of this exercise. First, we use the year of 2017 as a reference for many variables, since this is the year of the latest Agricultural Census in Brazil. Second, for the site-specific parameters and variables, indexed by i, we build regular grids based on our most granular information, which comes from the MapBiomas platform.[10] We aggregate the high-resolution pixels ($\tilde{3}$0m) to a regular grid of 1887 sites, each with 2,250 30m-pixels,m of which 1059 overlap with the 421M hectares of the Amazon biome. This is our baseline sample for the site-specific data. To allow us to treat the case of agricultural price uncertainty we will aggregate 16 sites of finer grid to produce sites of $\approx$ 268km $\times$ 268km. We obtain 81 sites after dropping a site with less than .2% in the Amazon biome. Appendix A describes in detail all of the data used for the calibration of parameters and initial conditions.

Figure 2 shows the initial land allocated to agriculture and the initial stock of absorbed carbon across the sites we construct. Figure 3 shows how the carbon sequestration parameter γ_i varies across the different sites, and Figure 4 does the same for the agricultural productivity parameter, θ_i. While there is negative correlation between θ_i and γ_i across sites, they are not perfectly aligned.

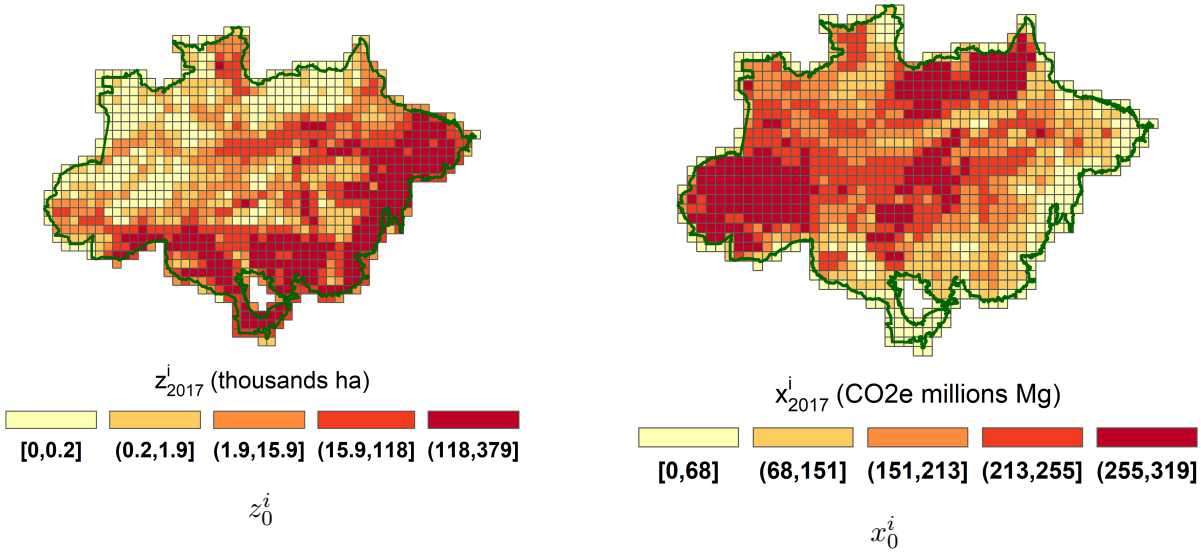

Figure 2: Initial values for agricultural area (z_0^i) and carbon stock (x_0^i)

[10] www.mapbiomas.org (Collection 5).

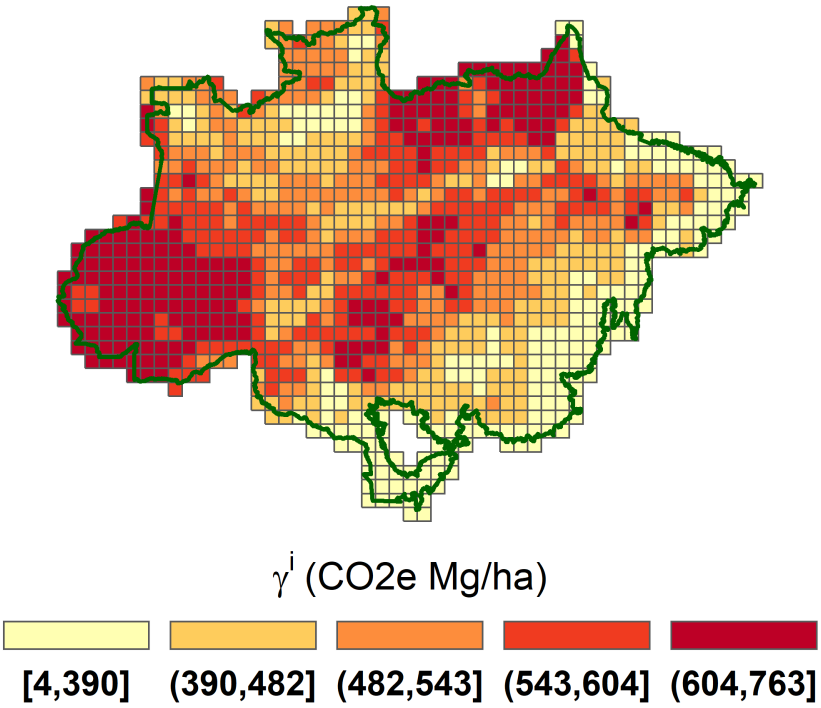

Figure 3: Carbon sequestration parameter heterogeneity.

