This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Carbon prices, forest conservation and reforestation 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 and Princeton Universities

September 4, 2026

Abstract

Deforestation for cattle ranching in the Brazilian Amazon emits carbon and reforestation absorbs carbon. The social productivities for these alternative activities vary across locations. We analyze a spatial/dynamic model of efficient land allocation to establish a benchmark for policies. We treat cattle prices as stochastic and location-specific productivities as uncertain when assessing the consequences of imposing alternative prices of carbon emissions. Modest increases in carbon prices would incentivize Brazil to choose policies that capture a significant amount of greenhouse gasses in the next 30 years. Our analysis pinpoints tropical forest management as an important contributor to climate change mitigation.

***Keywords*—** rainforests; reforestation, renewable resources; ambiguity aversion; land allocation

This is a substantial revision of an earlier manuscript entitled "Carbon prices and forest preservation over space and time in the Brazilian Amazon". We thank Pengyu Chen, Bin Cheng, Leonardo Gomes, Jiaying Liao, Patricio Hernandez, João Pedro Vieira, Zhaoyang Xu, and Daniel Zhao for their expert research assistance, Andy Atkeson, Joanna Harris, and Diana Petrova for their helpful comments. Carmen Quinn provided editorial assistance. Assunção's research received support from Norway's International Climate and Forest Initiative (NICFI) and the Amazônia 2030 Initiative through grants to Climate Policy Initiative and the Pontifical Catholic University of Rio de Janeiro; Hansen's research was supported by the Griffin Applied Economics Incubator Project on Policy-making in an Uncertain World, by an EPIC/Argonne National Laboratory collaboration award, and by the Haddad Fund for Economics Research at the Becker Friedman Institute for Economics at the University of Chicago; and Scheinkman's research was supported by the Columbia Climate School, by Princeton University and by the Haddad Fellowship Fund at Columbia University. Edited by Andrew Atkeson.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

1 Introduction

This paper investigates the potential social benefits of designing prudent policies of reforestation and conservation in the Brazilian Amazon Biome (henceforth, the Brazilian Amazon) through the lens of a spatial and dynamic model. We build the model to capture the trade-off in land allocation between agricultural production[1] and forest conservation or regeneration.

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 make up 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, which includes an area the size of Texas, has resulted in high emissions. If we were to conceptualize the Brazilian Amazon as a country on its own, it would emerge as an outlier in a plot of countries' emissions relative to GDP (see Figure 1).[2]

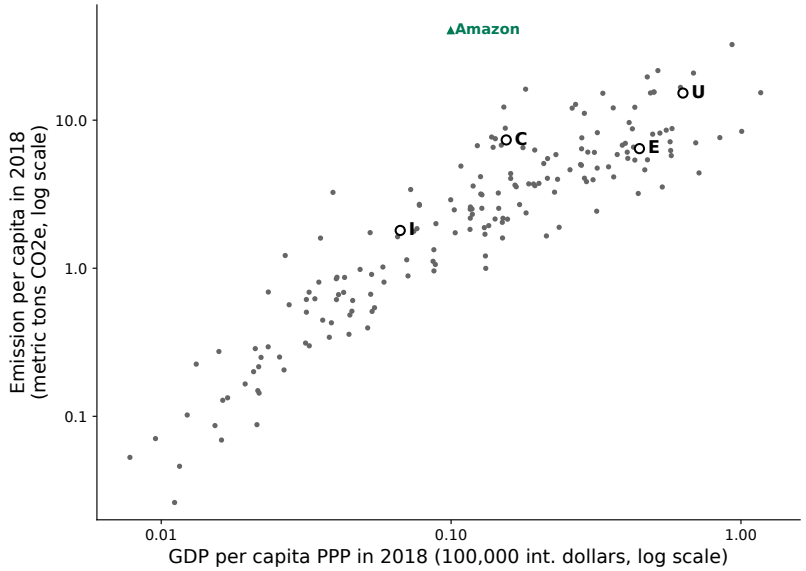


Figure 1: Each dot represents a country in 2018, except for the European Union and the Brazilian Amazon. Highlighted letters stand for (C)hina, (I)ndia, (E)uropean Union, and (U)nited States. Sources: World Bank Data, downloaded on March 2021; Fatos da Amazônia 2021 (www.amazonia2030.org).

We formulate the problem through the lens of a social planner. Massive deforestation of the Brazilian Amazon started with the military government's "Operação Amazonia," which offered land, infrastructure, subsidies, and other incentives for internal migrants to "occupy" the region. During this occupation, deforestation was a necessary condition for accessing subsidies and claiming property rights. The 1988 Consti-

[1] Since 85% of deforested and not abandoned land in the Amazon biome is currently used for pasture, we identify agriculture with cattle farming in this paper, and use the two terms interchangeably.

[2] The World Bank data do not include emissions from land use, land use change and forestry (LULUCF), which are the bulk of emissions in the Amazon. However, for many wealthy countries, including the US (see US LULUCF) and the EU (See EU LULUCF) these emissions are negative.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

tution, under democratization, reduced the ambiguity associated with the legality of deforestation during the military regime and the 1965 Forest Code (Chiavari, Lopes and de Araujo, 2021). Since then, the expansion of deforestation by private actors, even if in possession of a land-title, has often been illegal but tolerated by the authorities. Additional amnesties for illegal occupation of the forest were included in the revised 2012 Forest Code. However, Brazil has also successfully implemented 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 municipal priority list resulted in a reduction of more than 80% in deforestation rates between 2004 and 2012 (Assunção et al. (2023*b*))[3] Therefore, current deforestation is more likely to reflect the valuation of the forest and its alternative uses by governments than by decentralized land-owners. In addition, well-documented edge effects, including fires resulting from human activities at the boundaries (e.g., Nepstad et al. (1999)) and the lower carbon-absorption capacity of forest fragments of less than 100km^2 indicate that conservation/reforestation should be optimized at a scale far beyond typical private holdings in the Amazon.

Our optimization model is dynamic and quantitative and uses detailed location information. We explicitly account for the dynamics of carbon accumulation in the forest — a crucial ingredient in providing credible measures of the potential role of conservation and reforestation in the Amazon rainforest in moderate global warming at different horizons.[4]

With our dynamic model, it is possible to assess the potential for carbon sequestration in the Brazilian Amazon for any time-horizon (see *e.g.*, Figure 5). Later in the paper, we report tables and figures that display hypothetical land allocations for the next 15 or 30 years, when the world is almost certain to continue to produce net emissions[5]. Introducing forest restoration into the analysis not only provides a more complete assessment of optimal land-use cover, but also has critical policy implications. Although conservation and forest preservation are essential to protect the carbon stocks stored in the Brazilian Amazon, carbon removal from large-scale restoration can significantly contribute to offset other sources of global warming.[6] As underscored by United Nations Secretary-General António Guterres:

> We will not be able to contain global warming below 1.5 degrees in the next few years. Overshooting is now inevitable. (…) It is possible to anticipate as much as possible to reach net zero and then to be consistent with negative net in the future to allow temperatures to again go down and the 1.5 still remains, according to all the scientists I met.[7]

The potential contribution of carbon capture in forests is the reason why establishing credible carbon markets that include carbon capture in forests is a priority of the Presidency of COP30.

Our data document a large cross-sectional variability in cattle farming productivity and in the potential

[3] See also Gandour (2018), Assunção and Rocha (2019), Assunção et al. (2020, 2023*a*).

[4] For an online support notebook, see the link: https://lphansen.github.io/Amazon/intro.html.

[5] For instance, China announced in September 2025 a new Nationally Defined Contribution that promises net-zero CO_2 by 2060.

[6] See Assunção et al. (2025) for further information.

[7] Secretary-General's remarks to the High-level Event on Early Warnings for All at the Extraordinary Session of the World Meteorological Congress, 22 October 2025. Available here link.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

absorption of carbon in the Brazilian Amazon. This cross-sectional heterogeneity often leads to corner solutions in land allocation in which economic activity is concentrated in one activity. This drives our choices of numerical-solution methods. The productivity heterogeneity highlights the importance of incorporating a spatial dimension in the model. To account for these locational differences, we divide the Amazon region into various sub-regions or sites, each of which has its own capacity to support agriculture and forestry. Although the model has considerable cross-sectional richness, it is nevertheless highly stylized for reasons of tractability and transparency. Spatial correlations emerge because of similarities in the productive possibilities.

To set the stage for our analysis, we first infer a shadow price for emissions revealed by the change in forest cover during 1995-2008, assuming a robustly optimizing planner. The year 1995 is the first date at which we have reliable data on cattle prices.[8] and the year 2008 marks the beginning of the pay-for-performance Amazon Fund, financed primarily by the Norwegian and German governments.

We use the inferred shadow price for simulations designed to capture "business-as-usual." We also use this shadow price to measure the value to Brazilians of the "forest services" provided by preserved areas. These services include climate services, as well as the economic value of production that occurs without affecting the forest.[9]

We then study the impact of adding external payments for *net* capture of CO_2e in the Amazon. We assume no payments for simple conservation - forests are conserved because the incentives for net reforestation naturally engender a shadow tax on deforestation. We obtain our quantitative results in three steps.

Step 1 uses our finest spatial resolution (1,043 sites; about $67.5\text{ km} \times 67.5\text{ km}$ at the equator)[10] and abstracts from agricultural price uncertainty by fixing the agricultural output price at the stationary mean of the estimated two-state Markov process.

Step 2 incorporates location-specific productivity uncertainty. While Step 1 treats site-level productivities for carbon sequestration and agriculture as known, these objects are measured with substantial error and are not directly observed in the cross section. We therefore use regressions to recover the inputs for the planner's problem, so uncertainty in regression coefficients translates into uncertainty in implied site productivities; throughout, we continue to abstract from agricultural (cattle) price uncertainty. Although convenient, this baseline construction is *ad hoc* and uncertain. We therefore embed this uncertainty in the planner's problem by allowing for ambiguity about the baseline itself: the estimation outcome serves as a "prior," which the planner treats as a reference in a penalized sensitivity analysis. The penalty bounds deviations from the baseline and can be interpreted as the inverse of ambiguity aversion. Solving the resulting high-dimensional problem requires numerical methods, and we use a Markov chain Monte Carlo method based on Hamiltonian dynamics, which is well suited to high dimension and improves on the Metropolis–Hastings approach. We deploy it in a novel way to confront "deep uncertainty" about productivity parameters, equivalently implementing a robust Bayesian control problem. Accounting for ambiguity changes the shadow prices central to

[8] Until mid-1994, Brazil went through a period of very high and volatile inflation.

[9] These include forest products such as natural rubber, nuts, and açaí, along with sustainable timber.

[10] We constructed 1887 sites, each aggregating 2,250 30 m pixels of MapBiomas data considering the top-left corner as the origin, and dropped all sites with less than 3% overlap with the Amazon Biome.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

our analysis, with an offsetting effect on implied transfer payments.

Step 3 moves to a coarser resolution with 78 sites[11] to incorporate stochastic agricultural prices. At this resolution we compute optimal policies under stochastic prices and, for comparison, under a deterministic price. Cross-sectional heterogeneity and natural state constraints preclude standard recursive methods for the continuous-time Hamilton-Jacobi-Bellman equations, so we use and extend Modified Predictive Control (MPC) methods (*e.g.* Scokaert and Rawlings (1998), Bemporad et al. (2002), Thangavel et al. (2018)). MPC methods approximate state inequality constraints via an interior-point method and incorporate Markov uncertainty using a shorter horizon than the overall control horizon. Aggregated biome-level outcomes are close to those in the invariant-price case, indicating that stochastic agricultural prices play a minor role in our results.

Our simulations show that, in the scenario without additional international payments, the deforestation of the Amazon biome would reach 20% in the next 30 years. Lovejoy and Nobre (2018) suggests that this level of deforestation could trigger a tipping point for the forest, while the more recent and slightly less pessimistic analysis in Flores et al. (2024) indicates that this level of deforestation would be sufficient for the Amazon hydrological cycle to be unable to support the rainforest in substantial areas of the current biome.

However, additional payments of at least $25 USD/ton not only shield us from these disastrous possibilities, but would instead trigger forest restoration on a large scale, while providing sufficient transfers to more than fully compensate Brazil for the loss of agricultural income. The carbon dynamics of the model indicates a net carbon capture of 10 $GtCO_2e$ over 15 years, compared to net emissions of 12 $GtCO_2e$ in the scenario without international payments. During a 30-year period, net carbon capture would reach 18 $GtCO_2e$ instead of emitting 15 $GtCO_2e$, resulting in an effective carbon cost of less than $14 USD/ton when considering the total impact on Amazon emissions. In this sense, the carbon sink potential of secondary forests emphasized by the literature can be realized with sufficient additional carbon payments (Griscom et al. (2017), Heinrich et al. (2021), Liang et al. (2025) and Fesenmyer et al. (2025)).

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 shows how we confront parameter uncertainty. Section 5 summarizes how we use a large collection of relevant data sets to calibrate the model. Section 6 discusses the numerical methods used to compute solutions to the social planner's maximization problem. Our results are presented in Section 7, which includes subsection 7.5, where we present a simple bonding mechanism that would guarantee that Brazil would have no incentive to deforest previously reforested areas. This is followed by our conclusions and suggestions for further work.

2 Related Literature

A recent body of work uses discrete-choice models to study the link between agriculture and *deforestation*, while ignoring the potential for forest regeneration. The literature on deforestation is surveyed in Balboni et al.

[11] We aggregated 16 of our original sites produced by aggregating 30 m pixels as in footnote 1, and again dropped units with less than 3% biome area, leaving 78 sites.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

(2023), who note that many investigations of agricultural expansion and deforestation are static and unable to explore the transitional dynamics as they play out over heterogeneous regions. For example, Souza-Rodrigues (2019) and Dominguez-Iino (2021) develop static approaches, focusing on the role of the transportation network and trade in the design of policies.

Following Scott (2014), some recent papers build dynamic discrete-choice models to study deforestation. This modeling approach exploits the convenient assumption of type-I extreme value shocks that are independent over time and location-specific state variables. For example, Araújo, Costa and Sant'Anna (2025) build a dynamic discrete-choice model of land allocation to evaluate the impact of deforestation taxes. Farrokhi et al. (2025) study the impact of international trade costs on global deforestation.

The inclusion of forest carbon dynamics is a key distinguishing feature of our analysis and allows us to study the impact of reforestation, in contrast to the literature that focuses on deforestation. Consider, for example, the arguably most ambitious recent analysis in Araújo, Costa and Sant'Anna (2025) and the one closest to our investigation. Although in the introduction of that paper the authors claim that "...we estimate the carbon-efficient forest cover in the Brazilian Amazon", they do so by primarily focusing on controlling deforestation. This is due in part to the fact that the empirical measurements used in their analysis are limited area units (pixels) that were not deforested before 2000. The resulting empirical focus limits the scope of welfare analysis in ways that we seek to avoid, and in so doing is critical for us to assess more fully the potential of rainforests to capture carbon. Thus, the heterogeneous potential of alternative locations for nature-based carbon capture is a central measurement challenge for us, as it is a direct input into our welfare analyses. Moreover, Araújo, Costa and Sant'Anna (2025) find that the computations in their setup are sufficiently difficult that their hypothetical policy analyses only consider steady states, whereas transitional dynamics are central to our investigation.

Our intertemporal perspective allows us to study the transition dynamics in response to alternative carbon price trajectories. According to IPCC (2023), the next ten to thirty years are critical to the use of new approaches to address climate change, including the regeneration of tropical forests.

In the scientific literature, the potential for carbon capture through forest regeneration is widely acknowledged (*e.g.,* Strassburg et al. (2020)). However, we did not identify any papers that incorporate both the dynamics of carbon accumulation and the opportunity costs to determine the economic viability of reforestation on a large scale.

As an alternative, Busch and Engelmann (2017) estimate the marginal abatement cost curves for various rainforests around the world. This approach investigates local changes through the use of a Poisson regression model. Their approach deliberately sidesteps the use of an explicit dynamic model of competing land usages, rationalized in part by their marginal characterization. An essential input into their regression model is potential gross revenue as measured by FAOSTAT database provided by the Food and Agricultural Organization of the United Nations. Hypothetical changes in carbon price are implemented as shifts in this revenue measure. In contrast, we incorporate estimates of potential productivity for agriculture, from municipal agricultural census, and carbon absorption, from estimates derived from satellite data, into our analysis. Our input of productivity is central to our analysis as low agricultural productivity in many locations in

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

the Brazilian Amazon has a great impact on our conclusions. In contrast, Busch and Engelmann (2017) choose to drop cattle revenue, the activity that occupies the vast majority of deforested land that has not been abandoned, from their featured measurements, making their analysis of limited interest for the Amazon rain-forest. Our structural model with competing productivities of land use opens the door to global policy assessments, distinct from the local justifications for marginal abatement curve estimation.

Finally, there are applications of Global Timber Models (GTM), which feature the dynamics of wood products. Austin et al. (2020) is a recent example of this approach in the study of forest conservation, and Kim et al. (2018) provides a more complete description of the GTM models and also gives a specific implementation. GTM models are calibrated using data on timber production and prices to determine carbon stocks accumulation based on optimal forest management. However, wood production in the Amazon "rarely adopts forest management practices, ... [and is] extensive, predatory and unplanned",[12] making it inadequate for GTM calibration. In fact, Table 3 in Lentini et al. (2019) reports that production of timber from the Amazon forest between 1998 and 2018 fell by more than 2/3 and moved out of the exhausted areas in the "arc of deforestation" in the east to new areas in the southwest Amazon. Furthermore, most deforestation in the Amazon results from illegal activities rather than timber extraction, primarily occurring on public lands and driven by land grabbing, periodically incentivized, tolerated, facilitated and rewarded with amnesties by the federal and state governments.[13]

Quantification of uncertainty is important in many disciplines, including economics. We use some long-standing, as well as recent, constructs from decision theory, control theory, and statistics to extend the usual uncertainty analyses practice. A typical sensitivity analysis in applied econometrics and economic dynamics proceeds entirely from the perspective of the researcher and not from that of the economic agents inside the model. For example, one endows individuals, firms, or a policy maker with the policy-dependent probabilistic outcomes with full confidence, in spite of the uncertainties understood by an external researcher. In our case, this would mean that our planner fully embrace productivity measures, including ones that are not tied exclusively to how the land has been used in the past. We proceed differently.

As a baseline, we use a convenient Bayesian regression approach to produce an *ex ante* posterior distribution for the productivity measurements for all of the sites, both for agriculture and for carbon absorption. We then use this posterior as a *baseline prior* for the planner. We then explore the consequences of different degrees of *ambiguity aversion* on the part of the planner. This ambiguity reflects skepticism in baseline Bayesian computations, treating them instead as a convenient approximation. The formulation of ambiguity aversion we consider (a special case of what is called smooth ambiguity in the decision theory literature) is implemented as a robust approach that looks across alternative priors that deviate from the baseline. With this interpretation, a larger degree of ambiguity aversion effectively imposes utility costs for deviating from the baseline prior. A revealing outcome of the decision problem is to isolate which measurement uncertainties are most consequential to the decision problem, a computation that is not part of more typical approaches to

[12] See Lentini et al. (2019) reported with our translation.

[13] Assunção, Gandour and Rocha (2015), Assunção et al. (2023*b*). Laurance, Clements and Sloan (2014) are examples of a large literature documenting how the construction or paving of new roads without increasing police force facilitates illegal occupation in tropical forests.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

uncertainty quantification. When presenting our results, we include ambiguity neutrality or full confidence in the baseline Bayesian probabilities as a featured special case. In summary, we move the uncertainty concerns about productivity measurements "inside the decision problem" when deducing prudent policies instead of having the planner embrace our probabilistic measurement approximation with full confidence.

3 Model

We pose the problem of a fictitious social planner who considers the trade-off 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 $t \in [0, T]$ is 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 lowercase letters the potential state realizations. At date t,[14]

$$\begin{array}{lll} Z_t & \stackrel{\text{def}}{=} & (Z_t^1, Z_t^2, ..., Z_t^I) \quad \text{vector of agricultural areas expressed in hectares} \\ X_t & \stackrel{\text{def}}{=} & (X_t^1, X_t^2, ..., X_t^I) \quad \text{vector of carbon captured expressed in Mg } CO_2e \end{array}$$

We use the notation $Z \stackrel{\text{def}}{=} \{Z_t : 0 \leq t \leq T\}$ to denote the corresponding process that evolves over time, and similarly for other states and controls. In addition we assume that agricultural output-price is described by a Markov process P_t^a, an index of cattle prices in Brazil expressed in 2017 US dollars.[15] 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 represented as an intensity matrix $\mathbb{M}$ with non-negative entries off-diagonal entries $\mathsf{m}_{\ell\ell'} \geq 0$ 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, $\exp(\tau\mathbb{M})$ is computed using a matrix counterpart to a power series.

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

$$0 \leq Z_t^i \leq \bar{z}^i$$

where $\bar{z}^i$ is the amount of land in the Amazon biome available for agriculture at site i.[16] Let $\dot{Z}_t$ be the time

[14] CO_2e denotes equivalent CO2.

[15] 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).

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

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

derivative of Z at date t.

The evolution of X^i introduces an important asymmetry in our problem. We 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:

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

The first term on the right side of (2) connects deforestation to a loss in captured carbon. The site-specific parameter $\gamma^i > 0$ denotes the density of CO_2e that is present in a primary forest in site i.[17] The next term expresses the growth in captured CO_2e, when the size of the forest in site i is held constant. The mean-reversion coefficient $\alpha > 0$ guarantees that if one lets the forest grow undisturbed in a deforested area, it would reach $100[1-\exp(-\alpha 100)]\%$ of the maximum captured CO_2e per hectare in 100 years, as in Heinrich et al. (2021). In our case, we choose α such that $100[1 - \exp(-\alpha 100)] = 99\%$. In Remark 3.5, we argue that at the optimum, one of the controls is always zero, which introduces additional binding constraints into the analysis. We write this constraint as: $U_t^i V_t^i = 0$.