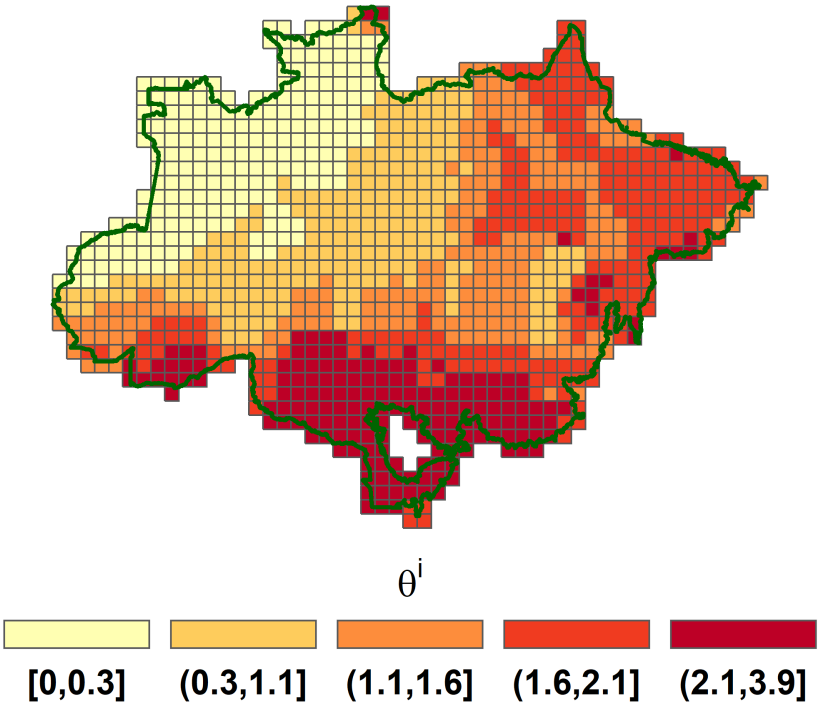

Figure 4: Agricultural productivity heterogeneity

7 Results

7.1 Shadow prices under business as usual

We infer a *shadow* value for the planner based on historical experience. To obtain this value, we first choose a period $[\underline{t}, \bar{t}]$ and then choose a price for emissions, that we denote by P^{ee} to match the aggregate deforestation predicted by the model to the actual aggregate deforestation from the initial period $\underline{t}$ to a final observation period $\bar{t}$. More precisely, we find the optimal trajectory from $\underline{t}$ to $\bar{t}$ given the observed initials condition $(X^o_{\underline{t}}, Z^o_{\underline{t}})$, and the realized history of agricultural prices $\{P^a_t\}^{\bar{t}}_{t=\underline{t}}$, and we choose a price of emissions $P^e(\underline{t}, \bar{t})$ so that $Z_{\bar{t}} = Z^o_{\bar{t}}$.[11] We use $\underline{t} = 1995$, the initial date for our price data and $\bar{t} = 2008$ the announcement of the Amazon fund that would pay for preservation projects in the Amazon, using funds contributed mostly by Norway and set $P^{ee} \stackrel{\text{def}}{=} P^e(1995, 2008)$.

To match the observed deforestation in 1985-2008, we obtain $P^{ee} = \$5.76$, for $\delta = .02$, while for $\delta = .03$, we obtain $P^{ee} = \$3.89$. The lower price associated to a higher δ is natural because the cost of deforestation is instantaneous while benefits from cattle farming have long duration. Thus a higher discount rate would lead to less cattle and to match a fixed deforestation requires a lower price for emissions.

7.2 Results for 1059 sites

In this section we discuss results for 1059 sites but with a constant price for cattle. We first examine the case when $P^a_t = P^{a,h}$ the 75th percentile of the series thus biasing our results away from conservation. We show results for simulations with $\rho = .02$, for a period of 50 years, and $P^e = P^{ee} + \$b$ with $b = 10, 15$ and 20. Recall that for $\rho = .02$, $P^{ee} = \$5.76$.[12]

As Figure 5 shows, with "business as usual" ($P^e = P^{ee}$), the optimal choice involves an increase in the agricultural area from 15 to 30% of the Biome, a fraction that may result on tipping of the eastern, southern and central Amazon to a non-forest system (Lovejoy and Nobre (2018)). The trajectories predicted are much different starting with an additional payment to the planner of \$10, \$15 or \$20 per unit of *net* captured emissions, leading to substantial decreases in agricultural area and corresponding increases in carbon captured. However as shown in Table 1 the discounted value to the planner of a commitment to receive (pay) \$10 of net transfers for each ton of CO_2 captured (emitted) of CO_2 *lowers* the value to the planner in the case of high agricultural by 16%, relative to the case of business as usual, reflecting the costs the planner would incur when producing net emissions. In the case when $P^a = P^{a,\ell}$, \$10 transfers increase the value by 1.5%. Since the stationary distribution puts higher probability on $P^{a,h}$, Table 1 indicates that transfers of \$10 would make Brazil worse off when agricultural prices fluctuate. In contrast the value to

[11] We obtain a similar value if instead we minimize the norm of the vector $(\frac{X^o_{\bar{t}} - X^o_{\underline{t}}}{X^o_{\underline{t}}}, \frac{Z^o_{\bar{t}} - Z^o_{\underline{t}}}{Z^o_{\underline{t}}})$.

[12] In fact, the results for $\rho = .03$ are very similar. The same duration differential between deforestation and cattle farming diminishes deforestation for any given P^e for $\rho = .03$, relative to the case $\rho = .02$ Since $P^{ee}(.03) < P^{ee}(.02)$ the trajectories end up being quite similar for $\rho = .03$ and $\rho = .02$.

the planner with $P^a = P^{a,h}$ and $P^e = P^{ee} + 15$ ($P^e = P^{ee} + 20$) is higher by \$8 billion (\$82 billion) than with no transfers. Table 1 illustrates the convexity of the value to the planner as a function of emissions pricing. The convexity of the value function is a mathematical property but the non-monotonicity in the interval $[P^{ee}, \infty)$ is particular to the calibration we obtained.

Table 1 also shows that $3.34 \times 10^{11} USD = V(P^{ee}, P^{a,h}) \sim V(P^{ee} + 15, P^{a,\ell}) = 3.34 \times 10^{11}$ USD. Given the strict monotonicity of V with respect to the price of agriculture, this indicates that when prices follow the Markov chain we estimated, Brazil is better off by accepting a contract that calls for \$15 per net ton of CO_2 captured, then going alone. We may also calculate a lower bound for the gain obtained from moving from business as usual to a contract with transfers of \$b dollars per ton, under the assumption that the distribution of agricultural prices is the stationary distribution $\tilde{P}^a$ associated with the estimated Markov chain. If we write $\bar{P}^a$, for the stationary expected value, than we know that for each initial conditions (X_0, Z_0) the linearity of the utility function on P^A guarantees that

$$V(\tilde{P}^a, P^e) > V(\bar{P}^a, P^e).$$

Our simulations show that $V(\bar{P}^a, P^{ee}+20) = 4.13 \times 10^{11} > 3.34 \times 10^{11} = V(P^{a,h}, P^{ee}) > V(\bar{P}^a, P^{ee})$. Thus if we assume the stationary distribution of agricultural prices the gain from a contract with \$20/ton transfers is at least 82 billion dollars or 25% of the value with zero transfers.

Figure 6 exhibits the initial distribution of occupation, and the distribution after 50 years for transfers per ton = \$0, \$10, \$15 and \$20. For reference, this figure again shows the parameter heterogeneity. Specifically, Figure 6 shows that for the case of transfers that exceed \$15 per ton of net emissions, the area of the Biome that is occupied by cattle farming after 50 years would be minimal. In fact, as shown in Table 1, when transfers change from \$0 to \$20 per ton, the present value of agricultural output goes down from \$536 billion to \$43 billion[13], a loss that is more than compensated by (i) the increase in the present value of net transfers from \$0 to \$ 305 billion, and (ii) the increase in climate services as valued by Brazilians (valued at the shadow price of \$5.76 per ton of net emissions captured) from the change in policy induced by the \$20 per ton transfers, which totals \$277 billion. Figure 7 illustrates, that for transfers exceeding \$10 /ton, much of the changes in land occupation would occur within the first 15 years. Of course, if adjustment costs are higher than we estimated (see Appendix A), this process would slow down.

Table 2 displays the effect of transfers per ton of net CO_2 captured (b), conditional on the price of cattle, on captured emissions in the first 30 years. When P^a = \$44.76, the top quartile in our observations, in the business a usual case, deforestation would produce carbon emissions of 31.6 billion tons in the next 30 years. With transfers of \$20/ton, optimal management induces capture of 16.1 billion tons. These transfers total \$323 billion. The effective cost per ton of the net amount of $C0_2$ emissions saved (47.7 billion tons of CO_2) is only \$6.76/ton. The leverage from transfers goes down as b increases since the business as usual emissions are a constant, but even in the case of $b = 25$ USD/ton, the effective cost of saved emissions is below \$10/ton. Lower agricultural prices increase effective costs, since emissions in the business as usual case drop proportionately more

[13]Recall however that we count the full output as value added, thus exaggerating the loss of agricultural output.

than capture for each b, but effective cost per ton changes are minimal. Table 2 illustrates the gains from trade in instituting a contract that pays Brazil $\$b$ per net ton captured.

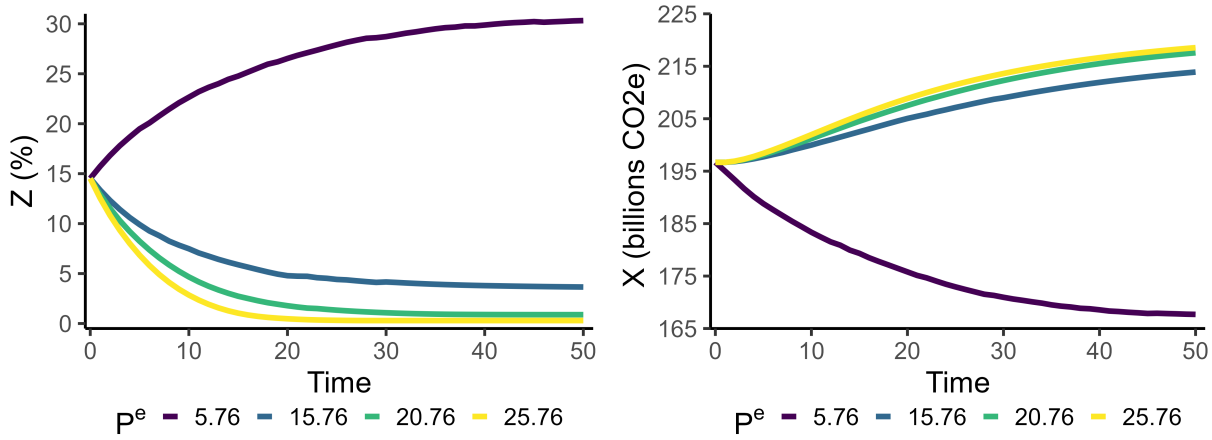

Figure 5: Agricultural Area and Carbon Stock Evolution (50 years)