We model the value of the cattle output on site i at time t, as $P_t^a A_t^i$, where A_t^i is proportional to the land allocated to cattle farming,

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

and θ^i is a site-specific productivity parameter that, in particular, incorporates the ratio of local prices to the index of national prices. Let

$$\varphi^i := (\gamma^i, \theta^i),$$

and φ the full vector of φ^i.

All locations contribute to emissions through the capture of carbon and emissions resulting from 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 parameter κ captures the emissions that result from cattle farming.[18] We include asymmetric, land-

[17] We thank Stephen Pacala for a discussion on the equation of carbon accumulation. For simplicity, equation (2) assumes that all deforestation occurs in primary forests, which is not far from what has been observed in the Brazilian Amazon.

[18] About 75 percent of the emissions from agricultural activity in the Amazon is the result of the natural digestive process of cattle. Another approximately 21 percent is from soil management; therefore, for simplicity, we assume that the cattle herd per hectare does not vary and that productivity variations come mainly from transportation costs and carcass weights.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

use-change costs with contributions from each site,

$$LC_t = \frac{\zeta_1}{2}\left(\sum_{i=1}^{I} U_t^i\right)^2 + \frac{\zeta_2}{2}\left(\sum_{i=1}^{I} V_t^i\right)^2.$$

Importantly, the adjustment cost depends on the aggregate change in land use. Thus, we assume that resources for deforestation or the removal of farms can move across the Amazon. The assumption that resource constraints are at the level of the biome favors (re) deforesting fully a site before starting on another one. Since our sites' area exceed 4,400 km^2 optimal trajectories would not create fragments.

The planner faces ambiguity in the parameters γ^i and θ^i that govern the land use productivity for each site i. The planner takes as input the *ex ante* (to the decision problem) posterior distribution constructed from data with a conveniently chosen likelihood and prior distribution.[19] This construction requires some cross-sectional extrapolation since we have limited direct evidence for some of the sites. The estimation and extrapolation induce dependencies in the posterior distribution for φ. The planner confronts the parameter ambiguity by performing a sensitivity analysis: minimizing the planner's objective by altering the posterior distribution of φ subject to a relative entropy or Kullback-Leibler penalty scaled by a parameter ξ. Larger values of ξ imply a larger penalty and, therefore, less aversion to ambiguity. Setting $\xi = \infty$ gives the planner full confidence in the baseline posterior distribution.

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.

To pose the robustly optimal planning problem for the Brazilian Amazon, we start by computing the intertemporal objective conditioned on the parameter vector φ, taking into account any pure risk considerations:

$$f(d,\varphi) = \mathbb{E}\left(\int_0^\infty e^{-\delta t}\left[P_t^a \sum_{i=1}^{I} A_t^i - P^e \sum_{i=1}^{I}\left(\kappa Z_t^i - \dot{X}_t^i\right) - LC_t\right] dt \middle| \varphi\right) \tag{5}$$

subject to (2) and $0 \leq Z^i \leq \bar{z}^i$, the total area of site i. Here, δ is the subjective discount rate and d denotes the entire sequence of hypothetical decisions contingent on the relevant agricultural prices.

The exogenously specified emissions price, P^e, is an input into the analysis that allows us to explore how costly it will be to make important changes in deforestation outcomes in Brazil. Its magnitude reflects the sum of the marginal value attributed by the planner to emission and any monetary transfers obtained from

[19] One limitation of our dynamic model is that it is not well suited to include potential real-time learning. To allow for this in an interesting way, we would have to disguise the parameters from the observations of the Markov states and controls in order that the learning not be degenerate. This would require adding stochastic complexity to the modeling. Even if we were to do this, we would need to include additional state variables necessary to capture learning in a recursive way, taking into account that the planner will have more evidence about land productivities in the future. Thus to capture this dynamic learning in a meaningful way would make model solutions and analysis even more computationally demanding. Also, rather than being purely passive, this learning could offer the potential for experimentation as a way to gain a better understanding of site-specific productivities. While dynamic learning potentially interesting, such an exercise is beyond the scope of our current analysis.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

others, such as sales in carbon emission markets.

We adopt an *ex ante* representation of the decision problem. Let π denote the baseline distribution over the parameter vector φ, constructed with computationally tractable Bayesian method. The ambiguity-averse planner ranks alternative decision processes by solving the minimization problem:

$$\min_{g} \int \left[f(d,\varphi) + \xi \log g(\varphi) \right] g(\varphi) d\pi(\varphi) \tag{6}$$

subject to $\int g(\varphi) d\pi(\varphi) = 1$ where f is given in (5). In this formulation, $g(\varphi)d\pi(\varphi)$ represents an altered distribution over the parameter vector φ and $\xi \int \left[\log g(\varphi)\right] g(\varphi) d\pi(\varphi)$ penalizes departures from the baseline posterior distribution $d\pi(\varphi)$. To construct a robustly optimal allocation of land in the Brazilian Amazon over time and across space, the planner solves:

Problem 3.1.

$$\max_{d \in \mathcal{D}} \min_{g \geq 0, \int g d\pi = 1} \int_{\mathcal{B}} f(d,\varphi) g(\varphi) d\pi(\varphi) + \xi \int_{\mathcal{B}} \log g(\varphi) g(\varphi) d\pi(\varphi).$$

Remark 3.2. *While it is not our job as external analysts to dictate how uncertainty averse society should be, we can draw on insights from robust Bayesian methods to assess the plausibility of the different ξ settings. Since we treat ξ as a penalization parameter, direct interpretation can be challenging. By applying the Min-max theorem, for each ξ we can deduce an implied worst-case probabilities, under which the robust planner solution optimizes, while being fully committed to this probability. Such distributions reveal insight as to the plausibility of the penalty parameter. We perform some ξ sensitivity in our calculations below.*

Remark 3.3. *The objective function* (5) *values agricultural output by the value of sales, thus assuming that inputs to production do not have alternative use. This choice is dictated by the lack of data on the cost of attracting or redeploying agricultural inputs, but it biases the results in favor of agricultural use.*

Remark 3.4. *The only explicit interaction across sites in the objective function* (5) *occurs through the land-use-change costs. This interaction is intended to be the result of a less than perfectly elastic supply of resources needed to change land use at the level of the whole Amazon. In addition, sites are spatially related by their similarity in productivities for alternative ways to use the land.*

Remark 3.5. *To show 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 that you 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 $\gamma^i \Delta$ 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 discount rate 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 3.6. *Optimization problem 5 does not involve the stocks of (extended) carbon in the atmosphere generated by activities in the Amazon biome. However, given emission trajectories from the optimal solution, one could use geo-science inputs to inform the mapping of emissions from the Brazilian Amazon and*

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

elsewhere into carbon in the atmosphere to compute the impact on the evolution of carbon stocks. This would require a much more comprehensive model that is beyond the scope of this particular exercise.

4 Implementing parameter uncertainty

We treat parameter uncertainty probabilistically by starting with a baseline subjective prior to the parameters. We form this baseline "prior" conditioned on available data. Rather than assuming a full commitment to this baseline distribution, we allow for some skepticism by exploring sensitivity to distributional changes. We limit the scope of the sensitivity analysis by penalizing deviations from the baseline prior using a relative entropy or Kullback-Leibler measure of divergence. This divergence is well known to have convenient mathematical and conceptual implications. We implement this sensitivity analysis by converting our one-person maximization problem into a two-player game where the sensitivity analysis is conducted via minimization. This delivers a form of ambiguity aversion consistent with two alternative representations of ambiguity aversion: smooth ambiguity and variational preferences.

Our model is dynamic and Markovian. As such, it could be formulated from the vantage point of an initial period or recursively. Even in the absence of parameter uncertainty, we find it convenient computationally to solve it from the perspective of an initial date. This converts the optimization problem into a "static problem." We adopt this same static perspective to explore the consequences of uncertainty. In contrast to the single-agent decision theory, this has conceptual implications beyond just computational considerations as the minimization is performed as well at the initial date. In effect, we treat the two-player formulation as a static max-min game where, as we noted, the minimizing player is used as a formal device to explore the sensitivity to changes in the distribution over parameters used in optimization.

We use a regression method to quantify productivity parameters because we do not have direct evidence for each of the site-specific productivities. Our productivity data are available at different resolutions than the sites within our model. Moreover, we need to fill in missing observations as we do not have direct measures of agricultural productivity for some regions in the Amazon. Given these information limitations, we use attributes as right-hand side variables in regressions and feed in site-specific attributes. Moreover, we used the regression approach to fill in the missing observations. Since the dependent variables in the regressions are expressed in terms of logarithms, we exponentiate the predictions implied by regressions. We include the random effects in this computation.

We implement this approach as follows. For each site, we consider the parameter pair (γ^i, θ^i) for $i = 1, 2, ..., I$. The full parameter vector including all sites is thus of dimension $2 \times I$. We use a regression approach to construct baseline estimates of the site-specific productivities given site attributes. Observations on agricultural production are at the level of municipalities and a site may intersect multiple municipalities. We write M^i for the set of municipalities that intersect site i. We construct site-specific productivities using

$$\begin{bmatrix} \gamma^i \\ \theta^i \end{bmatrix} = \begin{bmatrix} \exp\left(\beta_\gamma \cdot R^i_\gamma + \nu^i_\gamma\right) \\ \frac{1}{P^A_{2017}} \sum_{m \in M^i} w^i_m \exp\left(\beta_\theta \cdot R^m_\theta + \nu^m_\theta\right) \end{bmatrix}. \tag{7}$$

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

using coefficients $(\beta_\gamma, \nu_\gamma^i)$ and $(\beta_\theta, \nu_\theta^m)$ from two regressions as inputs. The left side variables for these regressions are the logarithm of CO_2e and the logarithm of the slaughter value, respectively, as regressands. Here, w_m^i is the area-based weight of the importance of municipality m on site i, and P_{2017}^A is the price of cattle in 2017.

In the first regression, R_γ^i is a vector of geographical variables used to construct baseline estimates for the carbon-absorption productivity for the site i. Similarly, in the second equation, R_θ^m is a vector of such variables used to construct baseline estimates for the value of agricultural output per hectare for a municipality. In turn, the ν's are vectors of random effects. The ν_γ^i's are assumed to be equal across sites that belong in the same coarser partition of 78-site, and ν_θ^m's will be assumed to be equal across municipalities belonging to the same Amazonian water-basin. Details of regression findings are reported in Appendix D.3 along with a Bayesian counterpart to the R^2's reported as posterior densities. Although both sets of regressors are informative for both equations, the R^2 fit is better for the first equation used to make inferences about the γ^i's than for the second equation.

Construct the composite regression parameter vector,

$$\rho' \stackrel{\text{def}}{=} (\beta_\gamma{}', \nu_\gamma{}', \eta_\gamma, \zeta_\gamma, \beta_\theta{}', \nu_\theta{}', \eta_\theta, \zeta_\theta) \tag{8}$$

where the η's are the precisions for the regression errors and the ζ's are the precisions for the random coefficients. Uncertainty about ρ induces uncertainty about the site-specific productivities, (γ^i, θ^i) for $i = 1, 2, ..., I$ via equation (7), which we write abstractly as:[20]

$$\varphi = \Phi(\rho). \tag{9}$$

The underlying dimension of the uncertainty is given by the sum of the number of the unknown regression parameters, the number of distinct random-effects, and the four precision parameters that determine residual and random-effects variances for each regression equation. We choose a parametrization such that this sum is substantially less than $2 \times I$. This reduction turns out to be important for implementation. With this in mind, we let $d\hat{\pi}(\rho)$ be the baseline posterior (the "prior" used in solving the planner problem) over the parameter space of ρ's. This posterior in turn implies a baseline posterior $d\pi(\varphi)$ over the space of possible productivity pairs for each of the locations induced by $\Phi(\rho)$ and the change of variables formula applied to probability measures. Rather than inferring the $d\pi(\varphi)$ through this change of measure computation, we achieve the same robustness adjustments for land allocation decisions by performing the minimization over probabilities $\hat{g}(\rho)d\hat{\pi}(\rho)$ for alternative densities $\hat{g}$'s in the space of ρ's. This approach requires that we substitute for φ in the objective, that is $f(d, \varphi) = f[d, \Phi(\rho)]$.

For the benchmark Bayesian implementation, we use the familiar conditional normal/inverse gamma prior for the unknown parameters in the two regressions. Specifically, the prior distribution for the unknown regression coefficients including the random effect is presumed to be normally distributed conditioned on the random effects and regression-error variances. The prior for the precision coefficients are posited to have

[20]For notational simplicity, we leave the data dependence as implicit.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

improper gamma distributions.[21] Our planner takes the implied posterior distribution, after conditioning on data, for the productivities as inputs into the decision problem and makes a robustness adjustment.[22]

Our measure of divergence used to explore the sensitivity restricts the alternative probabilities to be absolutely continuous with respect to the baseline distribution. With this restriction, it suffices to focus on alternative distributions to the baseline specification for the regression coefficients. This simplifies substantially the numerical computations.

The preferences described by the minimization problem (6) are recognizable as a special case of what are called variational preferences. The minimization problem has a well-known quasi-analytical solution:

$$\hat{g}^*(\rho) = \frac{\exp\left(-\frac{1}{\xi} f\left[d, \Phi(\rho)\right]\right)}{\int_{\mathcal{R}} \exp\left(-\frac{1}{\xi} f\left[d, \Phi(\rho)\right]\right) d\hat{\pi}(\rho)} \tag{10}$$

with a minimized objective:

$$-\xi \log \int_{\mathcal{R}} \exp\left(-\frac{1}{\xi} f\left[d, \Phi(\rho)\right]\right) d\hat{\pi}(\rho). \tag{11}$$

The minimizing $\hat{g}$ lower discounted utilities. The magnitude of ξ determines the strength of this tilt. We shall refer to the limiting $\xi = \infty$ case as *ambiguity neutrality*. For this limit, the decision problem uses the familiar expected utility objective:

$$\max_{d} \int_{\mathcal{R}} f\left[d, \Phi(\rho)\right] d\hat{\pi}(\rho).$$

Large values of ξ lead to outcomes that approximate this limiting case.

Remark 4.1. *The minimized objective given by* (11) *is a special case of a smooth ambiguity objective, first suggested by Klibanoff, Marinacci and Mukerji (2005). They deduced a rationale for an ambiguity adjustment represented using a concave function distinct from the one used to express risk aversion. While they take such a concave adjustment to be a starting point, we deduce a logarithmic-exponential representation from a starting point motivated by distributional robustness. Thus, their axiomatic motivation is different from the distributional robustness that interests us.*

Given the parameter ambiguity adjustment, the implied decision problem is a two-player zero-sum game. In our problem formulation, we may change the order of maximization and minimization in Problem 3.1, which is of interest in its own right.

Problem 4.2.

$$\min_{\hat{g}\geq 0, \int \hat{g} d\hat{\pi}=1} \max_{d\in\mathcal{D}} \int_{\mathcal{R}} f\left[d, \Phi(\rho)\right] \hat{g}(\rho) d\hat{\pi}(\rho) + \xi \int \hat{g}(\rho) \log \hat{g}(\rho) d\hat{\pi}(\rho).$$

Formally, we may invoke the Min-max Theorem and claim that the objective for Problems 3.1 and 4.2 will be same and the minimizing g evaluated in the maximized d for Problem 3.1 will agree with the minimizing g of Problem 4.2. Similarly, the optimized decision processes will agree.

[21] See Appendix D for details.

[22] Our static formulation of the two-player decision problem does not allow for the possibility of dynamic learning going forward. We abstract from this for computational tractability.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Consider the inner maximization for Problem 4.2:

$$\max_{\boldsymbol{d}} \int_{\mathcal{R}} f\left[d, \Phi(\rho)\right] \hat{g}(\rho) d\hat{\pi}(\rho)$$

where we are free to drop the relative entropy penalty, as it does not depend on the decision process d. Provided that this inner problem has a solution for the outer g minimization, 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. Following a common practice for robust Bayesian methods, we find it revealing to explore alternative choices of ξ and deduce their implications for implied worst-case probabilities.

5 Productivity measurement

We used two different spatial resolutions for the results that we report. At the most detailed level, we consider a regular grid of the Amazon region with pixels of 30 m $\times$ 30 m resolution from MapBiomas (Souza Jr et al., 2020). We then aggregate pixels to form 1887 sites that are 67.5 km $\times$ 67.5km. Many of these sites do not overlap the Amazon biome. We discard these and 20 others with less than 3% of their area in the Amazon biome. This reduced our number to 1043 sites.

For tractability reasons, when we consider the stochastic evolution of agricultural prices, we use a less refined grid of 130 sites that are approximately 270 km $\times$ 270 km. We obtain 78 sites after dropping sites that do not overlap the Amazon biome at all and four additional sites with less than 3% in the Amazon biome. We also use this 78-site partition to define the group of coarse sites for which random effects on carbon productivity, ν_{γ}^{i}, are equal.

As discussed in section 4, we construct site-specific productivity estimates from the output of regression equations. See formula (7). Appendix A describes in detail all data used for these regressions. What follows is a summary of the evidence we have gathered.

For the agricultural productivity regression, we use 2017 as the reference year, as it is the most recent Agricultural Census in Brazil. For regressands, this census provides information on the value of cattle sold for slaughter per hectare of pasture land at the level of a municipality. As regressors, we use geographical variables as stipulated in Appendix A. The census provides observations on the value of cattle sold for slaughter per hectare of pasture land for 466 of the 540 municipalities that intersect the biome. In addition to dimension-reduction of the estimates, the regression allows us to obtain data for these missing municipalities.

To measure the productivity of carbon sequestration, the γ^i's, we first use data from MapBiomas[23] to select pixels that can be considered primary forests. For these pixels, we used 2017 data from ESA Biomass[24] to obtain carbon per hectare. We then calculate average productivities γ^i for each site i. Analogously to the procedure we used for the θ^i's, we run a regression using the logarithms of the γ^i as regressands and

[23] Web address: www.mapbiomas.org (Collection 5).
[24] See Santoro and Cartus (2021).

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

geographical variables as regressors to obtain estimates.[25]

As implied by equation (2), the parameters γ_i govern the maximum carbon accumulation for site i and the parameter α the dynamic responses of the carbon accumulation path when previously deforested lands are restored. Because this modeling input is central to our quantitative analysis, we assess the plausibility of the implied reforestation dynamics by comparing its implications to available data.

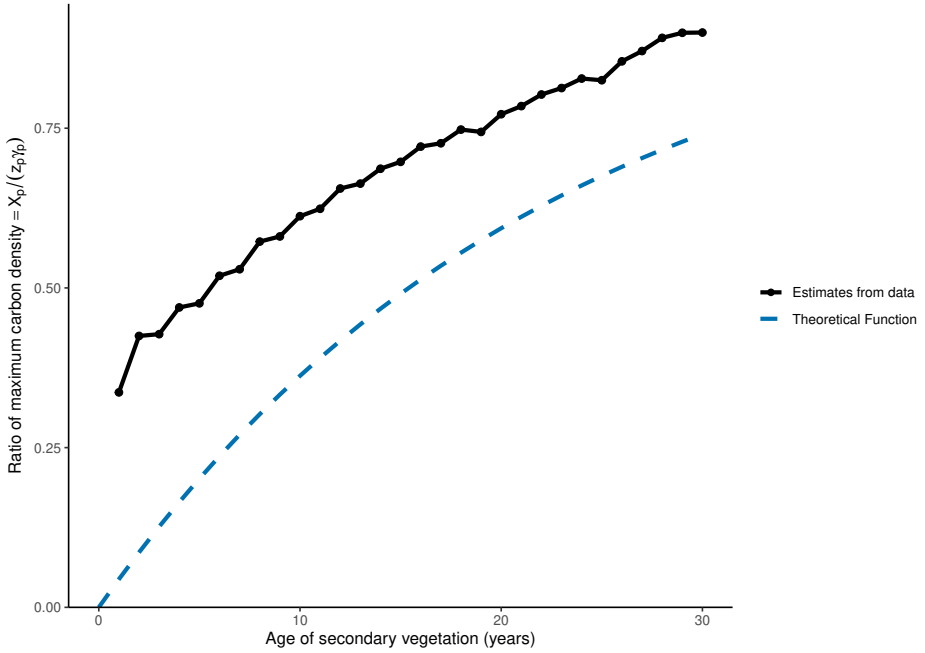


Figure 2: Percentage of maximum carbon density: actual vs. theoretical. The black solid curve gives the an empirically-based construction of the carbon recovery fraction as a function of the recovery time. The blue dashed curve plots $1 - \exp(-\alpha t)$.

Our approach measures observed carbon capture over the large area of secondary forests observed in the Brazilian Amazon.[26] The presence of extensive deforested areas that were subsequently abandoned and underwent natural regeneration provides a natural validation experiment. We identify all pixels p of size 0.001078° (approximately 120 m $\times$ 120 m at the equator) that, as of 2017, were classified as undergoing reforestation. The resulting set covers 21.3 million hectares.[27]. For each 120 m pixel p, $X_p/\bar{z}_p$ measures the carbon capture by secondary forest scaled by the available land. We divide this by γ_p, which is the average

[25]This procedure delivers estimates aligned with those presented in Liang et al. (2025) and Fesenmyer et al. (2025).

[26]For each pixel classified as secondary vegetation, MapBiomas defines age as the number of years since the pixel first exhibits a land-cover signature consistent with native vegetation in regeneration in that biome and region. Age thus measures time since the first year of MapBiomas secondary-vegetation classification, not time since ecological regeneration began. For the Amazon biome, the lag between ecological onset and this first classification is approximately five years. We therefore add five years to the MapBiomas age when mapping secondary-forest age into carbon sequestration over the restoration path. We also consider alternative curves, assuming a 3-year and a 7-year gap, to ensure robustness.