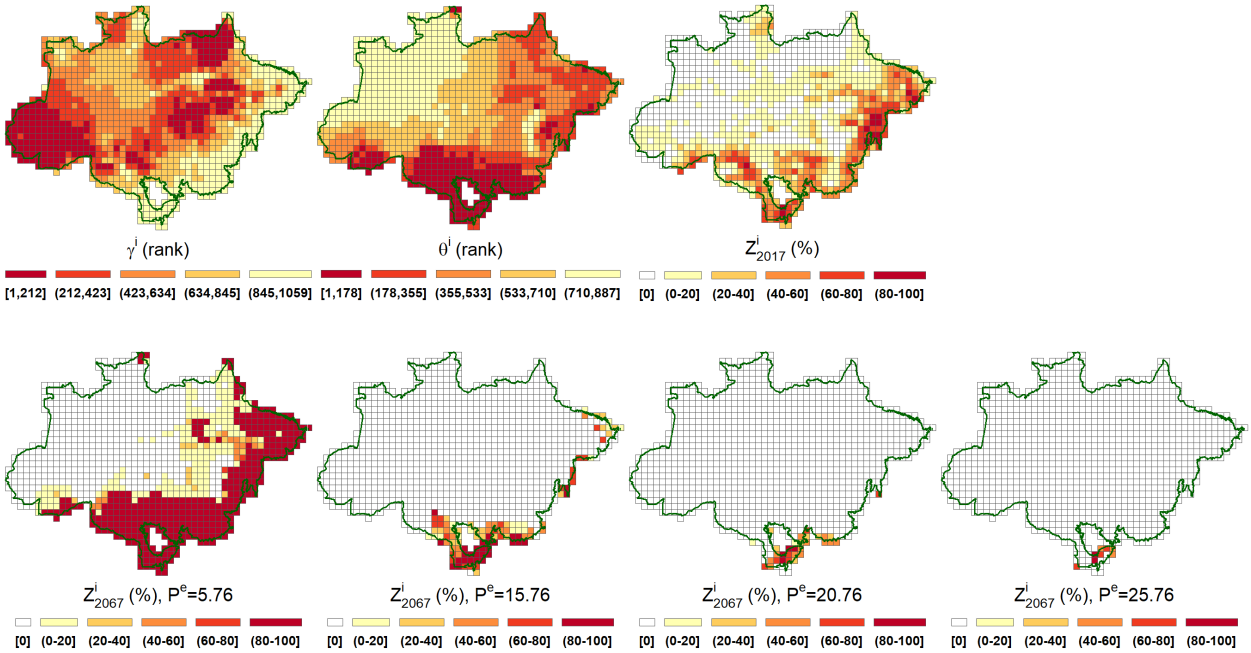

Figure 6

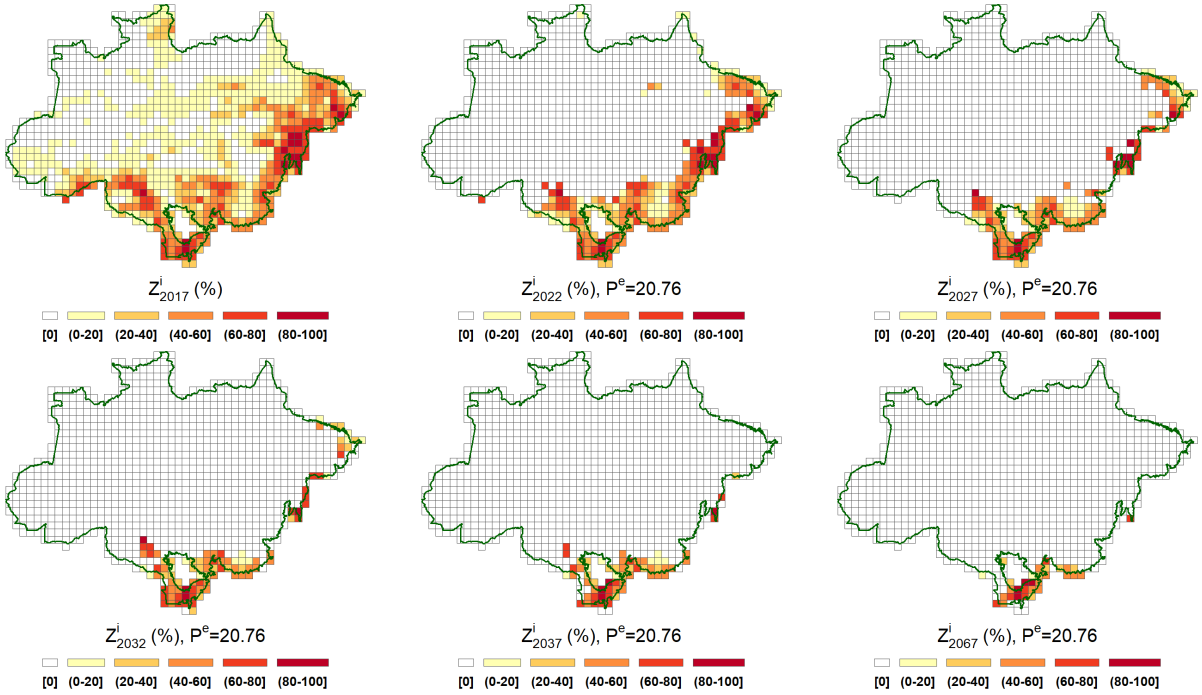

Figure 7

Table 1: Planner Value Decomposition (200 years)

P^a ($)	P^e ($)	b ($)	Agricultural Output ($ 10^{11})	Net Transfers ($ 10^{11})	Climate Services ($ 10^{11})	Adjustment Costs ($ 10^{11})	Planner Value ($ 10^{11})
44.76	5.76	0	5.36	0.00	-1.89	0.12	3.34
44.76	15.76	10	1.27	1.03	0.59	0.08	2.81
44.76	20.76	15	0.66	2.10	0.81	0.14	3.42
44.76	25.76	20	0.43	3.05	0.88	0.20	4.16
38.30	5.76	0	4.06	0.00	-1.37	0.08	2.61
38.30	15.76	10	0.82	1.23	0.71	0.10	2.65
38.30	20.76	15	0.44	2.21	0.85	0.17	3.34
38.30	25.76	20	0.29	3.13	0.90	0.22	4.10

Table 2: Transfer Cost (30 years)

P^a ($)	P^e ($)	b ($)	Net Captured Emissions (billion tons of CO2e)	Net Transfers ($ 10^{11})	Effective cost ($ per ton of CO2e)
44.76	5.76	0	-31.59	0.00	NaN
44.76	20.76	15	14.47	2.17	4.71
44.76	25.76	20	16.13	3.23	6.76
44.76	30.76	25	16.86	4.21	8.70
38.30	5.76	0	-22.26	0.00	NaN
38.30	20.76	15	15.52	2.33	6.16
38.30	25.76	20	16.59	3.32	8.54
38.30	30.76	25	17.05	4.26	10.84

7.3 Results for 81 sites

In this section we summarize results for the case of 81 sites and uncertainty on the price of agricultural output. In this case the trajectories depend on the realization of agricultural prices so that our results involving a single trajectory should be taken as indicative. Notice that the picture is much like in the previous case. As the trajectories in Figure 8 indicate, in the absence of transfers, the optimal choice involves an increase in the agricultural area that may result on tipping of the eastern, southern and central Amazonia to a non-forest system, whereas transfers of at least $15 per net ton, would lead to substantial reforestation. When transfers change from $0 to $15 per ton, the present value over the next 200 years of agricultural output is reduced from $492 billion to $50 billion, a loss that is more than compensated by (i) the increase in (expected) net transfers from $0 to $244 billion, and (ii) the increase in climate services as valued by Brazilians (at the shadow price of $5.76 per ton of net emissions captured) from the change in policy induced by the $ 15 per ton transfers, which totals $52 billion. Figure 9 shows the pattern of concentration is similar to the previous case - agricultural activities are restricted to portions of the southern Amazon. Figure 10 illustrates the spatial dynamics by showing that with the level of adjustment costs we imposed, much of the geographical change in land use occurs within the first ten years, and most of it in the first 20 years. Finally, Figure 11 illustrates the effect of the planner internalizing the potential for forest growth of the different sites. The site that appears in bold is the 5th most productive for cattle farming and at the end of 50 years 68% of its area would still have cattle farming activity when $p^e = 20.76$. Sites ranked 2,3 or 4 in agricultural productivity are completely reforested. The reason is that the site that appears in bold is ranked 78th of 81 sites for forest growth.

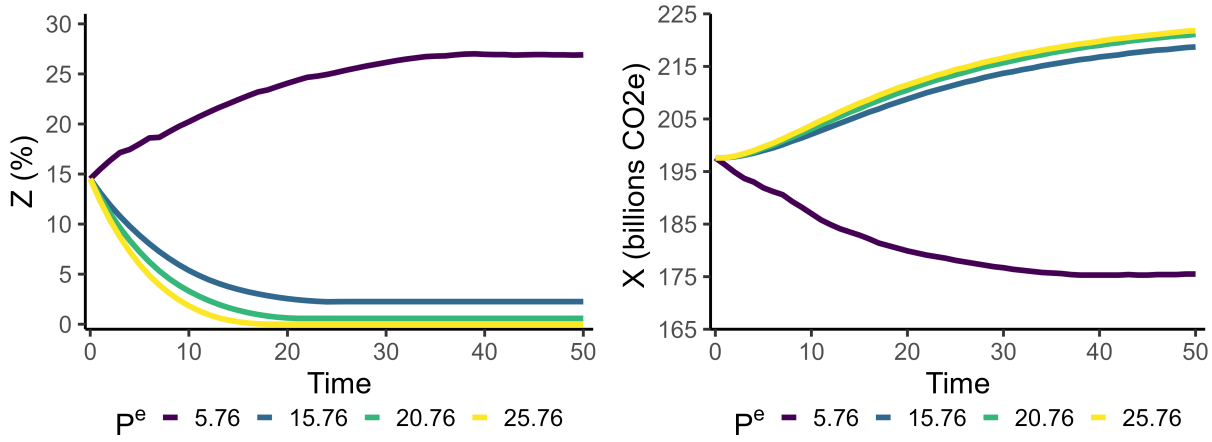

Figure 8: Agricultural Area and Carbon Stock Evolution (50 years)