[27]This calculation treats these pixels as naturally regenerating. Indeed, assisted restoration accounts for less than 1% of the 2017 secondary-forest area considered here (Alliance for Restoration in the Amazon (2020)).

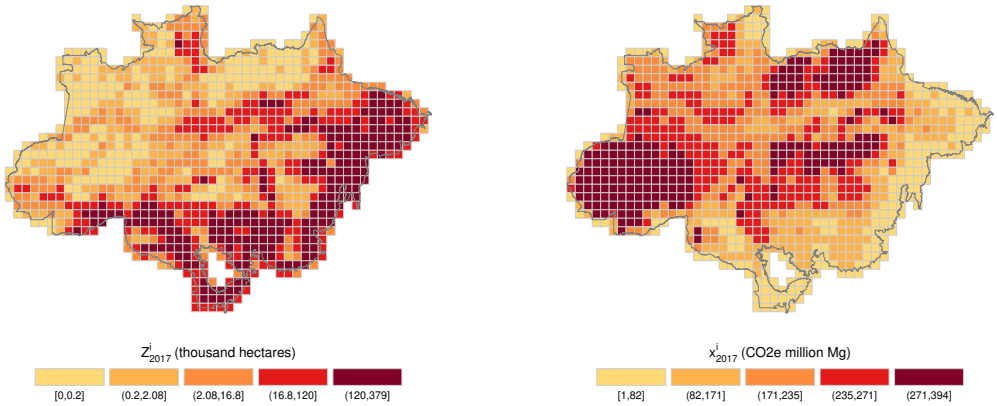
This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

value of γ's within the pixel and regress this onto dummy variables for the age of the secondary forest[28]

$$\frac{X_p}{\bar{z}_p \gamma_p} = \sum_t \beta_t \mathbf{1}_{\{t(p)=t\}} + \epsilon_p.$$

This quantity measures the fraction of the maximum possible carbon content in a pixel that has been attained by 2017, according to our model. Given the scaling of the left-hand side variable in the regression, including that of the γ_p's, the relevant comparison for the right-hand side regression curve is with the $1 - \exp(-\alpha t)$, consistent with equation (2). This comparison indicates that our assumed carbon accumulation under natural reforestation is not overly optimistic, although we understate the ability of newly regenerated forests to sequester carbon. [29]

Figure 3 shows the initial land allocated to agriculture and the initial stock of absorbed carbon across the 1043-grid sites. Figure 4 shows how the estimated carbon sequestration parameter γ^i and agricultural productivity parameter θ^i varies across the different sites. The correlation between θ^i and γ^i is -0.27 for the finer resolution. Thus, while agricultural productivity and carbon absorption capacity are negatively correlated, this relationship is imperfect.



z_0^i x_0^i

Figure 3: Initial values for agricultural area (z_0^i) and carbon stock (x_0^i)

[28]The regression $R^2 = .66$.

[29]Poorter et al. (2021) uses specialized data from 77 sites worldwide to predict the recovery of a multidimensional set of tropical forest attributes using cross-sectional analysis. While their use of detailed information on a relatively small area allows for the analysis of characteristics our data does not—such as species composition—we opted instead to measure actual carbon capture across the vast area of secondary forests observed in the Brazilian Amazon. For tractability and measurement limitations, our model does not to capture the full multi-dimensional dynamics of forest regeneration.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

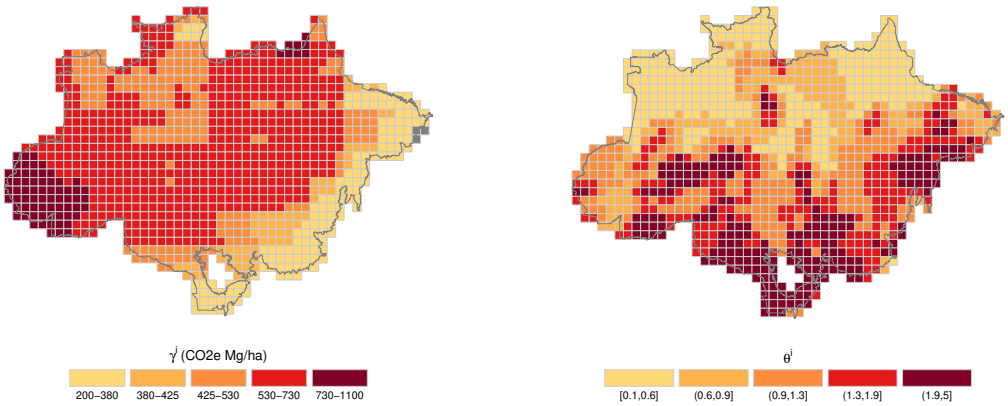


Figure 4: Carbon sequestration parameters and Agricultural productivity heterogeneity

6 Solving the maximization problem

To achieve the degree of economic and spatial richness needed, we use numerical methods to obtain model solutions. Given the number of locations that we use (1043 or 78), we necessarily have a large number of 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 imposes penalties on logarithms of variables that are constrained to be non-negative. 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.

6.1 Solution with parameter ambiguity neutrality

In the absence of stochastic prices and with ambiguity neutrality, we solve a static optimization problem over all possible trajectories for the next 200 years for 1043 sites. This problem is deterministic, evaluated at the *ex-ante* average parameters using the baseline distributions. Given obvious uniform bounds on possible utility flows and discount rates of at least two percent, the resulting trajectories give reasonable approximations for the infinite horizon optima.

6.2 Solution with parameter ambiguity aversion

For solving the robustly optimal problem numerically, we take an iterative approach, supported by the Min-Max Theorem. Specifically, we proceed as follows:

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

ii) For a given d, we solve the minimization problem with the relative entropy penalization to obtain a new candidate for g^*.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

iii) We repeat the steps until we achieve convergence.

Our computations in step ii) will take advantage of the quasi-analytical formula for g^* given in (10). We take a decision process as given, d, and evaluate the discounted objective to obtain the numerator for (10). As we show in the Appendix E, we may imitate a standard numerical Bayesian posterior calculation via simulation based on Hamiltonian dynamics, which is often more computationally efficient for high-dimensional problems than the familiar Metropolis-Hastings method. See, for instance, Neal et al. (2011) and Carpenter et al. (2017), with software support given by Stan Development Team (2023).[30] See Appendix E.1 for more details.

6.3 Solution with price stochasticity

To account for stochastic prices in a tractable way, we construct an approximate Markov chain with two states $P^a_\ell < P^a_h$ and transitions that match the empirical transitions from monthly data.[31] When considering a stochastic evolution of prices, we find "Modified Predictive Control" (MPC) methods (*e.g.* Scokaert and Rawlings (1998), Bemporad et al. (2002), and Thangavel et al. (2018)) to be particularly suitable for solving our planner's problem. Our MPC approximation is implemented as follows.[32] Starting from the current period, denoted as date zero, 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. Although the cattle price distribution follows a Markov chain, to simplify our computations, we set the prices in periods $\tau+1,\dots T$ equal to the value that prevails at τ. Prior to date $\tau+1$, we confront randomness in this problem by imposing appropriate "measurability" restrictions on the controls as functions of potentially realized states. We then apply the interior point method to find the optimal trajectory at date zero given P^a_0. We keep the optimal date $t=1$ states computed at time $t=0$ and repeat. That is, we consider the problem starting at $t=1$ with the new state vector and divide the future into two segments: an uncertainty segment of length τ and a remaining period of $T-\tau-1$. This step will determine an optimal state at period 2. We continue this procedure to produce the optimal state at periods 3, 4, ..., T supported by the corresponding optimal controls.

In practice, the dimensionality of the stochastic problem increases geometrically as a function of the uncertainty horizon, τ. Consequently, 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 $\tau=1,\dots,\tau^*$. In our coarse grid with 78 sites and price randomness, we chose $\tau^*=3$.

[30] From a mathematical standpoint, this calculation is equivalent to computing a Bayesian posterior where $-\frac{1}{\xi}f(d,\Phi(\rho))$ plays the role of a log-likelihood function and π plays the role of a prior.

[31] See Appendix A.9 for details.

[32] Related computational approaches have been proposed by Cai, Judd and Steinbuks (2017) and Cai and Judd (2023).

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

6.4 Solution with stochasticity and robustness

We extend the MPC method to include robustness to the misspecification of the transition probabilities for the agricultural price process. This form of uncertainty is distinct from the parameter uncertainty that we confronted in our use of productivity measurements. First, it is explicitly dynamic, and second it is not based on a particular parameterization of the Markov price dynamics. Instead of exploring *ex ante* uncertainty about a Bayesian prior used for measurement, we entertain potential departures in the baseline transition dynamics period-by-period. Since the potential departures in future periods can be different from those in the past, learning is off the table. The relevant measure of divergence used to penalize more extreme departures is based on the transition dynamics and not on divergence of a prior over uncertain parameters.

Anderson, Hansen and Sargent (2003) suggest a recursive way to include such an adjustment using an intertemporal notion of relative entropy for continuous-time jump processes. Conditioned on each state, there is a jump intensity for moving to a different price state that could be misspecified. (The jump probability is the intensity times the time interval.) We then make the robustness adjustment to the discounted conditional expected utility applied to the discrete-time approximation. More specifically, we start with a hypothetical decision rule contingent on the different state realizations over the uncertainty horizon. Very similar to the case with parameter ambiguity, there is an exponential tilting formula for adjusting the transition probabilities that we exploit. We distinguish the penalization parameter for the potential Markov chain misspecification with the notation $\widehat{\xi}$ in contrast to the parameter governing productivity ambiguity concerns, which we denoted by ξ.

We then make a robustness adjustment to the objective over the uncertainty horizon using this exponential tilting formula period-by-period working backwards and discounting. The outcome of this computation gives the objective that we maximize. As with ambiguity aversion, we iterate between maximization and minimization. We use the minimizing initial period transition distribution for the uncertainty horizon step of the computation outcomes to form uncertainty-adjusted probabilities for purposes of valuation.[33]

We allow for different choices of the value of penalty parameters for this robust adjustment and for the parameter ambiguity. Instead of making adjustments to probabilistic statements over time-invariant parameters, we now make robust adjustments to a baseline specification of the transition matrix over states taking the current price state as observable. We do this to allow for model misspecification, in contrast to parameter misspecification. Although the baseline transition probability matrix is time invariant, we do not impose this same invariance when we explore misspecification.

7 Results

In this section, we report our quantitative findings. We start by constructing a benchmark "business-as-usual" set of results. We accomplish this by deducing an implied social price of carbon that supports current aggregate implications. Then, in succession, we study i) deterministic solutions for higher social prices; ii)

[33] See appendix C.1 for details.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

solutions with parameter uncertainty; and iii) stochastic solutions with price randomness.

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 an interval $[\underline{t}, \bar{t}]$ and then select a time-invariant price for emissions, denoted by P^{ee}, to match the aggregate deforestation predicted by the model at a final observation period $\bar{t}$. We let $(X^o_{\underline{t}}, Z^o_{\underline{t}})$ denote the initial observed state vector. We also input the realized history of agricultural prices $\{P^a_t : \underline{t} \leq t \leq \bar{t}\}$. We then compute the optimal trajectory for the state variables implied by our model for alternative choices of P^e and find the P^{ee} that matches $\sum_{i=1}^{I}(Z^i_{\bar{t}} - Z^i_{\underline{t}})$ to the observed value of the aggregate deforestation in the period $[\underline{t}, \bar{t}]$.[34]

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 conservation projects in the Amazon, using money contributed primarily by Norway.

The business-as-usual price, P^{ee}, depends on the model specifications, and in particular on the ambiguity aversion of the planner. In what follows, we will consider solutions to the optimization problem starting in 2017 and a discount rate of two percent. As we document in Table 1, enhanced uncertainty concerns in the productivity parameters lead to a smaller business-as-usual price for reasons that will become clear in our subsequent discussion.

We will explore implications when the planner uses $P^e = P^{ee} + b$ for $b = 0, 10, 15, 20,$ and 25 where b represents transfers per ton of *net* captured emissions to the planner. Specifically, when net emissions total E tons of CO_2, the planner receives a transfer of bE. Note that even $b = 25$ corresponds to a social price that is low when compared to the prevailing price of emissions in some regions of the globe.

The forces that lead to changes in shadow prices have a direct and partially offsetting effect on deforestation. In the tables that follow, we report results in terms of the transfer payment b, which is an external input into our analysis. Notice that a given b implies a different social cost carbon under ambiguity aversion than under ambiguity neutrality. When our results are expressed in terms of b, there are offsetting impacts of ambiguity aversion what would be absent had we expressed results in terms of alternative social costs of carbon. We find b to be the economically relevant construct as it is tied directy to potential transfer payments.

Specifically, the implied robustly optimal trajectories for each choice of b will be less sensitive to the particular model specification. Essentially, the same argument applies to changes in the subjective discount rate.[35]

[34] We obtain a similar value if instead we minimize the norm of a vector with two components: The first is the percentage deviation of predicted forest-wide carbon capture from the observed carbon capture at $\bar{t}$; the second, the percentage deviation of predicted forest-wide deforestation from the observed deforestation at $\bar{t}$.

[35] Since emissions are a low duration asset relative to cattle, a larger discount rate implies less deforesting, thus lowering P^{ee}, approximating future trajectories for a given b. In fact, our simulations show that future trajectories do not change much for each b, when we move from two to three percent. See Appendix A.7.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

7.2 Results for case without stochasticity or ambiguity aversion

In this section, we discuss the results for a model with a constant price for cattle that equals the average price in the stationary distribution for the estimated 2-state Markov chain (41.1). We first discuss results for 1043 sites, and then include results for 78 sites for comparison.

As Figure 5 shows, with "business-as-usual" ($P^e = P^{ee} = 6.8$), the optimal choice involves an increase in the agricultural area from 15% to around 20% of the biome. This increase may actually cause sufficient deforestation for the hydrological cycle of the Amazon to degrade to the point of being unable to support rainforest ecosystems (Lovejoy and Nobre (2018)). The predicted trajectories are much different with an additional per ton payment to the planner of \$15 or \$25. Figure 6 reports the trajectories over time of the transfer payments for $b = 15$ and $b = 25$. The peak payments occur after about 10 years for both values of b. As expected the transfer payments for $b = 25$ are much larger than the corresponding payments for $b = 15$.[36]

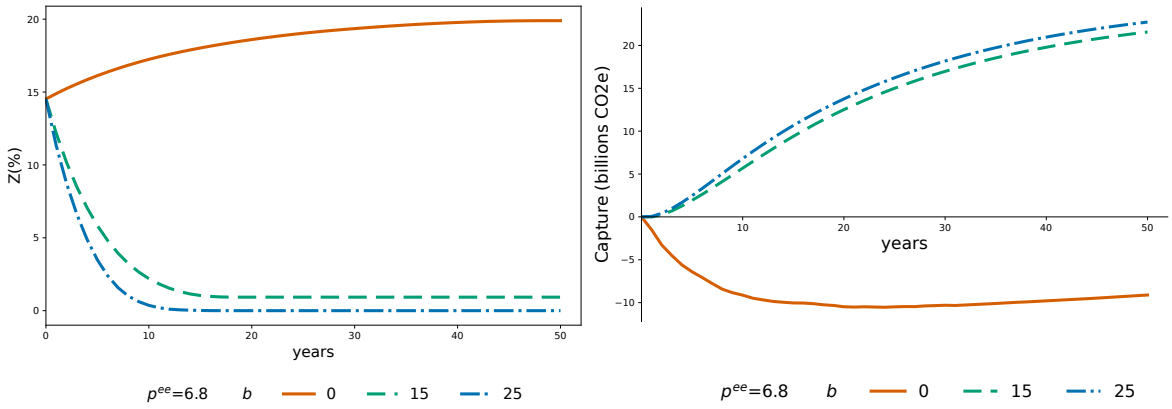


Figure 5: Agricultural area and cumulative CO_2 capture.

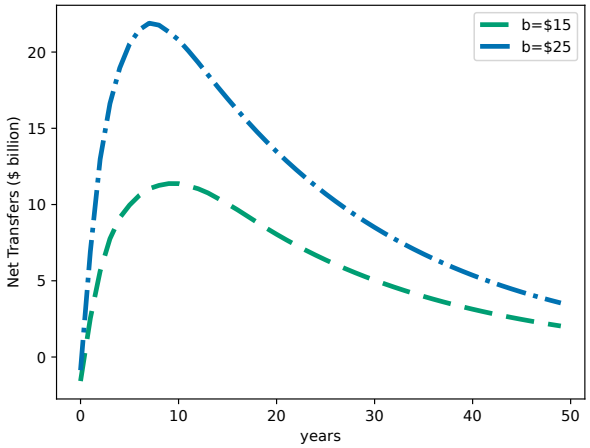


Figure 6: Evolution of transfer payments for two choices of b.

[36]Since reforestation is central to our quantitative analysis, we find that the Brazilian Amazon can be a viable nature-based contributor to carbon sequestration at substantially smaller social costs of carbon than Araújo, Costa and Sant'Anna (2025).

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

As we show, these transfer payments result in a substantial decrease in agricultural area and a corresponding increase in forested area. The rows of Table 2 give the discounted value to the planner of a commitment to receive $\$b$ of net transfers for each ton captured of CO_2e, when P^a is the stationary price. It also gives a decomposition of this present value to interpretable components. Among these components, "forest services" are measured at the implied Brazilian shadow price for business-as-usual. The net transfers to Brazil are reported separately. Even transfers of \$10 per ton are enough to compensate the losses of agricultural output, but the largest contributor to the gains is the increase in net transfers. The larger transfer of \$25 per ton of net captured CO_2e, almost doubles the value for the planner - a net gain of \$264 billion. This net gain is composed of a loss of \$360 billion in the value of cattle output,[37] which is more than compensated by \$417 billion in transfers and \$221 billion in forest services. Adjustment costs are only a small part of the story.

Table 3 displays the total effect of transfers per ton of net CO_2e captured in years 15 and 30. For the business-as-usual carbon price, the planner chooses deforestation that induces carbon emissions of about 12 and 15 billion tons per year in 15 and 30 years, respectively. This table uses this baseline to feature the "effective cost." We calculated this as the ratio of discounted net transfers to the difference between the net carbon captured and the corresponding baseline value when $b = 0$. With transfers of, say, \$15/ton, optimal management induce capture of about 8.8 billion tons by year 15 and an additional 7.5 billion tons by year 30. The effective costs are \$6.3 and \$7.9 , considerably less than the per ton subsidies captured by the b's. With transfers of \$25/ton, there are modest increases in the captured carbon with effective prices that are almost double, but still approximately 55% of the transfer payments per ton. Thus, the results in Table 3 illustrate the gains from trade in instituting a contract that pays Brazil per net ton of CO_2e captured.

Figure 7 exhibits the initial distribution of land allocation over 30 years for b = \$0, \$10, and \$25. It shows that for the case of transfers that exceed \$10 per ton of net emissions, the area of the biome that is occupied by cattle farming after 30 years would be substantially reduced in comparison to the 2017 allocation. This is in sharp contrast to what transpires in the $b = 0$ business-as-usual specification in which agricultural production becomes quite intense in the lower right sites.

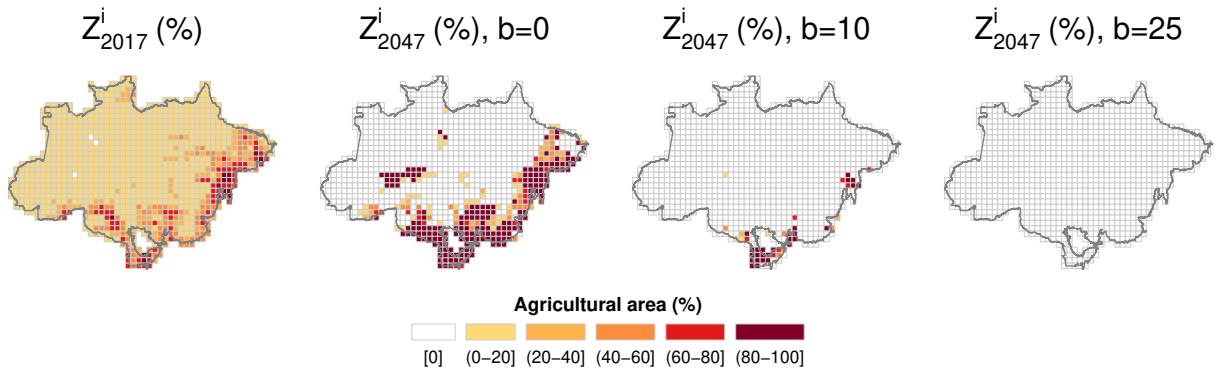


Figure 7: Agricultural area changes after 30 years.

[37] Recall, however, that we use a measure of full output as value added. Thus, we have exaggerated the loss of agricultural output.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Figure 8 provides a more complete spatial dynamic characterization for transfers of \$15/ton. In the optimal solution, much of the change in land occupation occurs within the first 15 years.