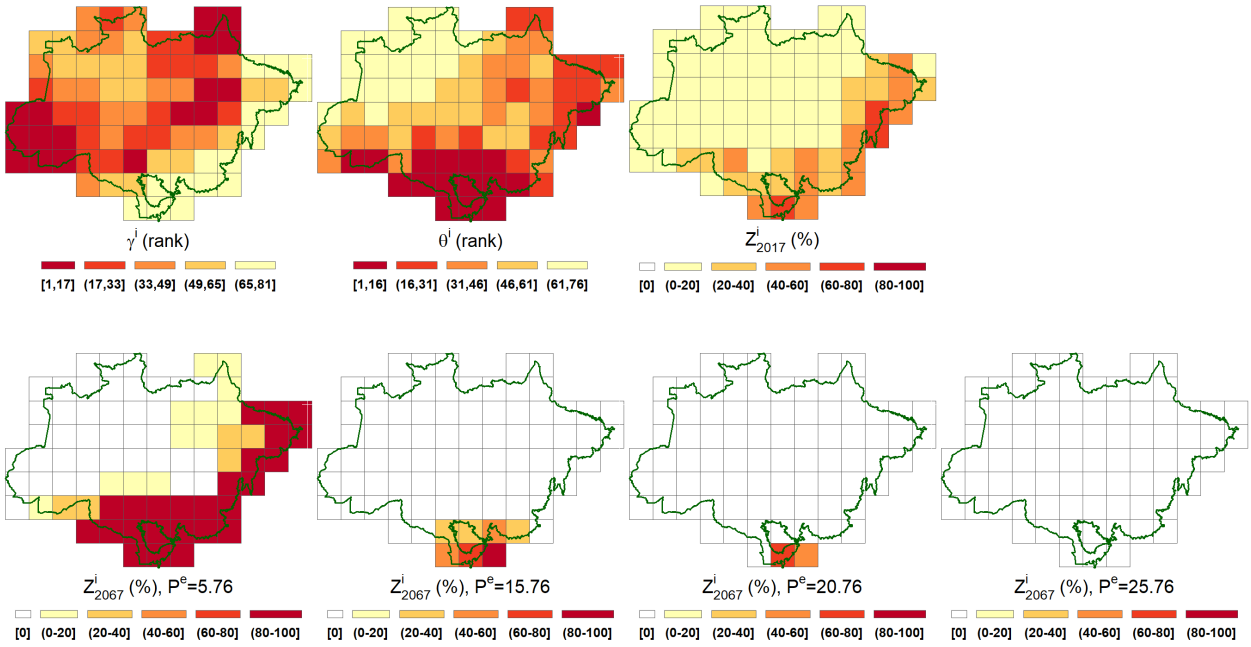

Figure 9: Agricultura Area Evolution by Agricultural Price

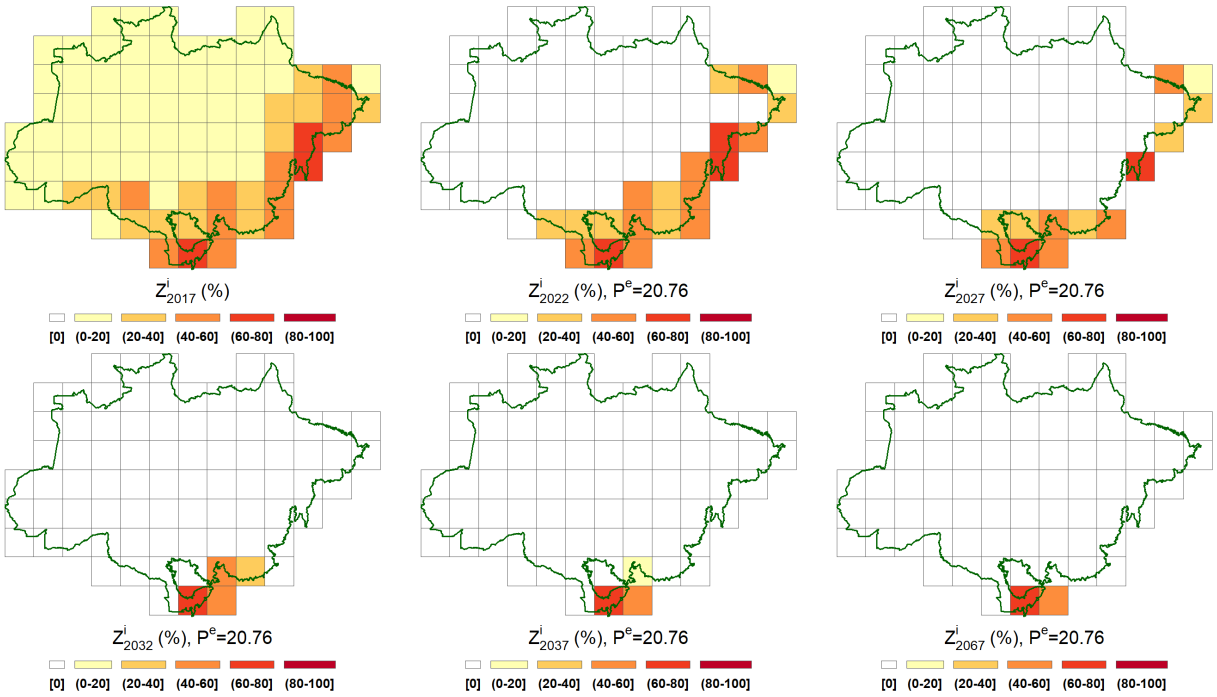

Figure 10: Agricultura Area Evolution by Time

Table 3: Planner Value Decomposition (200 years)

P^e ($)	b ($)	Agricultural Output ($ 10^{11})	Net Transfers ($ 10^{11})	Climate Services ($ 10^{11})	Adjustment Costs ($ 10^{11})	Planner Value ($ 10^{11})
5.76	0	4.92	0.00	-1.58	0.08	3.26
15.76	10	0.92	1.37	0.79	0.12	2.96
20.76	15	0.50	2.44	0.94	0.18	3.69
25.76	20	0.31	3.45	0.99	0.24	4.52