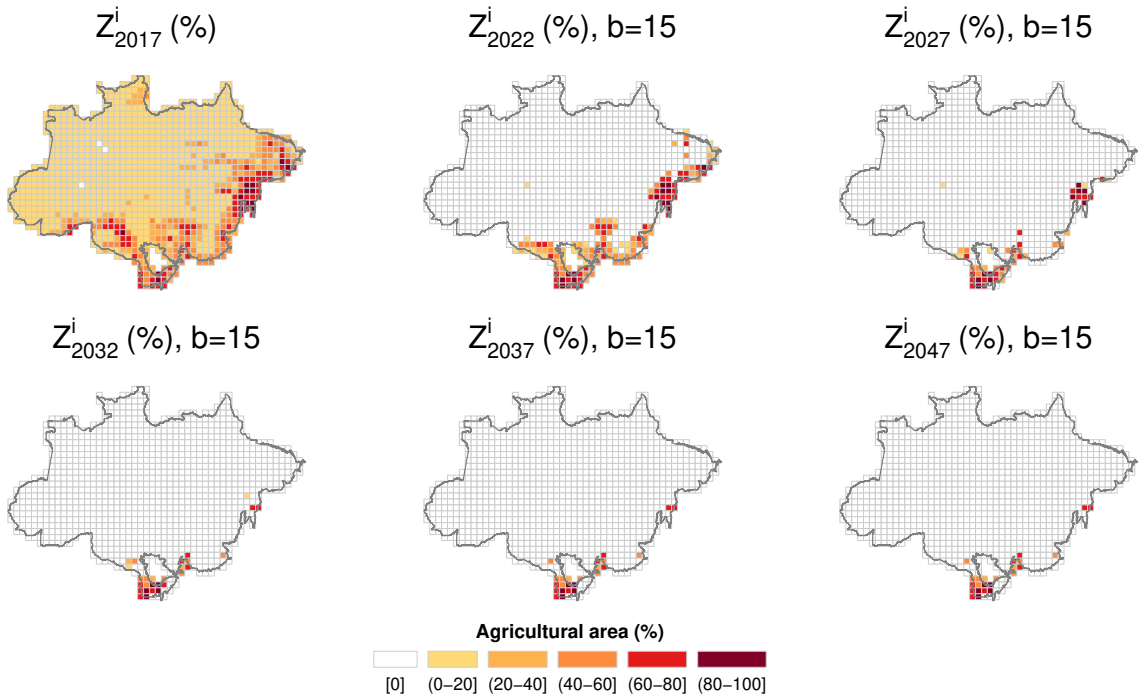


Figure 8: Agricultural area evolution over time

7.3 Results with robustness to parameter ambiguity

In this section, we present results when the planner is uncertain about the cattle productivity and CO_2e capture potential. Although it is revealing to perform robustness calculations for several values of ξ, here we report results only for $\xi = \infty$ and $\xi = 1$. We refer to the former as 'ambiguity neutral' and the latter as 'ambiguity averse.' We report results for other values of ξ in the appendix B.4. The implied ambiguity adjustments to the probabilities help us gauge the plausibility of different values of ξ. The calculated shadow price, as reported in Table 1, is \$6.8/ton under ambiguity neutrality and a considerably lower value of \$4.8/ton under ambiguity aversion. The shadow price reduction under ambiguity aversion compensates for the slower destruction of the almost virgin forest when there is ambiguity in site's productivity.

When looking across all sites, the ambiguity adjustments to the productivity parameter distributions are very heterogeneous; and for some sites there is very little difference between the two distributions. Under the business-as-usual benchmark, the adjustments are substantially more pronounced for the θ distributions than for the γ distributions. In other words, it is the uncertainty about agricultural productivity that the social planner finds most concerning since the planner is not incentivized to preserve much of the rainforest without transfer payments.[38] The uncertainty adjustment to the probabilities are substantially different when $b = 15$.

[38]Notice that the uncertainty adjusted γ distributions actually shift slightly to the right when $b = \$0$.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Now the adjustments are more substantial for the γ probability densities because reforestation becomes a prominent ambition for the planner.

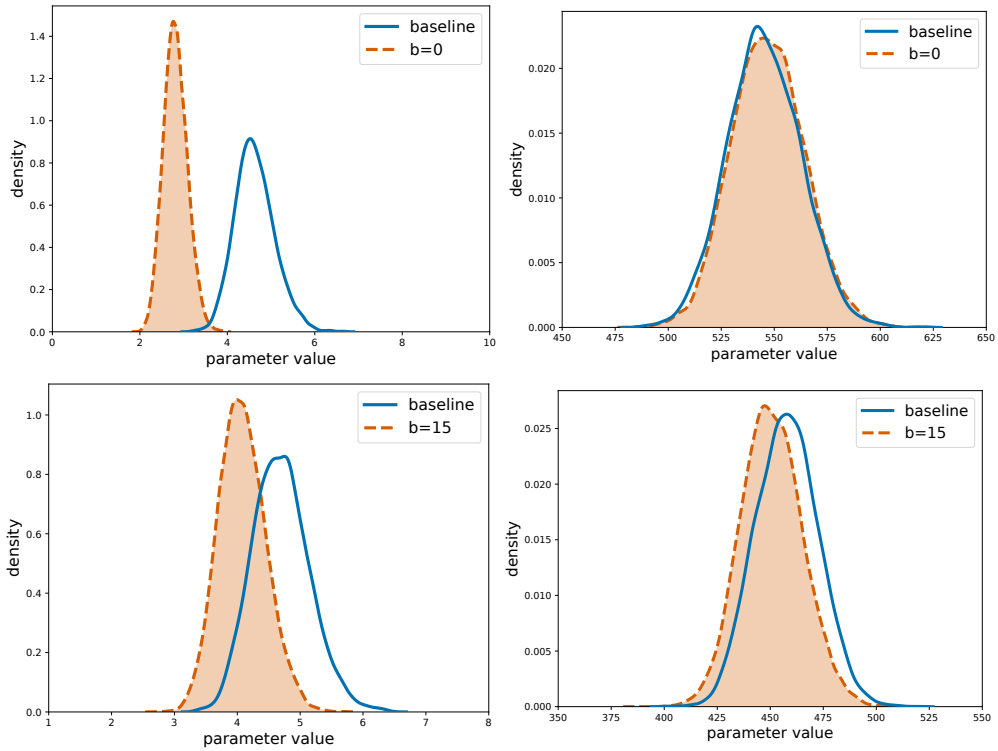


Figure 9: Ambiguity-adjusted densities for four sites. The left plots report sites with the largest entropies for θ and right plots report sites with the largest entropies for γ. The upper plots are $b = 0$ and the lower plots for $b = 15$.

Figure 9 illustrates these impacts by showing the baseline and ambiguity-adjusted densities for parameters γ and θ when $b = 0$ and $b = 15$. For the sake of illustration, we display results for the four sites with the largest relative entropy divergence between the baseline probability distributions and the ambiguity-adjusted counterparts.[39] There are four because of the two productivity parameters and the two values of b. The upper left and lower right plots show the notable density shifts toward more cautious productivity assessments depending on the choice of b.

These distributional shifts support robust adjustments to the land allocation decisions. They alter both the timing and the magnitude of the land allocations, as we now illustrate.

Cross-sectionally, some sites are deforested, and some sites are reforested by the planner in the absence of external transfer payments. Recall that $U_t^i > 0$ is when site i is being deforested, and $V_t^i > 0$ when the site i is being reforested. Only one of these can be strictly positive at any date t.

Figure 10 presents a histogram for the number of years it takes for one of the two controls to be maximal. Sites that do not reach their maxima within fifty years are not included in these histograms. This figure shows delays in deforestation or acceleration of reforestation for all sites under ambiguity aversion, in the absence of

[39] See Figure 17 in Appendix B.3 for a map of the relative entropy divergence of each site.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

transfer payments. It compares what happens when $b = 0$ under ambiguity neutrality and ambiguity aversion, using the *same* business-as-usual carbon emissions price. The ambiguity aversion leads to a modest shift to the right in this distribution. The overall quantity impacts are substantially different as shown in Figure 11. Under ambiguity neutrality, there is a substantial increase in the agricultural land over time, in stark contrast to the outcome under ambiguity aversion.

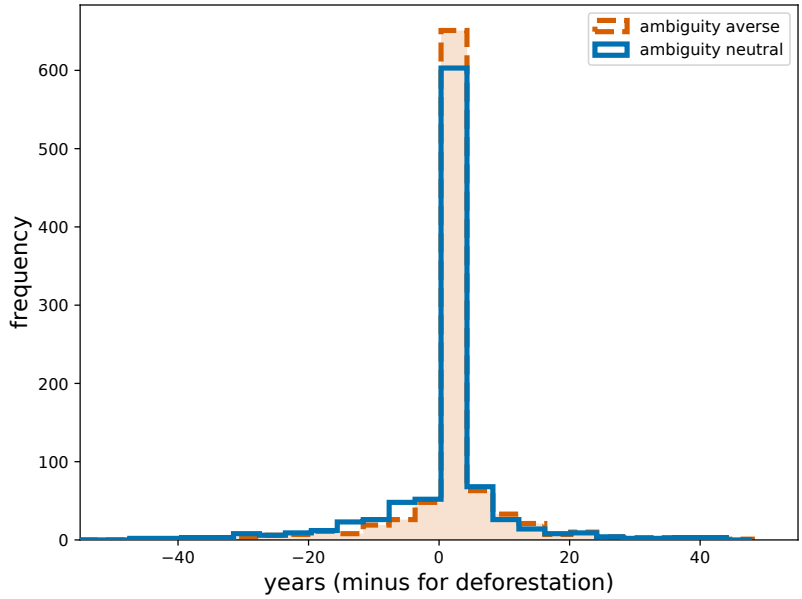


Figure 10: Histogram of the years in which one of the two controls is maximal, for $b = 0$ and common $P^{ee} = 6.8$. The controls remain at zero for fifty years in 10.1% of the sites under ambiguity neutral and in 10.2% of sites under ambiguity aversion.

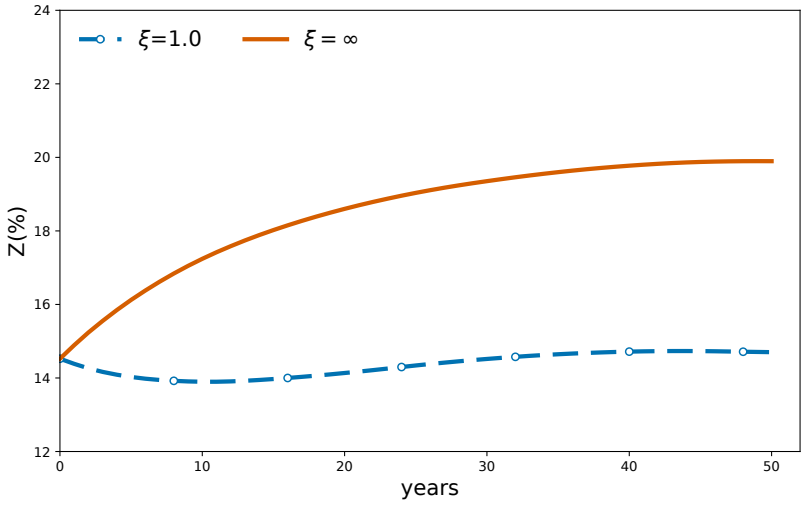
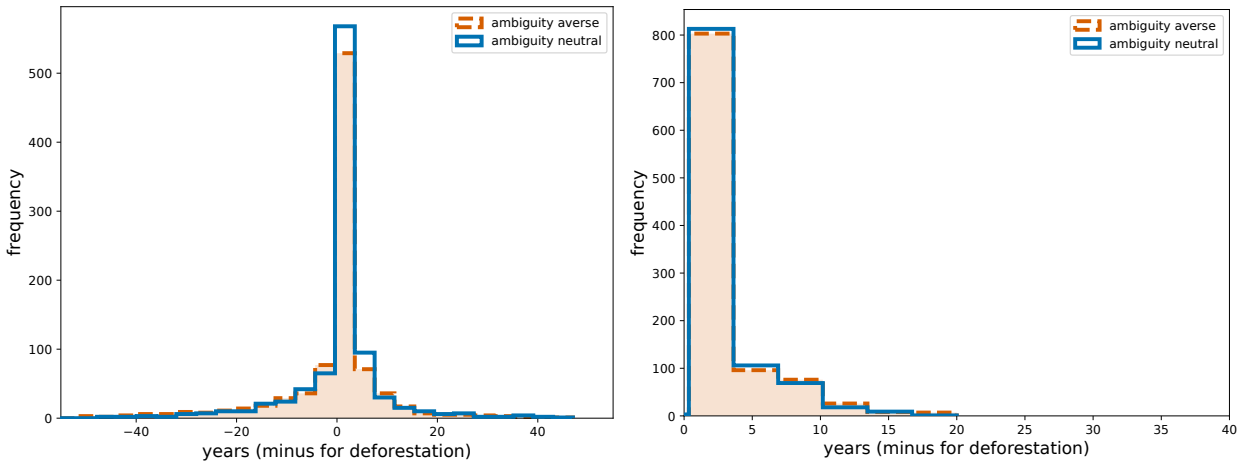


Figure 11: Evolution of agricultural area under ambiguity neutrality and ambiguity aversion for $b = 0$ at a common business-as-usual carbon price, $P^{ee} = 6.8$.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

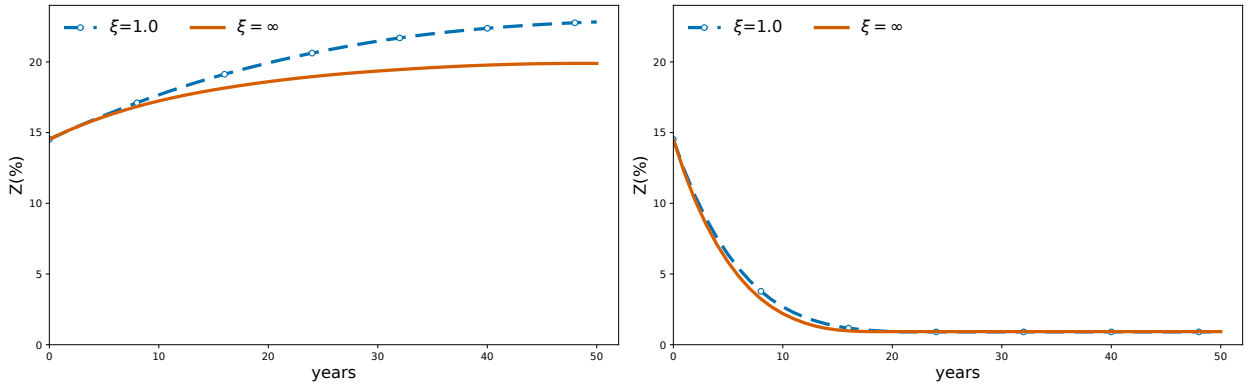
The previous results hold the business-as-usual price fixed as we introduce ambiguity aversion. It turns out that these quantity impacts are altered and somewhat muted when we take account of the impact of ambiguity aversion on the implied business-as-usual price. As we know from Table 1, this price decreases endogenously when we impose ambiguity aversion from 6.8 to 4.8. The impact on the spatial-dynamic land allocation is shown in the left panels of Figures 12 and 13. Although ambiguity aversion still alters some of the dates of the maximal responses, there is no longer a notable shift to the right under ambiguity aversion as in Figure 10. Instead, there is a little more dispersion in the distribution. The aggregate evolution of land allocated to agriculture now increases over time under ambiguity aversion in contrast to what is displayed in Figure 11, and it is actually somewhat higher.



(a) b=0 (b) b=15

Figure 12: Histogram of the years in which one of the two controls is maximal, $P^{ee} = 6.8$ for ambiguity neutral and $P^{ee} = 4.8$ for ambiguity averse. For $b = 0$, 10.1% of the sites remain unchanged under ambiguity neutrality, while 13.2% do so under ambiguity aversion. For $b = 15$, 2.0% of the sites do not change under ambiguity neutrality, and 2.1% remain unchanged under ambiguity aversion.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.



(a) b=0 (b) b=15

Figure 13: Evolution of agricultural area under ambiguity neutrality and ambiguity aversion, using the corresponding shadow prices.

The right panels of Figures 12 and 13 display the spatial-dynamic impacts when $b = 15$. According to the right panel of Figure 12 reforestation dominates with only limited sensitivity to the ambiguity aversion in this case. As shown in the right panel of Figure 13, while the allocation of land to agriculture diminishes substantially over time, there is no significant difference in the land allocations under ambiguity neutrality and aversion. The endogenous adjustment in the business-as-usual emissions price mutes all of the quantity responses.

Finally, we consider the present values under ambiguity aversion in comparison to ambiguity neutrality in Table 4. We report planner values and agricultural values. As should be expected, the ambiguity aversion induces smaller planner values since they are computed with uncertainty-adjusted probabilities. In contrast, the agricultural value decreases under ambiguity aversion for $b = 0, 10$ and 15, although it increases for $b > 15$. Recall that the uncertainty-adjusted probabilities provide a more conservative assessment of agricultural productivities that are prominent when b is low, inducing more conservative assessment of agricultural value. For other values of b, ambiguity aversion actually results in increases in agricultural value, although the overall magnitudes are small relative to the corresponding planner values. The reason for the increase is that, under ambiguity aversion, the planner makes a more cautious assessment of the ability of the Brazilian rainforest to absorb carbon, leading to a very small increase in the land allocated to agriculture. Notice that although for each b the presence of ambiguity lowers the value for the planner, the percentage gains from moving from business-as-usual to $b = 25$ are higher in the presence of ambiguity aversion.

Appendix B.4 displays some uncertainty-adjusted probabilities for $\xi = 0.5, 2$ that are counterparts to those in Figure 9. In addition, this appendix depicts evolutions for the agricultural land allocations that are versions of those in Figure 8 updated for ambiguity aversion.

Remark 7.1. *In contrast to land allocation process, Z, with parameter ambiguity, the state vector process, X, of captured carbon is disguised to the planner, because initial conditions and the dynamics of carbon captured depend on the value of γ. This has ramifications for policy since we presume transfer payments are based on carbon reduction. Under ambiguity aversion, our planner uses the ambiguity-adjusted probabilities*

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

to compute these payments. In ad hoc policy-making settings distinct from our fictitious planner formulation, one could imagine differences in perspective among providers and recipients of transfers, opening the door to explicit consideration of differences in their aversion to uncertainty.

7.4 Results with stochastic variation in agricultural prices

For our final set of results, we explore implications allowing for an explicit randomness in the agricultural price process. We generate these results using the MPC method described previously. To keep things tractable, we have the social planner assume a two-state Markov process for the price process. We obtained the inputs into this specification by estimating a hidden-state Markov process with Gaussian noise as we describe in Appendix A.9. This is a rather substantial reduction in the stochastic structure of agricultural prices, but it allows us to engage in an initial exploration of price randomness in a tractable way. Under the Markov chain, there are two price realizations: P^a = 35.7 and 44.3. The implied annual transition probabilities of staying in each state are: .71 for the low state and .83 for the high state.[40] Since the stochastic specification makes the computations more challenging, we use a coarser partitioning of the land area into 78 sites.

Table 5 reports the adjustment in the shadow-price of emissions for the business-as-usual case for a few different values of dynamic model misspecification aversion parameter, $\widehat{\xi}$, that applies to transition probabilities, as described in Section 6.4. Decreases in $\widehat{\xi}$ increase the misspecification aversion and induce the planner to consider the possibility that the probability of staying at (moving from) the least (best) advantageous state increases. We see modest drops in this price as for the values of $\widehat{\xi}$ that we consider.[41] The coarse grid alone reduces the shadow prices as we report in Table 13 of Appendix B.1.

Table 6 presents present-value decompositions for business-as-usual when agricultural prices are stochastic. Misspecification aversion diminishes both the agricultural value and the planner value, which is qualitatively similar to what we found with ambiguity aversion. In all cases, the agricultural value exceeds the planner values as forest services are negative. This adjustment is supported by uncertainty-adjusted changes in the transition distribution as reported in Table 7. The one-period transition probabilities from the low state to the low state are increased to near one, and the transition probabilities from the high state to the high state are substantially diminished. Overall, we find modest reductions in valuation due to robustness concerns. Not surprisingly, the agricultural output drops when the robustness concerns are enhanced with a partially offsetting change in forest service contributions. The aggregate implications for land allocation over time are depicted in Figure 14. There are very minor differences when comparing the trajectories for $\hat{\xi} = 1$ and the baseline $\hat{\xi} = \infty$.

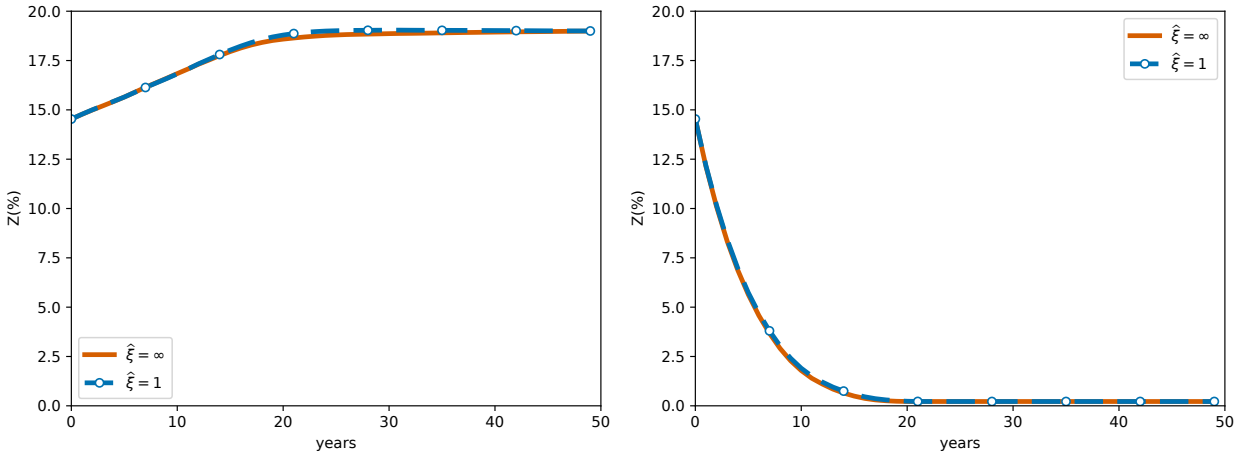
Next we report our findings when $b = 15$. As we have noted previously, with this transfer payment, the

[40] Appendix A.9 gives results for a second estimation of the hidden-state Markov process in which Gaussian shock variances are constrained to be the same. In this case, both realized states are lower and most of the time is spent in the higher of the two states.

[41] To construct the business-as-usual price for emissions when the agricultural prices are stochastic, we used the smoothed probabilities reported in left panel of Figure 16 in Appendix A.9 to assign the discrete states in our computations. While we used a probability .5 threshold for this assignment, many of the probabilities are actually close to zero or one.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

planner is much less concerned about uncertainty in potential agricultural productivities. This is evident in the right panel of Figure 14 which shows the relatively rapid decline over time in land allocated to agriculture along with perceptible changes in trajectory due to the potential misspecification concerns that the planner has about the price dynamics. Table 8 reveals the very small impact of such misspecification concerns on present-value contributions. Finally, the uncertainty adjustments to the transition probabilities for year 1 are reported in Table 9. While the robustness implications are qualitatively the same as those reported in Table 7 for $b = 0$, the magnitudes are now much more muted. Thus, uncertainty in the price dynamics becomes less consequential to the social planner when $b = 15$ in comparison to the case in which $b = 0$.[42]



(a) b=0 (b) b=15

Figure 14: Evolution of agricultural area under price uncertainty for $b = 0$ and $b = 15$.

7.5 Time-consistency

We have assumed a commitment of Brazil to carry out a plan that it finds advantageous at time zero. As restored forests mature, the rate of carbon sequestration and future expected payments decline. As a result, in future periods the seller may prefer to deviate rather than continue under the date-zero solution. We now explore how to structure a payment scheme to avoid the incentives for Brazil to deviate from this *ex ante* commitment. This is necessary to ensure the long-term forest protection and nurture reforestation are not undermined by the changes in Brazil's future perspectives. While it is beyond the scope of this paper to calculate an optimal mechanism that is immune to potential defection in the future, we now suggest a straightforward mechanism and show some numerical illustrations. For pedagogical simplicity, we presume ambiguity neutrality and constant cattle prices over time.

Our computations so far presume a fixed per-period payment of b for each net ton of CO_2 captured. We now modify this payment scheme to lower incentives for Brazil to defect. The mechanism is based on a

[42] As noted previously, fixing $b = 15$ across aversion specifications does not hold fixed the implicit social costs of carbon because the business-as-usual prices are lower for increases in misspecification concerns.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

conservation fund, controlled by a third party. It is defined by three parameters: a per-period payment $\underline{b}$ dollars for each net ton of CO_2 captured (as we have assumed previously), an additional amount b_f dollars that are paid into an interest-bearing financial account, and a threshold date τ_f after which the flow of interest of the financial account will be paid to Brazil, given that the planned reforestation is attained or exceeded. Specifically, the fund invests the full balance at the continuously-compounded interest rate δ. For $t < \tau_f$, accrued interest is added to the fund. For $t \geq \tau_f$, the fund pays to Brazil all interest accrued at t. Importantly, for each ton captured, Brazil continues to receive $\underline{b}$ and b_f continues to be added to the fund. If the net-reforestation area is below the agreed plan, buyers would stop paying $\underline{b}$ and stop depositing b_f, per ton of carbon captured, and Brazil would lose any claim to the fund's accumulated value. Note that if Brazil implements the planned net-reforestation, it would get a present value that is identical to that generated by a payment of $b = \underline{b} + b_f$ per net ton of CO_2 captured. Consequently, Brazil's optimal plan under no defection is identical to the optimal plan with payment per ton of b.

Formally, let B_t be the account balances at date t, which evolve as:

$$\dot{B}_t = \begin{cases} \delta B_t + b_f \sum_{i=1}^{I} \left(\dot{X}_t^i - \kappa Z_t^i \right) & \text{if } t < \tau_f \\ b_f \sum_{i=1}^{I} \left(\dot{X}_t^i - \kappa Z_t^i \right) & \text{if } t \geq \tau_f, \end{cases}$$

where Z and X are the optimal trajectory when transfer per-ton is b. The difference between accumulations before and after the threshold date is that after τ_f, δB_t is remitted to Brazil, but not prior to that date.

To investigate Brazil's incentive to defect from the original commitment, write $V(X, Z, b)$ for the value for Brazil if the stocks are (X, Z) and transfers per net-ton going forward are b. By construction, the discounted value of the flow b_f has the same value as the deferred payment under this proposed arrangement, which simplifies our computations. A deviation at t would occur only if

$$e^{\delta t} \int_t^{\infty} e^{-\delta s} \left[\hat{P}^a \sum_{i=1}^{I} A_s^i - (P^{ee} + b) \left(\sum_{i=1}^{I} \kappa Z_s^i - \dot{X}_s^i \right) - LC_s \right] ds < V(X_t, Z_t, 0) - B_t \tag{12}$$

where the trajectory (X_t, Z_t) is the optimal trajectory with transfers b. The term on the left side of the inequality reflects the value to continuing with the initial plan. The first term on the right reflects the fact that transfers for capture cease upon default (including not only b_f but also $\underline{b}$.) The second term is the amount accumulated in the conservation fund from past capture and for $t < \tau_f$ from accrued interest, which is lost if default occurs.

For a given $b = \underline{b} + b_f$, to guarantee that there is no defection at any $t > 0$, we need to have a sufficiently large b_f and τ_f. In addition, a larger b_f allows for a smaller τ_f, as we confirm in simulations. For instance, our calculations show that with $b = 25$, $b_f = .15b$, and $\tau_f = 15$ years, no defection would occur in 100 years. By increasing $\tau_f = 21$, we rule out defection permanently.[43] Increasing b_f to $.2b$, allows for τ_f to be

[43]This calculation was stopped at $t = 100$. However, we have three properties of the optimal path for $b = 25$, (a) Using the interior-point method we obtain, $\max_i z_{100}^i = 0.02$ (ha), which we equate to 0. This implies that the

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

zero and still not have a defection. The flow of deposits is enough to deter defection. Importantly, the ratio b_f/b can be quite modest when implementing our illustrative payment mechanism.

Remark 7.2. *While CO2 captured might seem a more directly relevant target, we suggest instead to use allocated land for performance verification. This immunizes the seller from the risk that carbon capture deviates from time zero estimates because of changes beyond Brazil's control e.g., changes in global temperature.*

Remark 7.3. *Scheinkman (2024) investigated the case in which a conservation fund is fully funded at time zero with the aim to ensure long-term conservation of the rainforest. Similarly, Harstad (2025) proposed that payments be treated as "loans" that become payable if and only if there are deviations in the trajectory of net reforestation. The results in this section suggest that requiring the payment mechanism to cover a relatively small fraction of total payments should suffice. Potentially, this will alleviate concerns from emerging market developing economies that the additional loans could hinder credit ratings.*[44]

8 Conclusions

Using a rich data set that informed us of the potential productivity of alternative land allocations, we study the impact of carbon prices on robustly optimal forest conservation and restoration over time and space in the Brazilian Amazon. We use an externally set price for emissions as an analytical tool in our analysis. As an important benchmark, we deduced a shadow price for emissions based on historical emissions to reflect "business-as-usual". When we impose this shadow price looking forward, we find that the resulting deforestation would eventually threaten the survival of large portions of the Amazon as a tropical forest. We then explored the resulting net forest restoration for robustly efficient dynamic land allocations when we augment this shadow price with alternative transfer payments. As a featured result, we found that transfer prices as low as \$25 per ton of CO_2e lead to substantial net restoration, carbon capture and economic gains for Brazil. Robustness concerns induced by productivity uncertainty on the part of the fictitious planner alter spatial and dynamic implications. Specifically, these concerns have a substantial impact on the baseline, "business-as-usual" price, but more modest implications for the implied alternative transfer payments.

To elaborate on our featured findings, we demonstrated that international carbon payments of \$25/ton can reduce emissions by about 22Gt CO_2e in 15 years and by about 32Gt CO_2e, in 30 years, relative to business-as-usual. This amount represents not only the total GtCO_2e of carbon captured by natural regeneration, for which Brazil will receive payments, but the avoided emissions from deforestation that would happen in the "business-as-usual" scenario. 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 11.3Gt of CO_2e per year globally with costs no greater than \$100 USD/ton. Our baseline simulation

left-hand side of (12) is positive and B_t increases for $t > 100$, (b) $\dot{X}_t$ is decreasing in t, for $t \geq 10$, what implies that $V(X_t, Z_t, 0)$ decreases for $t \geq 10$, and (c) $B^i_{100} > V(X_{100}, 0, 0)$. These three properties that inequality (12) holds for every $t > 100$.

[44] An example of this concern is the statement at COP30 by Richard Muyungi, chair of the African Group of Negotiators (AGN): "Africa is not ready to take an additional burden in terms of financing" (Muyungi (2025)).

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

in Table 3 suggests that the optimal management of the Brazilian Amazon can deliver about 10% of this total at a much lower effective cost.

Although the Amazon is the largest tropical forest in the world and the Brazilian Amazon occupies 60% of the Amazon, there are other substantial deforested areas in tropical forests across the world. The approach we took here could be used to calculate the potential for carbon capture and the economic costs involved in restoration across deforested areas in tropical forests.[45]

Our analysis is an initial step in a broader research agenda. To make our analysis tractable and revealing, this paper ignores some important costs of deforestation. For instance, we do not include the effect of deforestation on agricultural productivity in the Amazon (Leite-Filho et al. (2021)) or in other regions in Brazil, a country that is currently the fourth largest agricultural producer and third largest exporter in the world. We also do not take into account the loss of biodiversity or resiliency including the possibility that Amazon deforestation triggers a tipping point with broad-based consequences (Steffen et al. (2018) and Flores et al. (2024)). Finally, we do not account for the direct effect of deforestation or regeneration in one site on forests in other sites.[46] These are important considerations for future research, including assessments of which modeling extensions are most salient for the quantitative conclusions.

Data and Code Availability

Data and code for this research are available in the Harvard Dataverse replication package (Assunção et al., 2026), at `https://doi.org/10.7910/DVN/QIFRIB`. They were checked by the journal for their ability to reproduce the results presented in the paper. Given the highly demanding nature of the algorithms, the reproducibility checks were performed on a subset of results only. Those are clearly indicated in the accompanying README file of the replication package.

[45] Using the equations for carbon accumulation and the calibration methods for potential carbon content that we introduced here, Assunção et al. (2025), a report prepared for the Committee of economists advising the President of COP30, contains estimates of the carbon capture potential for restoration of areas of tropical forests deforested since 2001.

[46] See Araujo et al. (2023) for an estimate.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

References

Alliance for Restoration in the Amazon. 2020. "Forest Landscape Restoration in the Amazon: Overview and Paths to Follow." Alliance for Restoration in the Amazon Position Paper, Brazil.

Anderson, Evan W, Lars Peter Hansen, and Thomas J Sargent. 2003. "A quartet of semigroups for model specification, robustness,prices of risk, and model detection." *Journal of the European Economic Association*, 1: 68–123.

Araújo, Rafael, Francisco Costa, and Marcelo Sant'Anna. 2025. "Efficient Conservation of the Brazilian Amazon: Estimates from a Dynamic Model." *The Review of Economic Studies*. Advanced Access Publication.

Araujo, Rafael, et al. 2023. "Estimating the spatial amplification of damage caused by degradation in the Amazon." *Proceedings of the National Academy of Sciences*, 120(46): e2312451120.

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.

Assunção, Juliano, Clarissa Gandour, and Rudi Rocha. 2015. "Deforestation slowdown in the Brazilian Amazon: prices or policies?" *Environment and Development Economics*, 20(6): 692–722.

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, João Pedro Arbache, Joana Chiavari, Giovanna de Miranda, and Gabriela Zangiski. 2025. "The Forest-Climate Nexus: A Fit-for-Purpose Framework for Climate Impact." Climate Policy Initiative, Rio de Janeiro.

Assunção, Juliano, Lars Hansen, Todd Munson, and Jose Scheinkman. 2026. "Replication Data for: Carbon prices, forest conservation and reforestation in the Brazilian Amazon." *Harvard Dataverse*, Version 1. `https://doi.org/10.7910/DVN/QIFRIB`.

Assunção, Juliano, Robert McMillan, Joshua Murphy, and Eduardo Souza-Rodrigues. 2023*a*. "Optimal environmental targeting in the amazon rainforest." *The Review of Economic Studies*, 90(4): 1608–1641.

Assunção, Juliano, et al. 2023*b*. "DETER-ing Deforestation in the Amazon: Environmental Monitoring and Law Enforcement." *American Economic Journal: Applied Economics*, 15(2): 125–156.

Austin, KG, JS Baker, BL Sohngen, CM Wade, A Daigneault, SB Ohrel, S Ragnauth, and A Bean. 2020. "The economic costs of planting, preserving, and managing the world's forests to mitigate climate change." *Nature communications*, 11(1): 5946.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Balboni, Clare, Aaron Berman, Robin Burgess, and Benjamin A Olken. 2023. "The economics of tropical deforestation." *Annual Review of Economics*, 15: 723–754.

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

Busch, Jonah, and Jens Engelmann. 2017. "Cost-effectiveness of reducing emissions from tropical deforestation, 2016–2050." *Environmental Research Letters*, 13(1): 015001.

Cai, Yongyang, and Kenneth L Judd. 2023. "A simple but powerful simulated certainty equivalent approximation method for dynamic stochastic problems." *Quantitative Economics*, 14(2): 651–687.

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

Carpenter, Bob, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus A Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. "Stan: A probabilistic programming language." *Journal of statistical software*, 76.

Chiavari, Joana, Cristina L. Lopes, and Julia N. de Araujo. 2021. "Panorama dos Direitos de Propriedade no Brasil Rural." Climate Policy Initiative Relatório, Rio de Janeiro.

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

Farrokhi, Farid, Elliot Kang, Heitor S. Pellegrina, and Sebastian Sotelo. 2025. "Deforestation: A Global and Dynamic Perspective." CESifo 12077.

Fesenmyer, Kurt A, Erin E Poor, Drew E Terasaki Hart, Joseph W Veldman, Forrest Fleischman, Pooja Choksi, Sally Archibald, Mohammed Armani, Matthew E Fagan, Evan C Fricke, César Terrer, Natalia Hasler, Christopher A Williams, Peter W Ellis, and Susan C Cook-Patton. 2025. "Addressing critiques refines global estimates of reforestation potential for climate change mitigation." *Nature Communications*, 16(1): 4572.

Flores, Bernardo M., Encarni Montoya, Boris Sakschewski, Nathália Nascimento, Arie Staal, Richard A. Betts, Carolina Levis, David M. Lapola, Adriane Esquível-Muelbert, Catarina Jakovac, Carlos A. Nobre, Rafael S. Oliveira, Laura S. Borma, Da Nian, Niklas Boers, Susanna B. Hecht, Hans ter Steege, Julia Arieira, Isabella L. Lucas, Erika Berenguer, JoséA. Marengo, Luciana V. Gatti, Caio R. C. Mattos, and Marina Hirota. 2024. "Critical transitions in the Amazon forest system." *Nature*, 626(7999): 555–564.

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.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Gandour, Clarissa. 2018. "Forest Wars: A Trilogy on Combating Deforestation in the Brazilian Amazon." PhD diss. 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.

Harstad, Bård. 2025. "Improving Tropical Forest Financing: From TFFF and REDD+ to SSLB and FLL." Harvard Project on Climate Agreements Discussion Paper, Cambridge, Massachusetts.

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.

IPCC. 2023. "Synthesis Report." In *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. , ed. Core Writing Team, H. Lee and J. Romero. Cambridge, UK and New York, NY, USA:Cambridge University Press.

Kim, Sei Jin, Justin S Baker, Brent L Sohngen, and Michael Shell. 2018. "Cumulative global forest carbon implications of regional bioenergy expansion policies." *Resource and Energy Economics*, 53: 198–219.

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

Laurance, William F., G. R. Clements, and Sean et al. Sloan. 2014. "A global strategy for road building." *Nature*, 513: 229–232.

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.

Lentini, Marco, Leonardo Sobral, Marcos Planello, Robson Vieira, Felipe Cerignoni, Fernando Nunes, and Vinícius Guidoti IMAFLORA. 2019. "O que mudou no perfil da atividade madeireira na Amazônia nas últimas duas décadas (1998-2018)." *Piracicaba: Imaflora*.

Liang, Shijing, Alan D Ziegler, Peter B Reich, Kai Zhu, Dashan Wang, Xin Jiang, Deliang Chen, Philippe Ciais, and Zhenzhong Zeng. 2025. "Climate mitigation potential for targeted forestation after considering climate change, fires, and albedo." *Science Advances*, 11(15): eadn7915.

Lovejoy, Thomas E, and Carlos Nobre. 2018. "Amazon tipping point." *Science Advances*, 4(2): eaat2340.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

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.

Muyungi, Richard. 2025. ""Africa is not ready to take additional burden in terms of financing." — African Group demands grant-based climate finance at COP30." *Climate Change News*. COP30 Bulletin Day 1.

Neal, Radford M, et al. 2011. "MCMC using Hamiltonian dynamics." *Handbook of markov chain monte carlo*, 2(11): 2.

Nepstad, Daniel C., Adalberto Veríssimo, Ane Alencar, Carlos Nobre, Elisamara Lima, Paulo Lefebvre, Paulo Schlesinger, Christopher Potter, Paulo Moutinho, Emilio Mendoza, Mark Cochrane, and Viviane Brooks. 1999. "Large-scale impoverishment of Amazonian forests by logging and fire." *Nature*, 398(6727): 505–508.

Poorter, Lourens, Dylan Craven, Catarina C Jakovac, Masha T Van Der Sande, Lucy Amissah, Frans Bongers, Robin L Chazdon, Caroline E Farrior, Stephan Kambach, Jorge A Meave, et al. 2021. "Multidimensional tropical forest recovery." *Science*, 374(6573): 1370–1376.

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."

Scheinkman, José A. 2024. "Carbon prices and reforestation in tropical forests." Bank for International Settlements Working Paper 1223.

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." TSE Working Paper.

Souza Jr, Carlos M, Julia Z. Shimbo, Marcos R Rosa, Leandro L Parente, Ane A. Alencar, Bernardo FT Rudorff, Heinrich Hasenack, Marcelo Matsumoto, Laerte G. Ferreira, Pedro WM Souza-Filho, et al. 2020. "Reconstructing three decades of land use and land cover changes in brazilian biomes with landsat archive and earth engine." *Remote Sensing*, 12(17): 2735.

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.

Stan Development Team. 2023. "Stan Modeling Language Users Guide and Reference Manual." Version 2.33.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

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.

Strassburg, Bernardo BN, Alvaro Iribarrem, Hawthorne L Beyer, Carlos Leandro Cordeiro, Renato Crouzeilles, Catarina C Jakovac, André Braga Junqueira, Eduardo Lacerda, Agnieszka E Latawiec, Andrew Balmford, et al. 2020. "Global priority areas for ecosystem restoration." *Nature*, 586(7831): 724–729.

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.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

ambiguity aversion (ξ)	carbon price (P^{ee})
∞	6.8
2	5.6
1	4.8
0.5	2.8

Table 1: Business-as-usual prices for different levels of ambiguity aversion. The computations use 1043 sites. The agricultural price is presumed to be time invariant and is set to $P^a = 41.1$, which is the mean under the stationary distribution.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

P^e ($)	b ($)	agricultural output value ($ billion)	net transfers ($ billion)	forest services ($ billion)	adjustment costs ($ billion)	planner value ($ billion)
6.8	0	376	0	-108	6	261
16.8	10	88	128	87	10	293
21.8	15	48	228	103	14	364
26.8	20	26	325	110	18	443
31.8	25	16	417	113	22	525

Table 2: Present-value decomposition under ambiguity neutrality. We set $P^a = 41.1$, which is the mean agricultural price in the stationary distribution. Forest services are calculated by using baseline shadow price ($b = 0$). The present values are computed for two hundred years.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

		15 years			30 years		
P^e ($)	b ($)	net captured emissions (CO_2e Gt)	discounted net transfers ($ billion)	effective cost ($ per ton of CO_2e)	net captured emissions (CO_2e Gt)	discounted net transfers ($ billion)	effective cost ($ per ton of CO_2e)
6.8	0	-12.1	0	-	-14.7	0	-
16.8	10	6.7	56	3.6	13.7	102	4.8
21.8	15	8.8	112	6.3	16.3	186	7.9
26.8	20	9.8	166	8.9	17.4	266	10.8
31.8	25	10.3	221	11.5	17.8	344	13.7

Table 3: Transfer costs under ambiguity neutrality. We set $P^a = \$41.1$, which is the mean agricultural price in the stationary distribution.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

	agricultural output value ($ billion)			planner value ($ billion)		
b ($)	ambiguity neutral	ambiguity aversion	percent change	ambiguity neutral	ambiguity aversion	percent change
0	376	285	-24	261	190	-27
10	88	83	-6	293	255	-13
15	48	46	-4	364	327	-10
20	26	28	8	443	403	-9
25	16	18	13	525	482	-8

Table 4: Present-value decomposition under parameter ambiguity. The calculations impose $P^a = 41.1$, the average price under the stationary distribution. The valuations for ambiguity aversion are computed under uncertainty-adjusted probability measure.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

$\widehat{\xi}$	carbon price (P^{ee})
∞	6.9
1	6.5
0.5	6.1

Table 5: Business-as-usual prices for the stochastic specification of agricultural prices.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

	agricultural output value ($ billion)	forest services ($ billion)	adjustment costs ($ billion)	planner value ($ billion)
$\widehat{\xi} = \infty$	356	-111	5	239
$\widehat{\xi} = 1$	323	-97	5	221
$\widehat{\xi} = 0.5$	296	-83	5	209

Table 6: Present-value decomposition with stochastic agricultural prices for b = 0. The computations are based on the robust MPC approximations to the value functions in time period 1. The valuations for $\widehat{\xi} < \infty$ are computed using the implied uncertainty-adjusted probability measures.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

$\widehat{\xi}$	Prob from low to low	Prob from high to high
∞	0.71	0.83
1	0.91	0.52
0.5	0.98	0.16

Table 7: Uncertainty-adjusted transition probability (year 1), $b = 0$.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

	agricultural output value ($ billion)	net transfers ($ billion)	forest services ($ billion)	adjustment costs ($ billion)	planner value ($ billion)
$\widehat{\xi} = \infty$	33	237	109	15	364
$\widehat{\xi} = 1$	34	236	102	15	357
$\widehat{\xi} = 0.5$	34	235	96	14	351

Table 8: Present-value decomposition with stochastic agricultural prices for b = 15. The computations are based on the robust MPC approximations to the value functions in time period 1. The valuations for $\widehat{\xi} < \infty$ are computed using the implied uncertainty-adjusted probability measures.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

$\widehat{\xi}$	Prob from low to low	Prob from high to high
∞	0.71	0.83
1	0.72	0.81
0.5	0.74	0.80

Table 9: Uncertainty-adjusted transition probability (year 1), b = 15.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Figure legends

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).
Alt text: Log-scale scatterplot of emissions per capita against GDP per capita. Most countries form an upward-sloping cluster; the Brazilian Amazon lies well above countries with similar income and above the United States, European Union, China, and India.