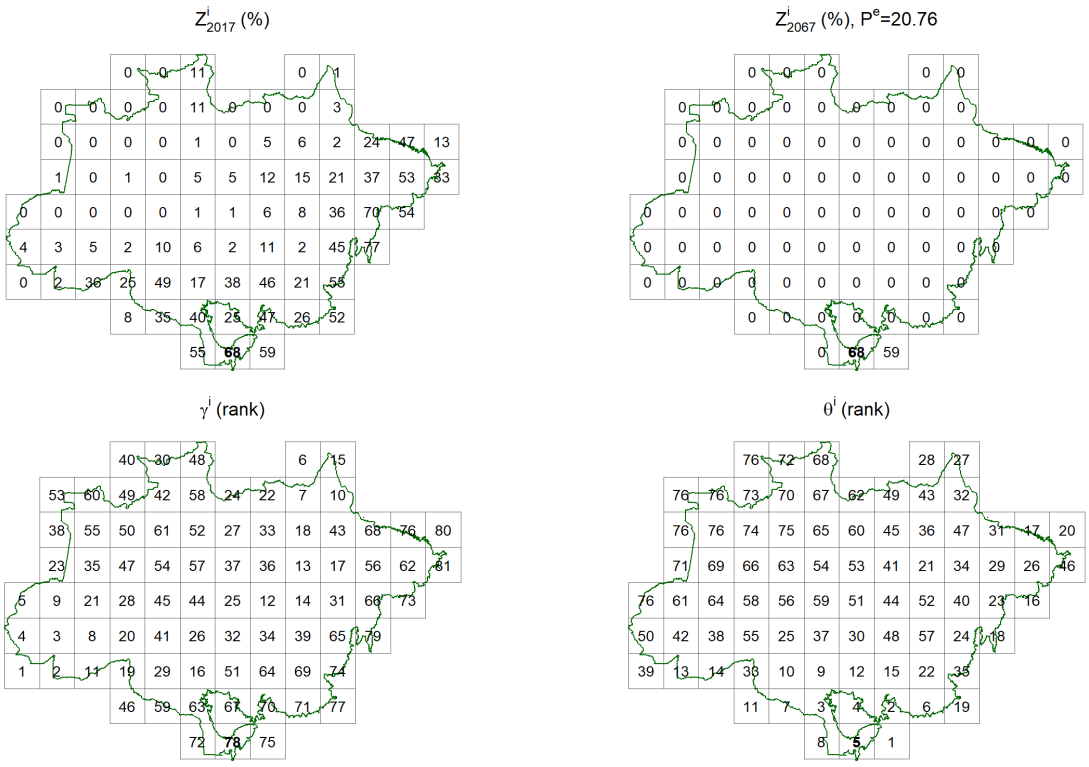

Figure 11

7.4 Results with robustness to parameter uncertainty

This subsection is under active reconstruction as we only recently figured out a much better way to perform the calculations along the lines described previously. We will very soon have results for this subsection.

8 Conclusions

We used a rich data set to study the impact of carbon prices on optimal forest preservation over time and space in the Brazilian Amazon. We produced results for three exercises. First we calculated the shadow prices for emissions that would justify historical deforestation. We then calculated the impact of payments for net emissions captured. The optimal strategy under the calculated shadow price would eventually produce enough deforestation to threaten the survival of the Amazon as a tropical forest. However transfer prices above \$15 per ton of $C0_2e$, increase Brazil's welfare and lead to substantial net reforestation and carbon capture in the optimal trajectory. As a third application, we considered optimal robust controls when the capacity of each site to capture carbon and and the corresponding agricultural productivity is uncertain to the planner.

Our results suggest international carbon payments of \$20 USD/ton can reduce emissions in 47.7 GtCO2e, in 30 years, considering a 2% discount rate and a price for cattle that represents

the highest quartile of observed past prices. This amount represents not only the 16.1 GtCO2e of carbon captured by natural regeneration, for which Brazil will receive the payments, but also 31.6 GtCO2e of avoided emissions from the deforestation that could happen in the 'business as usual' scenario. This combination implies a leveraged carbon mechanism for which the effective carbon cost is $6.76 USD/ton. If we consider a price that represents the lowest quartile of historically observed cattle prices, the effective carbon cost is a only a bit higher - $6.89 USD/ton.

According to Griscom et al. (2017), nature-based solutions such as forest restoration, avoided land conversion, forest management and other practices have the potential of capturing about 11.3 Gt of CO2 per year globally with costs no greater than $100 USD/ton. Our baseline simulation in 2 suggests that optimal management of the Brazilian Amazon can deliver at least 11% of this total at an effective cost that would not exceed $8.54 USD/ton.

Our calculations in this paper ignore some important costs of deforestation. We do not include, for instance, the effect of deforestation on agricultural productivity (Leite-Filho et al. (2021)). We also do not take into account the loss of biodiversity or the possible role of Amazon deforestation as a global tipping point (Steffen et al. (2018)). These are important considerations for future research.

A Data construction

Total available area: $\bar{z}^i$

The amount of available area for the planner's choice (forest or cattle farming) in each site i, $\bar{z}^i$, is computed as the fraction of 30m-pixels from the MapBiomas classified as agriculture (crops + pastures) or forests in 2017. We then multiply this fraction by the area of the site in order to provide a measure in hectares. Notice $\bar{z}^i$ comprises the total site area, excluding areas such as rivers, roads, cities and etc.

Parameters related to carbon dynamics: γ^i, α, *and* κ

We extract a random sample of 1.2M 30m-pixels and selected 893,753 pixels that could be considered primary forests in the year 2017 (pixels with no deforestation at least since 1985). We added *above ground* biomass density data for the year 2017 from ESA Biomass[14]. The biomass data also comes in a grid format with ∼100m resolution, so we spatially matched it to our sample. The original data is measured in biomass density (Mg per ha) but we converted to carbon per hectare, by dividing by 2 (carbon is approximately 50% of the biomass), and then converted to CO2 equivalent, multiplying by 44 and dividing by 12 (based on atomic mass).

ESA Biomass produces 2017 values of carbon captured, which we denote as $co2eHa2017$ by (after conversion), and standard deviations that reflect uncertainty, which we denote as $co2eHa2017SD$ for each pixel. We first run the following model to the pixel data:

$$\begin{aligned}\log(\text{co2eHa2017SD}) = \beta_0 + \beta_1(\log(\text{co2eHa2017}) \; + \\ \beta_2(\text{lon}) + \beta_3(\text{lat}) \; + \\ \beta_4(\text{lon\^{}2}) + \beta_5(\text{lat\^{}2}) \; + \\ \beta_6(\text{lon} \times \text{lat}) + \epsilon\end{aligned}$$

[14] (Santoro and Cartus, 2021)

Table 4: Title

	log(co2eSD_ha_2017)
	(1)
(Intercept)	2.591***
	(0.0206)
log(co2e_ha_2017)	0.9108***
	(0.0008)
lon	0.0743***
	(0.0007)
lat	0.1840***
	(0.0007)
lon square	0.0004***
	(5.54×10^{-6})
lat square	0.0006***
	(1.36×10^{-5})
lon × lat	0.0031***
	(1.16×10^{-5})
Observations	893,753
R^2	0.77794
Adjusted R^2	0.77794

The contribution of the geographic variables to R^2 is small, amounting in total to only 6.5% of the total R^2.

To calculate values for aggregate sites, we first spatially match all pixels in a site and calculate the average CO2e density (Mg/ha) on primary forest (γ^i). We then exponentiate the fitted value from the regression, in Table 4, to obtain a std (Mg/ha) of the site. Notice that this assumes that aggregation does not change the standard deviation relative to the mean of sites.