Figure 2: Percentage of maximum carbon density: actual vs. theoretical. The black solid curve gives the an empirically-based construction of the carbon recovery fraction as a function of the recovery time. The blue dashed curve plots $1 - \exp(-\alpha t)$.
Alt text: Line chart of the fraction of maximum carbon density reached against age of regrowing forest, where a solid data line rises steeply and concavely from one-third at year zero to near nine-tenths by year thirty, staying above a lower, more gradual dashed theoretical line throughout.

Figure 3: Initial values for agricultural area (z_0^i) and carbon stock (x_0^i)
Alt text: Two maps of 1,043 Amazon sites. Agricultural area is concentrated along the southern and eastern edges, while stored carbon is greatest in western, northern, and central areas.

Figure 4: Carbon sequestration parameters and Agricultural productivity heterogeneity
Alt text: Two maps of site-level productivity parameters. Carbon sequestration capacity is generally higher in western and central sites, while agricultural productivity is concentrated along the southern and eastern margins.

Figure 5: Agricultural area and cumulative CO_2 capture.
Alt text: Two line charts over fifty years for transfer prices of zero, fifteen, and twenty-five dollars. Agricultural area rises toward one-fifth of the biome without transfers but falls to near zero with them; cumulative capture turns negative without transfers and climbs steadily with them.

Figure 6: Evolution of transfer payments for two choices of b.
Alt text: Line chart of net transfer payments over fifty years, where both curves rise sharply to a peak near year eight to ten, then decline gradually. The higher transfer level peaks around twenty-two billion dollars, the lower around eleven billion.

Figure 7: Agricultural area changes after 30 years.
Alt text: Four Amazon maps compare agricultural shares in 2017 and after thirty years under transfers of zero, ten, and twenty-five dollars. Agriculture expands along the southeast without transfers, contracts sharply at ten dollars, and nearly disappears at twenty-five dollars.

Figure 8: Agricultural area evolution over time

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Alt text: Six Amazon maps track agricultural shares from 2017 to 2047 under a fifteen-dollar transfer. Agriculture retreats rapidly from most sites during the first fifteen years and remains concentrated in a small southern and eastern fringe.

Figure 9: Ambiguity-adjusted densities for four sites. The left plots report sites with the largest entropies for θ and right plots report sites with the largest entropies for γ. The upper plots are $b = 0$ and the lower plots for $b = 15$.
Alt text: Four density plots compare baseline and ambiguity-adjusted productivity distributions. The agricultural-productivity densities shift substantially left for both transfer cases; the carbon-sequestration densities change little at zero transfers and shift modestly left at fifteen dollars.

Figure 10: Histogram of the years in which one of the two controls is maximal, for $b = 0$ and common $P^{ee} = 6.8$. The controls remain at zero for fifty years in 10.1% of the sites under ambiguity neutral and in 10.2% of sites under ambiguity aversion.
Alt text: Overlaid histogram of the years when deforestation or reforestation is greatest across sites. Both specifications peak near year zero, but ambiguity aversion places somewhat more mass at later positive years and less at earlier negative years.

Figure 11: Evolution of agricultural area under ambiguity neutrality and ambiguity aversion for $b = 0$ at a common business-as-usual carbon price, $P^{ee} = 6.8$.
Alt text: Line chart of agricultural land share over fifty years at a common carbon price without transfers and two aversion settings. Under ambiguity neutrality the share rises toward twenty percent; under ambiguity aversion it dips slightly and then returns near its initial level.

Figure 12: Histogram of the years in which one of the two controls is maximal, $P^{ee} = 6.8$ for ambiguity neutral and $P^{ee} = 4.8$ for ambiguity averse. For $b = 0$, 10.1% of the sites remain unchanged under ambiguity neutrality, while 13.2% do so under ambiguity aversion. For $b = 15$, 2.0% of the sites do not change under ambiguity neutrality, and 2.1% remain unchanged under ambiguity aversion. Panels: (a) $b = 0$; (b) $b = 15$.
Alt text: Two overlaid histograms compare the timing of peak land-use adjustments. With zero transfers, observations cluster near year zero with long tails; with fifteen dollars, most sites peak within the first several years and differences across ambiguity specifications are small.

Figure 13: Evolution of agricultural area under ambiguity neutrality and ambiguity aversion, using the corresponding shadow prices. Panels: (a) $b = 0$; (b) $b = 15$.
Alt text: Two line charts compare agricultural land shares under ambiguity neutrality and aversion at their respective shadow prices. With zero transfers both rise, and the ambiguity-averse path is higher; with fifteen dollars both fall rapidly and nearly overlap.

Figure 14: Evolution of agricultural area under price uncertainty for $b = 0$ and $b = 15$. Panels: (a) $b = 0$; (b) $b = 15$.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Alt text: Two line charts compare agricultural land shares with and without price-model ambiguity. For zero transfers, both paths rise and nearly overlap; for fifteen dollars, both fall rapidly toward zero with only small differences.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Online Appendix for "Carbon prices, forest conservation and reforestation in the Brazilian Amazon"

Juliano Assunção Lars Peter Hansen Todd Munson
José A. Scheinkman

September 2, 2026

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

A Data construction

This appendix describes how we measure all variables and parameters used in the model.

A.1 Data sources and data citations

All external data used in the empirical construction are public administrative, market, climate, remote-sensing, or carbon-pricing data. The analysis uses no confidential or human-subject microdata. The replication archive contains the exact raw extracts used by the authors under `data/raw/`, together with the cleaned, processed, and calibrated files generated by the replication code under `data/clean/`, `data/processed/`, and `data/calibration/`. The machine-readable data inventory and source-permission audit are mirrored in `data/codebook.csv` and `data/source_permissions.csv`. Source-specific license and redistribution terms remain with the original providers.

External data sources. The dataset-level citations, uses, and archived inputs are as follows.

1. World Bank. *World Development Indicators*. URLs: `https://data.worldbank.org/` and `https://databank.worldbank.org/source/world-development-indicators`. Used for GDP per capita and CO_2 emissions in Figure 1. Figure 1 World Bank files were downloaded in March 2021 and are archived under `data/raw/worldbank/emission_kuznets/` and `replication/figure1/input/`; metadata are under `replication/figure1/documentation/`.
2. World Bank. *Carbon Pricing Dashboard*. URL: `https://carbonpricingdashboard.worldbank.org/`. Used as carbon-pricing source inputs in data-cleaning scripts and contextual outputs. Archived under `data/raw/worldbank/carbon_price/`.
3. World Resources Institute. *Climate Watch*. URL: `https://www.climatewatchdata.org/`. Used with the World Bank inputs for the Figure 1 emissions context. Exact extracts and metadata are retained under `replication/figure1/` and `data/raw/worldbank/`.
4. Amazonia 2030. 2021. *Fatos da Amazônia 2021*. Amazonia 2030. URL: `https://amazonia2030.org.br/fatos-da-amazonia-2021/`. Used for the Brazilian Amazon point in Figure 1. Reference files and the derived source-data CSV are archived under `replication/figure1/reference/` and `replication/derived/figure1_source_data.csv`.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

5. MapBiomas Project. *Collection 5 Annual Land Cover and Land Use Maps of Brazil* and related MapBiomas products. URL: `https://mapbiomas.org/`. Used for land-use and land-cover classifications, total available area, agricultural area, historical deforestation, secondary-vegetation age, pasture quality, and basin-related inputs. Raw files are archived under `data/raw/mapbiomas/`. The methodological article is Souza Jr. et al. (2020), but the dataset itself is cited here as MapBiomas Collection 5.
6. Agência Nacional de Águas e Saneamento Básico (ANA). 2006. *National Water Resources Plan basin layers, level 2*. Distributed through MapBiomas basin inputs. URL: `https://metadados.snirh.gov.br/geonetwork/srv/api/records/6141f37f-f15d-42e7-8495-ae9ddad0846f/attachments/GEOFT_PNRH_SUB2.zip`. Used for basin groups in the agricultural-productivity random effects and archived under `data/raw/mapbiomas/basin/`.
7. European Space Agency Climate Change Initiative Biomass Project. 2021. *ESA Biomass Climate Change Initiative (Biomass_cci): Global Datasets of Forest Above-Ground Biomass for the Years 2010, 2017 and 2018, v3*. NERC EDS Centre for Environmental Data Analysis. DOI: 10.5285/5f331c418e9f4935b8eb1b836f8a91b8. Project URL: `https://climate.esa.int/en/projects/biomass/`. The 2017 above-ground biomass rasters are used to construct carbon absorption and carbon-productivity inputs and are archived under `data/raw/esa/above_ground_biomass/`.
8. Sistema de Estimativas de Emissões e Remoções de Gases de Efeito Estufa (SEEG). *Greenhouse-gas emissions and removals data*. URL: `https://seeg.eco.br/`. Used to calibrate agricultural net emissions. Archived under `data/raw/seeg/emission/`; raw filenames record a 2020.11.05 extract.
9. Instituto Brasileiro de Geografia e Estatística (IBGE). 2017. *Censo Agropecuário 2017*, including Tables 6882 and 6911, and associated municipal and biome geographies. URL: `https://censoagro2017.ibge.gov.br`. Used for municipal boundaries, the Amazon biome geography, agricultural census land-use inputs, pasture area, and cattle outputs. Archived under `data/raw/ibge/`.
10. Secretaria da Agricultura e do Abastecimento do Paraná, Departamento de Economia Rural (SEAB-PR/DERAL). 2021. *Preço Médio – Recebido pelo Agricultor: Boi Gordo, Arroz (em Casca), Cana-de-Açúcar, Milho, Mandioca, 1990–2021*. Distributed by Instituto de Pesquisa Econômica Aplicada (IPEA), Ministério da Economia. URL: `https://www.ipeadata.gov.br/`. Used for the monthly cattle-price process and farm-gate price inputs. Accessed through IPEA on February 22, 2021 and archived under `data/raw/seabpr/commodity_prices/` and

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

data/raw/ipea/farm_gate_price/.

11. Instituto de Pesquisa Econômica Aplicada (IPEA). *Ipeadata geographic and economic tables*. URL: https://www.ipeadata.gov.br/. Used for the distance-to-state-capital regressor in the agricultural-productivity regression. Archived under data/raw/ipea/distance_to_capital/; the raw filename records an August 21, 2023 extract.
12. Fick, Stephen E., and Robert J. Hijmans. 2017. "WorldClim 2: New 1-km Spatial Resolution Climate Surfaces for Global Land Areas." *International Journal of Climatology* 37(12): 4302–4315. DOI: https://doi.org/10.1002/joc.5086. Project URL: https://www.worldclim.org/. WorldClim 2.1 historical temperature and precipitation rasters are used as productivity-regression covariates; the 2.5-minute historical rasters are archived under data/raw/worldclim/.
13. Fundação Getulio Vargas, Instituto Brasileiro de Economia (FGV IBRE). *Deflator series used for price preparation*. URL: https://portalibre.fgv.br/. Used to express price series in real terms and archived under data/raw/fgv/deflator_ipa/.

Generated analysis inputs are included with open-format companions where practical. In particular, calibration_1043_sites.Rdata and calibration_78_sites.Rdata have CSV and GeoJSON companions for the site-level calibration panels and grids; carbon and agricultural productivity objects have corresponding gamma_fit_* and theta_fit_* CSV or GeoJSON files. The main generated files used by the model and post-processing scripts are the 1,043-site and 78-site calibration panels, the spatial grids, the fitted productivity files, and the distribution-parameter summaries under data/calibration/.

The subsequent subsections describe the construction of the model inputs from these sources. Briefly, MapBiomas 2017 land-cover shares define available and initial agricultural area; ESA Biomass CCI 2017 above-ground biomass, matched to MapBiomas primary-forest pixels, defines site-level carbon productivity; SEEG state-level agricultural net emissions divided by MapBiomas agricultural area calibrate κ; IBGE 2017 Agricultural Census data define municipal cattle productivity; SEAB-PR/DERAL prices distributed through IPEA, deflated with the FGV IBRE series, define the cattle-price process; and WorldClim, IPEA distance-to-capital data, and ANA/PNRH basin layers supply covariates and grouping variables for the productivity regressions.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

A.2 Total available area

To compute $\bar{z}^i$, the amount of available area for the planner's choice (forest or cattle farming) in each site i, we first calculate the fraction of 30 m-pixels in site i classified as agriculture (crops + pastures) or forests in MapBiomas 2017 (Souza Jr et al., 2020). We then multiply this fraction by the area (within the biome) of the site, to obtain a measure in hectares. Notice $\bar{z}^i$ comprises the total site area, excluding areas such as rivers, roads, cities and etc.

A.3 Carbon absorption

We first extract a random sample of 1.2M 30 m-pixels and select 893,753 pixels with no deforestation during 1985-2017, which we treat as primary forests as of 2017. We add *above ground* biomass density data for the year 2017 from ESA Biomass[1]. The biomass data also comes in a grid format with $\sim$100m resolution, so we spatially match it to our sample. The original data is measured in biomass density (Mg per ha) but we convert it to carbon per hectare, by dividing by 2 (carbon is approximately 50% of the biomass), and then obtain CO_2e equivalent by multiplying by 44 and dividing by 12 (based on atomic mass). In Appendix D we exposit how we use the data to obtain a baseline distribution of the vector of site-specific carbon absorption productivities, $(\gamma^1, \ldots, \gamma^I)$.

A.4 Carbon depreciation

The parameter α is a carbon depreciation parameter, assumed to be constant across sites. It is set so that the 99% convergence time of the carbon accumulation process is 100 years (see Heinrich et al. (2021)), that is $\alpha = 1 - (1 - 0.99)^{1/100} = .045$.

A.5 Emissions contributed by agriculture

The parameter κ is calibrated using the agricultural net annual emission data at the state level available from the system SEEG.[2] We use $\kappa = 2.0942$, which is the average of agricultural net emission divided by the agricultural area from MapBiomas for all states within the Amazon biome, weighting by the area of each state overlap with the biome, from 1990 to 2019.

[1] (Santoro and Cartus, 2021)

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

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

A.6 Cattle farming productivities

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 for each site. Since we do not have measurements concerning the cost of attracting or redeploying variable inputs to 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.[3] We consider the value of cattle sold for slaughter per hectare of pastureland at the municipal level, from the 2017 Agricultural Census (IBGE, 2017). In Appendix D we exposit how we derive a baseline distribution for the vector of site-specific cattle farming productivities, $(\theta^1, \ldots, \theta^I)$.

A.7 Discount rate and adjustment costs

We use the discount rate $\delta = 0.02$. To calibrate ζ_1, we compute the average marginal cost of deforestation implied by our model using data from MapBiomas on annual historical deforestation between 2008 – 2017 (Souza et al., 2020) and match this to the difference in prices for forested and clear land (Araújo, Costa and Sant'Anna, 2025). To calibrate ζ_2, we compute the average marginal cost of natural reforestation using data from MapBiomas on annual historical secondary vegetation age (Souza et al., 2020) and match this to natural reforestation costs in Benini and Adeodato (2017).

Figure 15 illustrates that the aggregate agricultural allocations over time are not very sensitive to if the subjective discount rate is raised from two to three percent. There is very modest sensitivity when $b = 0$ and essentially no sensitivity when $b = 25$.

[3] In contrast to other areas in Brazil, average value of slaughter per hectare of pasture in the Amazon, decreased between 2006 and 2017, making it doubtful that future productivity will increase.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

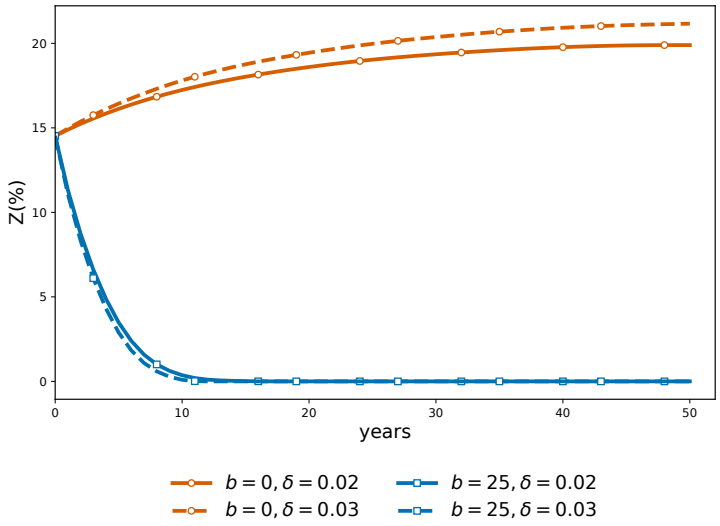


Figure 15: Agricultural area trajectories for two different discount rates.

A.8 Initial values for location-specific land allocation and stored carbon

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 assumed to be given by $X_0^i = \gamma^i(\bar{Z}^i - Z_0^i)$, i.e., the carbon stock per hectare of forest times the forest area. Notice that X_0^i is measured in CO_2e (Mg). Notice that we assume that all forest at the initial point is primary, which is compatible with equation (2).

A.9 Agricultural prices

We use a data series on monthly deflated cattle prices (reference date 01/2017),[4] from 1995, the year in which the Real Plan stabilized the Brazilian currency, until 2017.

[4]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)

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

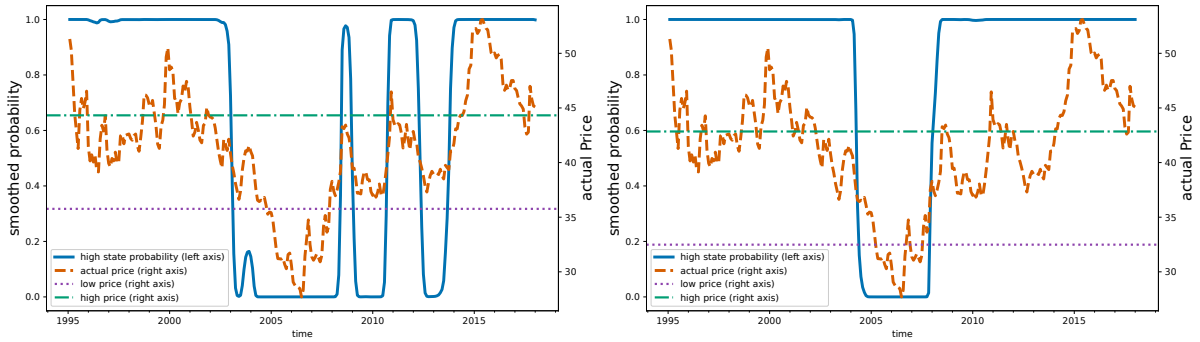
For the model inputs, we fit a two-state Markov process as a hidden-state Markov chain with with Gaussian noise. We estimated two versions of this model using the **hmmlearn** package in python. This package provides a collection of software tools for analyzing Hidden State Markov Models. In estimation, the hidden states were initialized in the implied stationary distribution of the transition probabilities through an iterative process. The implied calibration we used for results reported in the main body of the paper allowed for the normally distributed variances to be different depending on the state. We also considered a specification in which the variances are the same. The state realizations and transition probabilities for the two specifications are given in Table 10.

Table 10: Estimates for the hidden-state Markov models

	distinct variances		a common variance	
	low price	high price	low price	high price
	35.76	44.32	32.49	42.85
s.d.	0.106	0.075	0.089	0.089
	transition probabilities			
	low	high	low	high
low	0.707	0.293	0.762	0.238
high	0.174	0.826	0.041	0.959

The smoothed probabilities for both models are given in Figure 16. The more constrained estimation picks lower values for both states but assumes the process spends most of its time in the higher of the two states.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.



(a) distinct variances (b) common variance

Figure 16: Smoothed probabilities for the two hidden state Markov chain models

Table 11 reports the likelihoods and AIC and BIC model selection diagnostics for both models. The AIC criterion picks the less constrained of the two models and the BIC criterion just the opposite.

Table 11: Likelihood ratios and information criteria for the hidden state Markov chain estimation

	distinct variances	common variance
log likelihood	270.16	268.04
aic	-528.32	-526.09
bic	-502.98	-504.37

In Table 12, we report the counterpart to Table 6 constructed using the implied calibration the same variance for each state. The differences between results are modest.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

	agricultural output value ($ billion)	net transfers ($ billion)	forest services ($ billion)	adjustment costs ($ billion)	planner value ($ billion)
b = 0					
$\widehat{\xi} = \infty$	371	0	-121	6	244
$\widehat{\xi} = 1$	350	0	-113	6	231
$\widehat{\xi} = 0.5$	288	0	-84	5	199
b = 10					
$\widehat{\xi} = \infty$	63	140	94	10	286
$\widehat{\xi} = 1$	64	138	86	9	279
$\widehat{\xi} = 0.5$	73	132	74	9	270
b = 15					
$\widehat{\xi} = \infty$	29	240	107	15	360
$\widehat{\xi} = 1$	32	237	98	14	352
$\widehat{\xi} = 0.5$	37	233	87	14	343
b = 20					
$\widehat{\xi} = \infty$	23	327	110	19	441
$\widehat{\xi} = 1$	23	327	101	18	433
$\widehat{\xi} = 0.5$	24	326	91	18	423
b = 25					
$\widehat{\xi} = \infty$	15	419	112	22	525
$\widehat{\xi} = 1$	15	419	104	22	516
$\widehat{\xi} = 0.5$	15	418	94	21	506

Table 12: Hidden state Markov chain model with a common variance. The computations are based on the robust MPC approximations to the value functions in time period 1. The valuations for $\widehat{\xi} < \infty$ are computed using the implied uncertainty-adjusted probability measures.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

B Sensitivity to Grid Resolution and Ambiguity Aversion

B.1 Coarse-grid sites

In table 13 we present value decomposition results for coarse 78-sites. The shadow price decreases from \$6.8 for 1043 sites to \$6.7 for 78 sites, accounting for the slight decline in agricultural output and planner value.

P^e (\$)	b (\$)	Agricultural output value (\$ billion)	Net transfers (\$ billion)	Forest services (\$ billion)	Adjustment costs (\$ billion)	Planner value (\$ billion)
6.7	0	355	0	-114	5	236
16.7	10	61	140	94	10	285
21.7	15	28	240	107	15	360
26.7	20	22	327	110	19	441
31.7	25	14	420	112	22	524

Table 13: Present-value decomposition under ambiguity neutrality for 78 sites. We set $P^a = 41.1$, which is the mean agricultural price in the stationary distribution. Forest services are calculated using baseline shadow price ($b = 0$). The present values are computed for two hundred years.

B.2 Transfer costs under ambiguity aversion

In Tables 14 and 15, we report transfer costs under ambiguity aversion with $\xi = 1$ for 15 and 30 year horizons, respectively.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Table 14: Transfer costs under ambiguity - 15 years

	ambiguity neutral			ambiguity aversion		
b ($)	net captured emissions (CO_2e Gt)	discounted net transfers ($ billion)	effective cost ($ per ton of CO_2e)	net captured emissions (CO_2e Gt)	discounted net transfers ($ billion)	effective cost ($ per ton of CO_2e)
0	-12.1	0	-	-10.7	0	-
10	6.7	56	3.6	6.1	51	3.6
15	8.8	112	6.3	8.4	106	6.6
20	9.8	166	8.9	9.3	159	9.3
25	10.3	221	11.5	9.9	212	12.0

Notes: Agricultural price $P^a = 41.1$, which is the mean of the agricultural price in the stationary distribution. Shadow prices are $P^{ee} = 6.8$ under ambiguity neutral and $P^{ee} = 4.8$ under ambiguity aversion.

Table 15: Transfer costs under ambiguity - 30 years

	ambiguity neutral			ambiguity aversion		
b ($)	net captured emissions (CO_2e Gt)	discounted net transfers ($ billion)	effective cost ($ per ton of CO_2e)	net captured emissions (CO_2e Gt)	discounted net transfers ($ billion)	effective cost ($ per ton of CO_2e)
0	-14.7	0	-	-16.9	0	-
10	13.7	102	4.8	13.0	96	4.3
15	16.3	186	7.9	15.9	180	7.3
20	17.4	266	10.8	16.8	257	10.0
25	17.8	344	13.7	17.4	334	12.7

Notes: Agricultural price $P^a = 41.1$, which is the mean of the agricultural price in the stationary distribution. Shadow prices are $P^{ee} = 6.8$ under ambiguity neutral and $P^{ee} = 4.8$ under ambiguity aversion.

B.3 Relative entropy of the sites

Figure 17 presents the ranking of relative entropy for all 1043 sites under $\xi = 1$. The four sites with the highest relative entropy are circled, one in each corresponding map. The numbers index from low value to high value.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

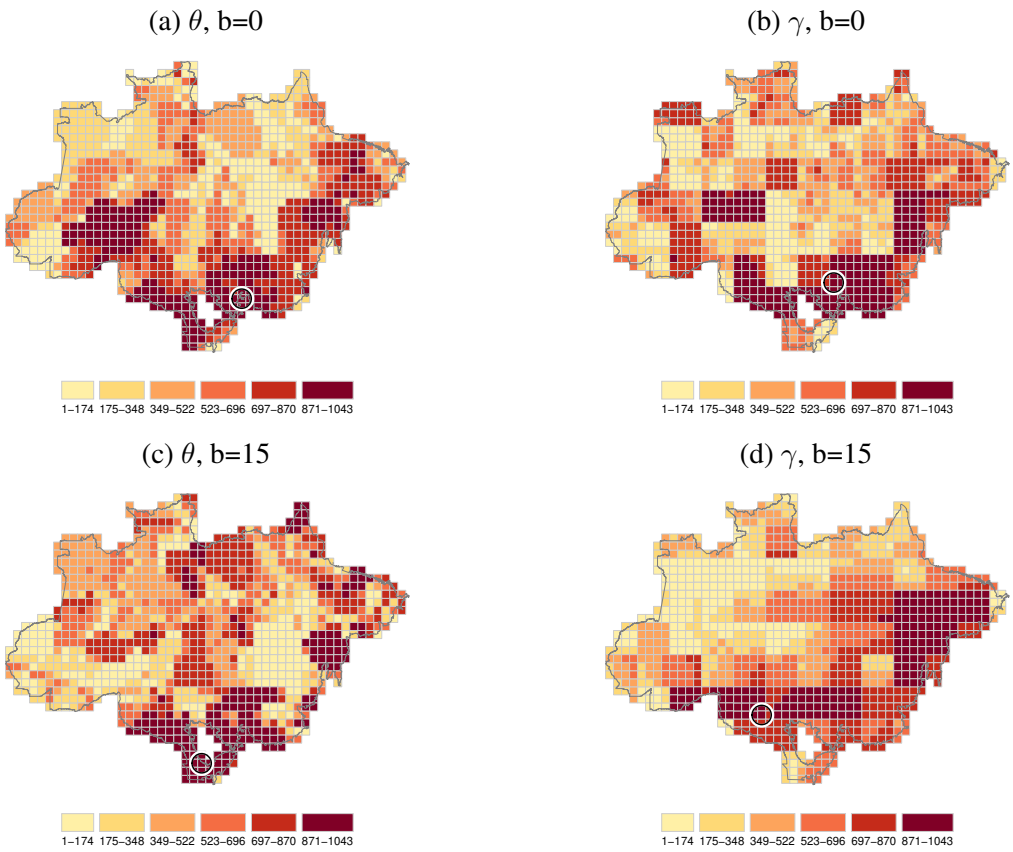


Figure 17: Relative entropy for $\xi = 1$. Sites with the highest relative entropy are circled.

B.4 Alternative values for the ambiguity aversion parameter

In this section, we report results for parameter uncertainty with $\xi = 2$ and $\xi = 0.5$. The calculated business-as-usual price is \$5.6 for $\xi = 2$ and \$2.8 for $\xi = 0.5$. Table 16 and Table 17 show the present values under $\xi = 0.5, 2$ in comparison to $\xi = \infty$. Figure 18 and 19 show the baseline and ambiguity-adjusted distributions, which are counterparts to Figure 9. Sites chosen are the same as $\xi = 1$ for comparison.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

	agricultural output value ($ billion)			planner value ($ billion)		
b ($)	ambiguity neutral	ambiguity aversion	percent change	ambiguity neutral	ambiguity aversion	percent change
0	376	321	-15	261	217	-17
10	88	87	-1	293	271	-8
15	48	47	-2	364	342	-6
20	26	27	4	443	419	-5
25	16	18	13	525	500	-5

Table 16: Present-value decomposition under parameter ambiguity, $\xi = 2$

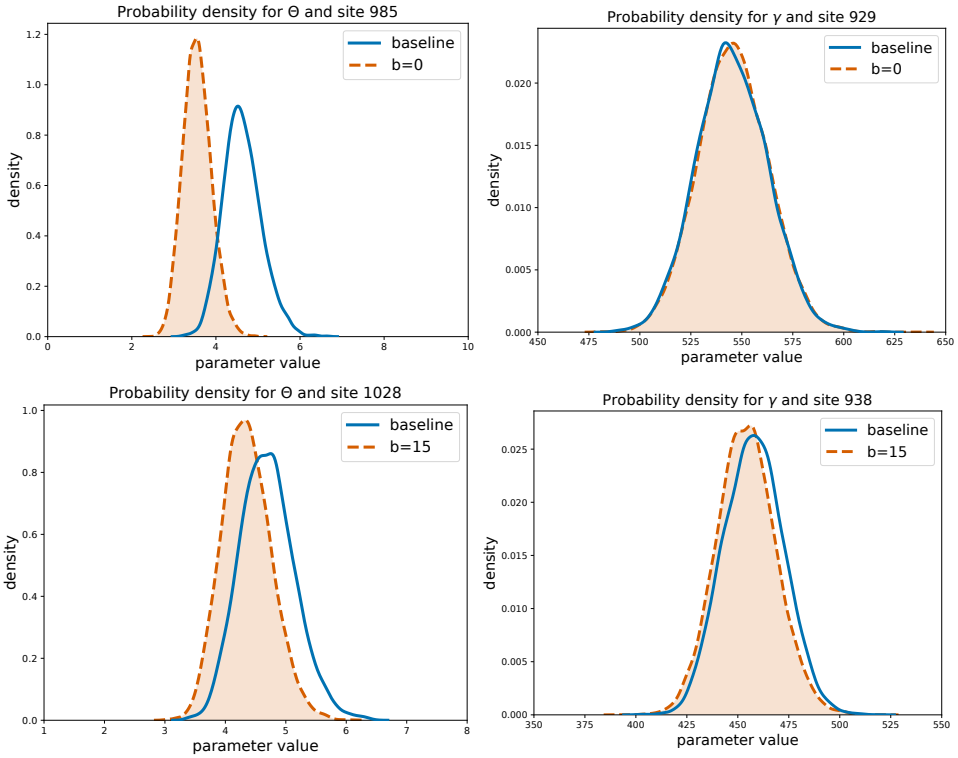


Figure 18: Ambiguity-adjusted densities for four sites, with $\xi = 2$

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

	agricultural output value ($ billion)			planner value ($ billion)		
b ($)	ambiguity neutral	ambiguity aversion	percent change	ambiguity neutral	ambiguity aversion	percent change
0	376	186	-51	261	119	-54
10	88	80	-9	293	220	-25
15	48	45	-6	364	291	-20
20	26	30	15	443	366	-17
25	16	19	19	525	442	-16

Table 17: Present-value decomposition under parameter ambiguity, $\xi = 0.5$

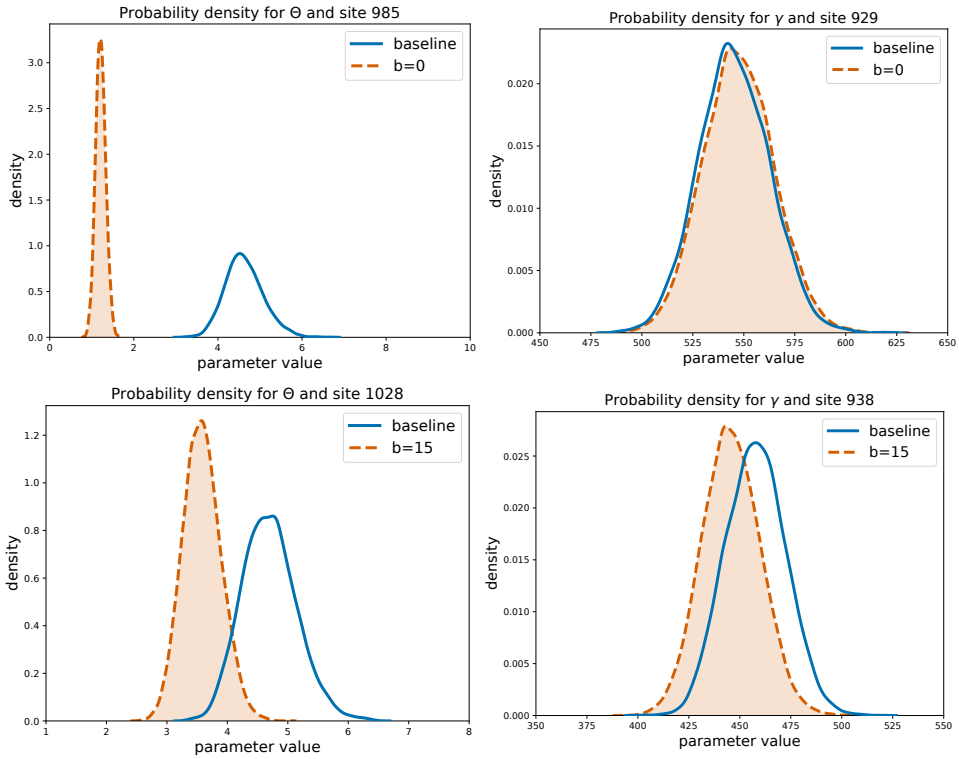


Figure 19: Ambiguity-adjusted densities for four sites, with $\xi = 0.5$

Figure 20 and 21 present evolution of agricultural land allocations under $\xi = 1$ and $\xi = 0.5$ respectively, which are counterparts to figure 8. Notice that the differences in agricultural area after 30 years are minor, but there is slightly less reforestation when $\xi = 0.5$.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

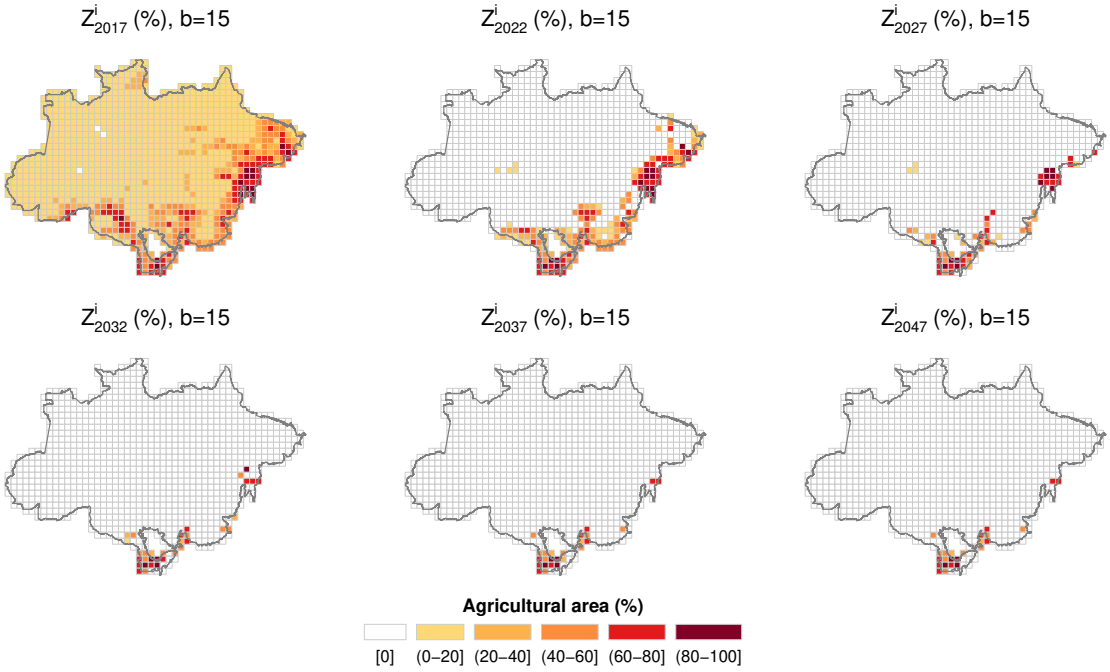


Figure 20: Agricultural area evolution over time with ambiguity aversion, $\xi = 1$

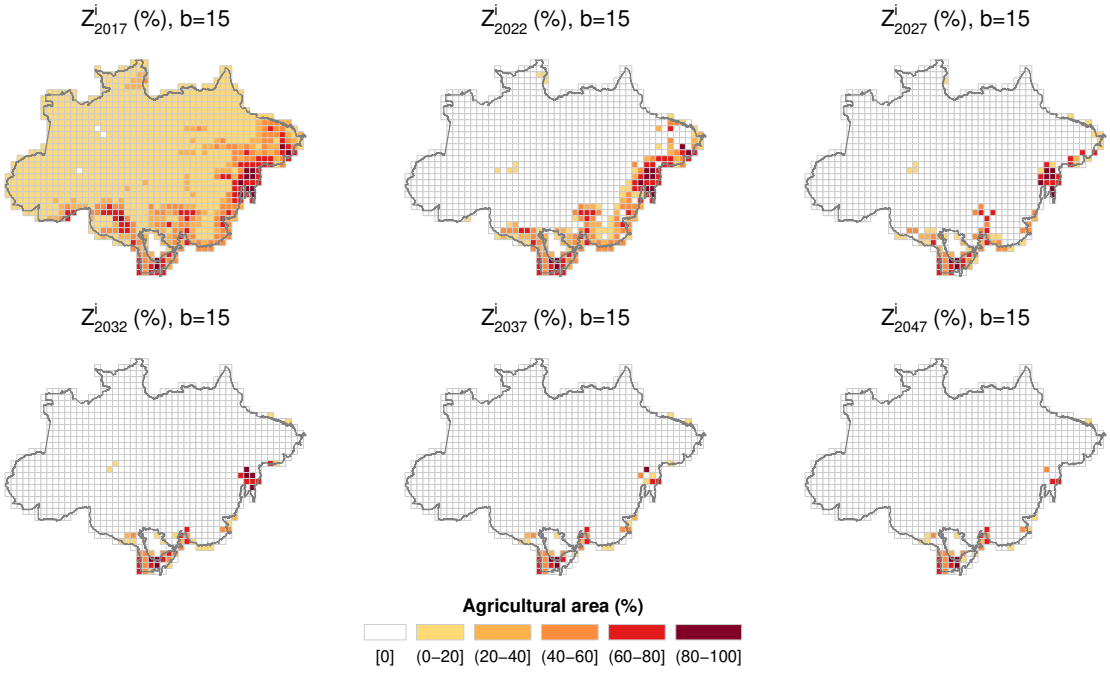


Figure 21: Agricultural area evolution over time with ambiguity aversion, $\xi = 0.5$

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

C Model discretization

In order to obtain numerical solutions for the social planner problem, we solve the following discrete-time approximation, for a finite horizon of $T = 200$ years[5]:

$$\max_{\{U_t,V_t\}_{t=0}^{T}} \sum_{t=0}^{T} e^{-\delta t}\left(P^a\sum_{i=1}^{I}A_{t+1}^i - P^e\left[\sum_{i=1}^{I}\kappa Z_{t+1}^i - \left(X_{t+1}^i - X_t^i\right)\right] - \frac{\zeta_1}{2}\left(\sum_{i=1}^{I}U_t^i\right)^2 - \frac{\zeta_2}{2}\left(\sum_{i=1}^{I}V_t^i\right)^2\right), \tag{13}$$

subject to the initial conditions in A.8 and the constraints:

$$X_{t+1}^i = X_t^i - \gamma^i U_t^i - \alpha\left[X_t^i - \gamma^i(\bar{z}^i - Z_t^i)\right] \qquad \forall i = 1,\dots I \text{ and } t = 0,\dots,T \tag{14}$$

$$Z_{t+1}^i = Z_t^i + U_t^i - V_t^i \qquad \forall i = 1,\dots I \text{ and } t = 0,\dots,T \tag{15}$$

$$U_t^i \geq 0, \quad V_t^i \geq 0 \qquad \forall i = 1,\dots I \text{ and } t = 0,\dots,T. \tag{16}$$

C.1 MPC robustness

Label the potential state realizations as $j = 1, 2, ..., n$. Let $\pi_{j\tilde{j}}$ be the probability of going to state $\tilde{j}$ given the current state j. Let h denote a finite (typically short) prediction horizon. Over h time periods, there are $N = n^h$ potential realizations of the Markov state. Let

$$k^0 = \begin{bmatrix} k_1 & k_2 & .... & k_h \end{bmatrix}$$

where $k_\ell \in \{1, 2, ..., n\}$ for $\ell = 1, 2, ..., h$. For each k^0 let:

$$k^{-\ell} = \begin{bmatrix} k_1 & k_2 & .... & k_{h-\ell} \end{bmatrix}$$

[5]Since period-payoff can be bounded by a constant, given the discount rates we use, the loss in precision for trajectories in the first 30 years, which is our period of interest, is small.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

for $1 \leq \ell \leq h-1$. Notice that $k^{-\ell}$ is a $1 \times (h-\ell)$ row vector for $0 \leq \ell \leq h-1$. With this construction, we think of actions at $t+1$ depending on the date $t+1$ (predetermined) endogenous state variables and the date $t+1$ realization of the agricultural price k^{-h+1}.

Let $U_{t+1,1}(k^{-h+1})$ denote the implied date $t+1$ contribution to utility. The date $t+2$ actions depend on the date $t+1$ endogenous state variables and the two-period agricultural price realizations k^{-h+2}. We let $U_{t+1,2}(k^{-h+2})$ denote the date $t+2$ utility contribution, and so on until we reach the final period of the uncertainty horizon, $t+h$.