The parameter α is a carbon depreciation parameter, assumed to be constant across sites. It is set so the convergence time of the carbon accumulation process is 100 years, following Heinrich et al. (2021). Assuming the 100 years-period and a convergence threshold of 99%, we set $\alpha = 1 - (1 - 0.99)^{1/100}$.

Finally, the parameter κ is calibrated through the agricultural net annual emission data at the state level available from the system SEEG.[15] We compute κ as the average of the agricultural net emission divided by the agricultural area from MapBiomas for all States within the Amazon biome weighting by the area of overlap from 1990 to 2019.

Cattle farming productivity: θ^i

Since almost 90% of the historically deforested land in the Amazon Biome that was used for agricultural activities in 2017 was used for pasture, we focus on the productivity of cattle farming

[15]Sistema de Estimativas de Emissões e Remoções de Gases de Efeito Estufa. Available in http://seeg.eco.br/.

for each site. Since we do not have measurements concerning the cost of attracting or redeploying variable inputs in the cattle farming sector, we focus on revenue per hectare. This choice leads to an overvaluation of the contribution of cattle farming in the Amazon to the Brazilian economy.[16] We consider the value of cattle sold for slaughter per hectare of pastureland at the municipal level, from the 2017 Agricultural Census. In order to build a smoother representation of technology, and fill missing values, we used the predicted value of the following specification[17]:

$$\begin{aligned}\text{Slaughter value per hectare} =& \beta_0 + \beta_1(\text{pasture area}) + \beta_2\text{rainfall} + \beta_3\text{rainfall}^2 + \beta_4\text{temperature} + \\ & \beta_5\text{temperature}^2 + \beta_6\text{longitude} + \beta_7\text{latitude} + \beta_8\text{longitude}^2 + \\ & \beta_9\text{latitude}^2 + \beta_{10}\text{longitude} \times \text{latitude} + \epsilon\end{aligned}$$

where *slaughter value per hectare* is the value of cattle sold per pasture area (USD/ha), *rainfall* and *temperature* are the average annual precipitation (mm) and temperature (degrees Celsius), respectively, for the period of 1970-2000, and *longitude* and *latitude* are the geographical coordinates of the municipality centroids. To convert from the municipal level to the site level, we calculated a weighted mean of the predicted value of this regression, with weights based on the share of the municipal area inside the site.

Discount factor (ρ) and adjustment costs cost (ζ)

We use discount rates $\rho = .02$ and .03. We calibrate ζ using the difference in price between forested land and cleaned land and the amount of annual deforestation that occurred from 2008 to 2017 based on Araujo et al. (2022). Notice that the difference in price should reflect both the cost of deforestation and any value of wood obtained in the process. Unfortunately we did not have data that would allow us to compute a separate adjustment cost for decreasing (as opposed to increasing) deforestation, so we opted for symmetry.

Initial values: z_0^i, x_0^i

The approach for computing the initial value for the agricultural area, z_0^i, is similar to that used for the total available area $\bar{z}^i$. The only difference is that we focus only on the fraction of pixels classified as agriculture (crops + pastures) in 2017 before multiplying by the site's area in order to obtain a measure in hectares.

The initial value for the carbon stored in the forests x_0^i is calibrated as $x_0^i = \gamma^i(\bar{z}^i - z_0^i)$, i.e., the carbon stock per hectare of forest times the forest area. Notice that x_0^i is measured in CO2e (Mg). Notice that we assume that all forest at the initial point is primary, which is compatible with equation (2).

Agricultural prices

[16] Thus our calibration is likely to continue to be reasonable even if technological changes discussed in Cardoso et al. (2016) are widely implemented.

[17] Negative predicted values were converted into the minimum positive predicted value - 0.674.

We use a data series on monthly deflated cattle prices (reference date 01/2017).[18] from 1995, the year in which the Real Plan stabilized the Brazilian currency, until 2017. In this paper, we fit a two-state Markov process. Let $\tilde{P}_t^a$ denote observed prices. We calculated a high (low) price $P^{a,h}$ ($P^{a,\ell}$) as the 75th (25th) percentile of the series $\tilde{P}_t^a$. We discretized the prices to $P^{a,h}$ and $P^{a,\ell}$ based on the following criteria. If the initial price is above (below) the median, we start the series at $P^{a,h}$ ($P^{a,\ell}$). If at time t, $P_t^a = P^{a,h}$ then $P_{t+1}^a = P^{a,h}$ unless $\tilde{P}_{t+1}^a \leqslant (P^{a,\ell}$. Symmetrically, if at time t, $P_t^a = P^{a,\ell}$ then $P_{t+1}^a = P^{a,\ell}$ unless $\tilde{P}_{t+1}^a \geqslant P^{a,h}$. We calculate the discrete transition probabilities using observed frequencies. The probability from high to high, for example, is defined as the number of consecutive periods with high prices divided by the number of periods with high prices.

10 sites aggregation with clustering

For considering uncertainty on the value of γ^i, we use the K-means clustering method to group the 1,059 sites into 10 clusters based on γ^i and θ^i. This method aims to partition the data into groups such that sites within the same group are as similar as possible and sites from different groups are as dissimilar as possible.

Before applying the method, we standardize the variables γ and θ to mean zero and standard deviation one. We apply the algorithm from Hartigan and Wong (1979) that can be summarized in five steps. First, we pick the number of clusters ($k = 10$). Second, 10 sites from the 1,059 are randomly selected as the initial centroids (γ^i x θ^i). Third, the remaining 1,049 sites are clustered to the closest centroid according to the Euclidean distance. Fourth, the cluster centroids are updated by calculating the variables' means across all sites in the cluster. Fifth, the third and fourth steps are repeated until the cluster assignment stops changing, reaching convergence. This process results in non-contiguous but more homogeneous regions.

[18] Commodity prices from SEAB-PR. Secretaria da Agricultura e do Abastecimento do Estado do Paraná (SEAB-PR). 2021. "Preço Médio - Recebido pelo Agricultor: boi gordo, arroz (em casca), cana-de-açúcar, milho, mandioca, 1990-2021." Secretaria da Agricultura e do Abastecimento do Estado do Paraná, Departamento de Economia Rural [publisher], Instituto de Pesquisa Econômica Aplicada, Ministério da Economia [distributor]. http://www.ipeadata.gov.br (accessed February 22, 2021)

Table 5: Cluster Summary Statistics

Cluster	θ	γ	γ_{SD}
1	1.91	264.39	134.00
2	3.23	374.48	181.83
3	1.92	395.27	186.90
4	0.29	515.34	205.48
5	2.60	525.40	249.13
6	1.55	528.37	238.11
7	1.00	629.56	271.38
8	0.06	647.84	258.88
9	1.99	655.24	291.55
10	1.81	775.48	326.39

References

Angelsen, Arild. 2017. REDD+ as result-based aid: General lessons and bilateral agreements of Norway. *Review of Development Economics* 21 (2):237–264.

Araujo, Rafael, Francisco Costa, and Marcelo Sant'Anna. 2022. Efficient forestation in the Brazilian Amazon: Evidence from a dynamic model.

Assunção, Juliano, Clarissa Gandour, Romero Rocha, and Rudi Rocha. 2020. The effect of rural credit on deforestation: evidence from the Brazilian Amazon. *The Economic Journal* 130 (626):290–330.

Assunção, Juliano, Clarissa Gandour, and Romero Rocha. 2022a. DETERring deforestation in the Brazilian Amazon: environmental monitoring and law enforcement. *American Economic Journal: Applied Economics (Forthcoming)* .

Assunção, Juliano, Robert McMillan, Joshua Murphy, and Eduardo Souza-Rodrigues. 2022b. Optimal environmental targeting in the amazon rainforest. *The Review of Economic Studies (Forthcoming)* .

Assunção, J. and R. Rocha. 2019. Getting greener by going black: the effect of blacklisting municipalities on Amazon deforestation. *Environment and Development Economics* 24 (2):115–137.

Balboni, Clare, Aaron Berman, Robin Burgess, and Benjamin A Olken. 2022. The Economics of Tropical Deforestation. *The Annual Review of Economics, forthcoming* .

Bemporad, Alberto, Manfred Morari, Vivek Dua, and Efstratios N. Pistikopoulos. 2002. The explicit linear quadratic regulator for constrained systems. *Automatica* 38:3–20.

Cai, Yongyang and Kenneth L. Judd. 2023. A Simple But Powerful Simulated Certainty Equivalent Approximation Method for Dynamic Stochastic Problems. *Quantitative Economics* Forthcoming.

Cai, Yongyang, Kenneth Judd, and Jevgenijs Steinbuks. 2017. A nonlinear certainty equivalent approximation method for dynamic stochastic problems. *Quantitative Economics* 8:117–147.

Cardoso, Abmael S, Alexandre Berndt, April Leytem, Bruno JR Alves, Isabel das NO de Carvalho, Luis Henrique de Barros Soares, Segundo Urquiaga, and Robert M Boddey. 2016. Impact of the intensification of beef production in Brazil on greenhouse gas emissions and land use. *Agricultural Systems* 143:86–96.

Correa, Juliano, Richard van der Hoff, and Raoni Rajão. 2019. Amazon Fund 10 years later: lessons from the world's largest REDD+ program. *Forests* 10 (3):272.

Dominguez-Iino, Tomas. 2021. Efficiency and redistribution in environmental policy: An equilibrium analysis of agricultural supply chains. Tech. rep., Working Paper.

Friedlingstein, Pierre, Matthew W Jones, Michael O'Sullivan, Robbie M Andrew, Dorothee CE Bakker, Judith Hauck, Corinne Le Quéré, Glen P Peters, Wouter Peters, Julia Pongratz, et al. 2022. Global carbon budget 2021. *Earth System Science Data* 14 (4):1917–2005.

Gandour, Clarissa. 2018. *Forest Wars: A Trilogy on Combating Deforestation in the Brazilian Amazon.* Ph.D. thesis, Economics Department, Pontifícia Universidade Católica do Rio de Janeiro.

Griscom, Bronson W, Justin Adams, Peter W Ellis, Richard A Houghton, Guy Lomax, Daniela A Miteva, William H Schlesinger, David Shoch, Juha V Siikamäki, Pete Smith, et al. 2017. Natural climate solutions. *Proceedings of the National Academy of Sciences* 114 (44):11645–11650.

Hartigan, John A and Manchek A Wong. 1979. Algorithm AS 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)* 28 (1):100–108.

Heinrich, Viola HA, Ricardo Dalagnol, Henrique LG Cassol, Thais M Rosan, Catherine Torres de Almeida, Celso HL Silva Junior, Wesley A Campanharo, Joanna I House, Stephen Sitch, Tristram C Hales, et al. 2021. Large carbon sink potential of secondary forests in the Brazilian Amazon to mitigate climate change. *Nature communications* 12 (1):1–11.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. A Smooth Model of Decision Making Under Uncertainty. *Econometrica* 73:1849–1892.

Leite-Filho, Argemiro Teixeira, Britaldo Silveira Soares-Filho, Juliana Leroy Davis, Gabriel Medeiros Abrahão, and Jan Börner. 2021. Deforestation reduces rainfall and agricultural revenues in the Brazilian Amazon. *Nature Communications* 12 (1):2591.

Lovejoy, Thomas E and Carlos Nobre. 2018. Amazon tipping point.

Malhi, Yadvinder, Daniel Wood, Timothy R. Baker, James Wright, Oliver L. Phillips, Thomas Cochrane, Patrick Meir, Jerome Chave, Samuel Almeida, Luzmilla Arroyo, Niro Higuchi, Timothy J. Killeen, Susan G. Laurance, William F. Laurance, Simon L. Lewis, Abel Monteagudo, David A. Neill, Percy Nunez Vargas, Nigel C. A. Pitman, Carlos Alberto Quesada, Rafael Salomao, Jose Natalino M. Silva, Armando Torres Lezama, John Terborgh, Rodolfo Vasquez Martinez, and Barbara Vinceti. 2006. The regional variation of aboveground live biomass in old-growth Amazonian forests. *Global Change Biology* 12 (7):1107–1138.

Santoro, Maurizio and Oliver Cartus. 2021. ESA Biomass Climate Change Initiative (Biomass_cci): Global datasets of forest above-ground biomass for the years 2010, 2017 and 2018, v3.

Scokaert, Pierre O.M. and James B. Rawlings. 1998. Constrained linear quadratic regulation. *IEEE Transactions on Automatic Control* 43:1163–1169.

Scott, Paul. 2014. Dynamic discrete choice estimation of agricultural land use. Tech. rep., TSE Working Paper.

Souza-Rodrigues, Eduardo. 2019. Deforestation in the Amazon: A unified framework for estimation and policy analysis. *The Review of Economic Studies* 86 (6):2713–2744.

Steffen, Will, Johan Rockström, Katherine Richardson, Timothy M Lenton, Carl Folke, Diana Liverman, Colin P Summerhayes, Anthony D Barnosky, Sarah E Cornell, Michel Crucifix, et al. 2018. Trajectories of the Earth System in the Anthropocene. *Proceedings of the National Academy of Sciences* 115 (33):8252–8259.

Thangavel, S., S. Lucia, R. Paulen, and S. Engell. 2018. Dual robust nonlinear model predictive control: A multi-stage approach. *Journal of Process Control* 72:39–51.