To approximate the date $t+1$ objective for a given prediction, we work backwards from N possible terminal value functions given the date $t+1$ endogenous state variables, the decisions constrained by the informational case just described. We start with N possible terminal value functions $\overline{\mathcal{C}}_{t+1,h}(k^0)$ pertinent for $t+h$. These are constructed as the deterministic discounted sums scaled by $1-\exp(-\delta)$. Compute

$$\begin{aligned}\mathcal{C}_{t+1,h-1}(k^{-1}) = &[1-\exp(-\delta)]\, U_{t+1,h-1}(k^{-1}) \\ &- \exp(-\delta)\widehat{\xi}\log\sum_{j=1}^{n}\pi_{k_{h-1},j}\exp\left[-\frac{1}{\widehat{\xi}}\overline{\mathcal{C}}_{t+1,h}(k^{-1},j)\right].\end{aligned}$$

Note that there are n^{h-1} value functions. Proceeding recursively, we form:

$$\begin{aligned}\mathcal{C}_{t+1,h-\ell}(k^{-\ell}) = &[1-\exp(-\delta)]\, U_{t+1,h-\ell}(k^{-\ell}) \\ &- \exp(-\delta)\widehat{\xi}\log\sum_{j=1}^{n}\pi_{k_{h-\ell},j}\exp\left[-\frac{1}{\widehat{\xi}}\mathcal{C}_{t+1,h-\ell+1}(k^{-\ell},j)\right]\end{aligned}$$

Repeat this backward induction computation going all the way back to obtain $\mathcal{C}_{t+1,1}(k^{-h+1})$.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

$\widehat{\xi} = \infty$	agricultural output value ($ billion)	net transfers ($ billion)	forest services ($ billion)	adjustment costs ($ billion)	planner value ($ billion)
$b = 0$					
initial period solution	356	0	-111	5	239
simulation	331	0	-93	5	233
$b = 10$					
initial period solution	59	142	98	10	289
simulation	62	140	96	10	288
$b = 15$					
initial period solution	33	237	109	15	364
simulation	28	240	110	15	363
$b = 25$					
initial period solution	15	420	116	22	528
simulation	14	420	116	22	528

Table 18: Value decomposition of $\hat{\xi} = \infty$. In the results reported in the paper, we based the decomposition on using the MPC outcomes at the initial date. For comparison, we followed the usual MPC practice of resolving the problem at each subsequent date to obtain the date-specific contributions to the present value using simulated time paths for the price process and computed discounted present values averaged over the simulated price paths. The reported numbers are based on 50 simulations of the prices under the baseline Markov specification. This table reports only modest differences across numerical methods, with discrepancies that are most notable when $b = 0$.

Table 19, 20, 21 and 22 present value decomposition and distorted transition probabilities for b=10 and 25, as counterparts to table 6 and 7.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

	agricultural output value ($ billion)	net transfers ($ billion)	forest services ($ billion)	adjustment costs ($ billion)	planner value ($ billion)
$\widehat{\xi} = \infty$	59	142	98	10	289
$\widehat{\xi} = 1$	62	140	91	10	283
$\widehat{\xi} = 0.5$	69	135	82	9	277

Table 19: Present-value decomposition with stochastic agricultural prices for b = 10. The computations are based on the robust MPC approximations to the value functions in time period 1. The valuations for $\widehat{\xi} < \infty$ are computed using the implied uncertainty-adjusted probability measures.

$\widehat{\xi}$	Prob from low to low	Prob from high to high
∞	0.71	0.83
1	0.75	0.79
0.5	0.79	0.75

Table 20: Representative distorted transition probability (Year 1), b = 10

	agricultural output value ($ billion)	net transfers ($ billion)	forest services ($ billion)	adjustment costs ($ billion)	planner value ($ billion)
$\widehat{\xi} = \infty$	15	420	116	22	528
$\widehat{\xi} = 1$	15	419	109	22	521
$\widehat{\xi} = 0.5$	15	419	102	22	515

Table 21: Present-value decomposition with stochastic agricultural prices for b = 25. The computations are based on the robust MPC approximations to the value functions in time period 1. The valuations for $\widehat{\xi} < \infty$ are computed using the implied uncertainty-adjusted probability measures.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

$\widehat{\xi}$	Prob from low to low	Prob from high to high
∞	0.71	0.83
1	0.71	0.82
0.5	0.72	0.82

Table 22: Representative distorted transition probability (Year 1), b = 25

D Baseline distributions

Equation 7 gave the formula for constructing measurement of site-specific productivities from regression coefficients with lower dimensionality. In what follows, we first outline the regression models used for γ and θ, and then we describe the procedure that we used constructing baseline Bayesian posteriors for the regression coefficients.

D.1 Measuring agricultural productivities

To construct a measurement of the θ's, we take data on Amazonian water basins from MapBiomas [6], which we denote by $b = 1, ..., B$, and match each municipality with the basin that has the largest overlap with the municipality. There are in total 466 municipalities and 70 basin groups. We then fit a random effects regression with probabilistic output. Denoting the basin assigned to municipality m by $b[m]$, we consider the specification below.[7]

$$y_\theta^m = R_\theta^m \cdot \beta_\theta + \nu_\theta^{b[m]} + \epsilon_\theta^m \tag{17}$$

where y_θ^m is $\log$(Slaughter value) and

$$\begin{aligned} R_\theta^m \cdot \beta_\theta \overset{\text{def}}{=} \quad & \beta_\theta^0 + \beta_\theta^1(\text{historical_precip}) + \beta_\theta^2(\text{historical_temp}) + \beta_\theta^3(\text{historical_temp}^2) \\ & + \beta_\theta^4(\text{lat}) + \beta_\theta^5(\text{lat}^2) + \beta_\theta^6 \log(\text{cattleSlaughter_farmGatePrice}) + \beta_\theta^7(\text{distance}) \end{aligned}$$

[6] Sub-basins of the National Water Resources Plan (level 2), by ANA 2006, which is available here: https://plataforma.brasil.mapbiomas.org/

[7] We standardize the regressors prior to the posterior estimation.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

where slaughter value is the value of cattle sold per hectare of pasture area in 2017 (USD/ha), precipitation and temperature are the average annual precipitation (mm) and temperature (degrees Celsius), respectively, for the period of 1970-2000 (Fick and Hijmans, 2017), latitude is the geographical coordinates of the municipality centroids, farm gate price is the price of cattle slaughter (SEAB-PR, 2021), and distance is measured the distance from the municipality to the state capital. Since the area dedicated to agriculture varies substantially across municipalities, we opted for weighting observations by the 2017 pasture area in each municipality.

The inclusion of farm gate prices on the right side of this regression is reasonable because variations in farm gate prices across municipalities mostly reflect unobserved costs to bring cattle to stockyards and meat to markets such as proximity to roads or rivers, which are not fully controlled by our geographical variables.

D.2 Measuring carbon absorption potential

To calculate the average CO_2e density (MG/ha) for each site, we first take the set of large (270km $\times$ 270km) sites, denoted by $k = 1, ..., K$, and match each small site i (67.5km $\times$ 67.5km) to the large site with centroid closest to the centroid of site i. In this case, there are 1043 small sites and 78 large sites. Writing $k[i]$ for the large site matched to i, we fit the following random effects regression specification:[8]

$$y_\gamma^i = R_\gamma^i \cdot \beta_\gamma + \nu_\gamma^{k[i]} + \varepsilon_\gamma^i$$

where y_γ^i is $\log(\text{CO2e_ha})$ and

$$\begin{aligned} R_\gamma^i \cdot \beta_\gamma \stackrel{\text{def}}{=} \quad & \beta_0^\gamma + \beta_1^\gamma \log(\text{historical_precip}) + \beta_2^\gamma \log(\text{historical_temp}) + \\ & \beta_3^\gamma \text{latitude} + \beta_4^\gamma \text{longitude} + \beta_5^\gamma (\text{latitude} \times \text{longitude}). \end{aligned}$$

D.3 Baseline posterior estimation

To estimate the benchmark posterior distribution π, we consider the random effects regression models for θ and γ separately. Below, we present the general specification used for these regression models, noting throughout the differences specific to the θ and γ cases.

[8] We again standardize the regressors.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Given observations $i = 1, \ldots, N$ and groups $b = 1, \ldots, B$, we model both carbon accumulation and agricultural productivity using the following specification:

$$Y = R\beta + Z\nu + U \tag{18}$$

where the observed outcome Y is N-dimensional, the vector of coefficients β is K dimensional, the group indicator matrix Z is $N \times B$ dimensional, and the vector of random coefficients ν is B-dimensional. We model the residual vector U as independent but heterogeneous, and so

$$U \sim \text{Normal}\left(0, \frac{1}{\eta}\Sigma\right),$$

where Σ is a prespecified diagonal matrix. Let the weighting matrix W be such that $WW' = \Sigma^{-1}$. We transform (18) by pre-multiplying all terms in (18) by W. With this transformation, we have the following distribution for WY conditioned on regressors and parameters:

$$WY|W, R, Z, \beta, \nu, \eta \sim \text{Normal}\left(WR\beta + WZ\nu,\ \frac{1}{\eta}I\right).$$

We impose normal priors on both vectors of coefficients including the random effects, with precision parameters η and ζ, respectively:

$$\begin{bmatrix}\beta \\ \nu\end{bmatrix} \mid \eta, \zeta \sim \text{Normal}\left(\begin{bmatrix}0 \\ 0\end{bmatrix}, \begin{bmatrix}\frac{1}{\eta}\Omega^{-1} & 0 \\ 0 & \frac{1}{\zeta}I\end{bmatrix}\right)$$

This implies the well known posterior distribution for

$$\begin{bmatrix}\beta \\ \nu\end{bmatrix} \mid R, Z, \eta, \zeta$$

as normal with mean

$$\begin{bmatrix}\eta R'\Sigma^{-1}R & \eta R'\Sigma^{-1}Z \\ \eta Z'\Sigma^{-1}R & \eta Z'\Sigma^{-1}Z + \zeta I\end{bmatrix}^{-1} \begin{bmatrix}\eta R'\Sigma^{-1}Y \\ \eta Z'\Sigma^{-1}Y\end{bmatrix} \tag{19}$$

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication. The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

and covariance matrix

$$\begin{bmatrix} \eta R'\Sigma^{-1}R & \eta R'\Sigma^{-1}Z \\ \eta Z'\Sigma^{-1}R & \eta Z'\Sigma^{-1}Z + \zeta I \end{bmatrix}^{-1} \tag{20}$$

where we take limits as Ω tends to zero.

Finally, we impose improper uniform priors on $\log(\eta)$ and $\log(\zeta)$ and thus the implied improper prior for η and ζ is:

$$\eta^{-1}\zeta^{-1}$$

Incorporating this, gives the implied logarithm of the density for η and ζ.

$$\begin{aligned} &\frac{N-2}{2}\log\eta + \frac{K}{2}\log\eta + \frac{B-2}{2}\log\zeta - \frac{1}{2}\log\det\begin{bmatrix} \eta R'\Sigma^{-1}R & \eta R'\Sigma^{-1}Z \\ \eta Z'\Sigma^{-1}X & \eta Z'Z + \zeta\mathbb{I} \end{bmatrix} \\ &-\frac{\eta}{2}Y'\Sigma^{-1}Y + \frac{1}{2}\begin{bmatrix} \eta Y'\Sigma^{-1}R & \eta Y'\Sigma^{-1}Z \end{bmatrix}\begin{bmatrix} \eta R'\Sigma^{-1}R & \eta R'\Sigma^{-1}Z \\ \eta Z'\Sigma^{-1}R & \eta Z'\Sigma^{-1}Z + \zeta\mathbb{I} \end{bmatrix}^{-1}\begin{bmatrix} \eta R'\Sigma^{-1}Y \\ \eta Z'\Sigma^{-1}Y \end{bmatrix} \end{aligned} \tag{21}$$

One way we produce the baseline posterior probabilities for the parameters of interest is first to random sample of (η, ζ)'s numerically using the exponential the expression in (21). Next draw a corresponding sample of regressions coefficients using a normal distribution with mean (19) and covariance matrix (20). Appendix E.1, which includes an adjustment for ambiguity aversion, provides a second way in the limiting case in which $\xi = \infty$.

In the following tables, we present quantiles for the posterior distributions described above:

Table 23: Quantiles for θ posterior estimation

	β_0^θ	β_1^θ	β_2^θ	β_3^θ	β_4^θ	β_5^θ	β_6^θ	β_7^θ
10%	3.965	-0.691	-0.571	-0.606	-3.433	-0.186	-0.074	0.447
50%	4.046	-0.512	-0.411	1.330	-1.505	-0.121	-0.026	0.490
90%	4.126	-0.325	-0.242	3.258	0.452	-0.054	0.020	0.533

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Table 24: Quantiles for γ posterior estimation

	β_0^γ	β_1^γ	β_2^γ	β_3^γ	β_4^γ	β_5^γ
10%	6.201	0.011	-0.072	0.069	-0.219	0.120
50%	6.228	0.025	-0.064	0.093	-0.193	0.144
90%	6.259	0.040	-0.055	0.119	-0.167	0.167

Table 25: Quantiles for σ posterior estimation

	σ_η^γ	σ_ζ^γ	σ_η^θ	σ_ζ^θ
10%	0.122	0.179	0.372	0.293
50%	0.125	0.199	0.389	0.369
90%	0.129	0.223	0.408	0.459

In Figure 22 we report Bayesian R-squared for the regression equations, computed following the methodology of Gelman et al. (2019).

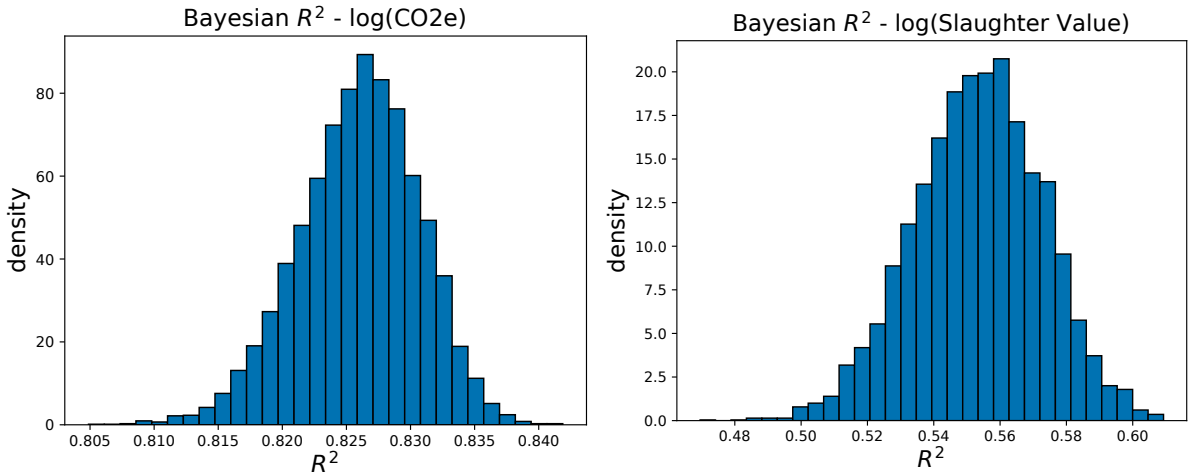


Figure 22: The densities of Bayesian R^2 for the regression equations we use to recover the γ^i's and θ^i.

In what follows, let $\log \mathcal{L}(Y \mid R, Z, \rho)$ denote the log-likelihood function, $\pi(\rho)$ the baseline prior density.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

E Solving the planners problem

We solve the robustly optimal planner's problem by iterating between computing action d by maximizing and an ambiguity-adjusted distribution over unknown productivity parameters by minimizing subject to penalization.

What matters for the maximization step is the mean of the productivity parameters under the ambiguity-adjusted distribution for ρ. Let $\bar{\varphi}^*$ denote this mean as computed using the robustly optimal decision, d^*.

To find the robustly optimal, we initialize the algorithm by setting $\bar{\varphi}_{(0)}$ as the mean of φ implied by the baseline distribution π over ρ and setting $d_{(0)}$ as the maximized solution taking this choice of $\bar{\varphi}_{(0)}$ as given. We then iterate as follows;

1. Given $\bar{\varphi}_s$, solve the planner's problem for decision vector $d_{(s)}$.
2. Given $d_{(s)}$, compute the ambiguity-adjusted distribution over ρ, and form the mean, $\tilde{\varphi}_s$, of the implied ambiguity-adjusted distribution over φ. Form:

$$\bar{\varphi}_{s+1} = .75\bar{\varphi}_s + .25\tilde{\varphi}_s$$

3. If $||\bar{\varphi}_{s+1} - \bar{\varphi}_s||_\infty < 0.005$, stop. Otherwise return to step one with a new $\bar{\varphi}_{s+1}$.

E.1 Computing ambiguity-adjusted distributions

We approximate the distributions of interest using Hamiltonian methods. These methods are typically used to compute Bayesian posteriors numerically. We extend these methods to compute ambiguity-adjusted probabilities. Recall that the productivity vector, φ is presumed to be a function of ρ given by equation (9). We treat $-\frac{1}{\xi} f\left[d, \Phi(\rho)\right]$ as an additional contribution to the log-likelihood function of the parameters, even though it is mathematically distinct. Such a treatment will give us uncertainty-adjusted probability distribution depicted as a "Bayesian posterior" and allow us to use existing computational algorithms designed for such problems.

To construct the Hamiltonian simulations, we form the potential energy term $\mathcal{U}$:

$$\mathcal{U}(\rho) = \left(\frac{1}{\xi}\right) f\left[d, \Phi(\rho)\right] - \log \mathcal{L}(Y|R, Z, \rho) - \log \pi(\rho). \tag{22}$$

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

so we can rewrite potential energy into objective contribution, log likelihood and log prior density (up to a constant).

HMC relies on an auxiliary momentum vector ω of the same dimension as ρ, where $\omega \sim \mathcal{N}(0, M)$ and M is a symmetric, positive-definite mass matrix. The Hamiltonian is then defined as:

$$\mathcal{H}(\rho, \omega) := \mathcal{U}(\rho) + \frac{1}{2}\omega' M^{-1}\omega \tag{23}$$

The HMC algorithm then consists of:

1. Initialize $\rho_{(0)}$.
2. Sample momentum $\omega_{(0)} \sim N(0, M)$.
3. Generate a state proposal $(\tilde{\rho}_{(0)}, \tilde{\omega}_{(0)})$ by evolving its position according to Hamilton's equations, using the leapfrog integrator with step size ϵ and a number of steps L:

$$\frac{d\rho}{dt} = \frac{\partial \mathcal{H}}{\partial \omega} \tag{24}$$

$$\frac{d\omega}{dt} = -\frac{\partial \mathcal{H}}{\partial \rho} \tag{25}$$

4. Perform a Metropolis test to accept or reject the state update $(\rho_{(1)}, \omega_{(1)}) \leftarrow (\tilde{\rho}_{(0)}, \tilde{\omega}_{(0)})$, with the acceptance probability given by:

$$\min\left\{1,\ \exp\left(\mathcal{H}(\rho_{(0)}, \omega_{(0)}) - \mathcal{H}(\tilde{\rho}_{(0)}, \tilde{\omega}_{(0)})\right)\right\}$$

5. Repeat steps 2-4 until to generate a total of 4000 samples by running HMC simultaneously across 4 independent chains, each producing 1000 samples with 500 burn-in-samples per chain.

Notice that these computations take as given a contingent decision theory sequence. To address this, we iterate as follows:

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

- given a contingent decision process, compute the implied uncertainty-adjusted means of the productivities and update with a weighted average of the previous mean and the new one with weights $.75$ and $.25$, respectively;
- given the uncertainty-adjusted means of the productivities, compute a new contingent decision process

We repeat this until the maximum percentage change in the uncertainty-adjusted mean of productivities across all sites falls below 0.005. Then we generate a final sample with 16000 points.

E.2 Computational implementation

We rely on the Stan software (Carpenter et al., 2017, Stan Development Team, 2023) for high-performance statistical computation. The Stan implementation for HMC makes a few adaptations to the algorithm described above to improve computation speed and sampling efficiency. We summarize these below:

- To ensure convergence onto the stationary target distribution, Stan discards the pre-specified number of burn-in samples at the start of the sampling process.
- Stan utilizes the No U-turn sampling (NUTS) variant of HMC, which adaptively determines the number of leapfrog steps L at each iteration to avoid U-turns in the state trajectory (Hoffman and Gelman, 2014, Betancourt, 2016).
- Stan determines the leapfrog step size ϵ using the dual averaging Nesterov algorithm (Nesterov, 2009).
- By default, Stan utilizes a diagonal matrix for M which is estimated using the burn-in samples collected at the start of the algorithm.
- Stan uses reverse-mode automatic differentiation to compute the Hamiltonian gradient.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

References

Araújo, Rafael, Francisco Costa, and Marcelo Sant'Anna. 2025. "Efficient Conservation of the Brazilian Amazon: Estimates from a Dynamic Model." *The Review of Economic Studies*. Advanced Access Publication.

Benini, Rubens de Miranda, and Sérgio Adeodato. 2017. *Economia da Restauração Florestal.* São Paulo, Brazil:The Nature Conservancy.

Betancourt, Michael. 2016. "Identifying the Optimal Integration Time in Hamiltonian Monte Carlo."

Carpenter, Bob, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus A Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. "Stan: A probabilistic programming language." *Journal of statistical software*, 76.

Fick, Stephen E, and Robert J Hijmans. 2017. "WorldClim 2: new 1-km spatial resolution climate surfaces for global land areas." *International journal of climatology*, 37(12): 4302–4315.

Gelman, Andrew, Ben Goodrich, Jonah Gabry, and Aki Vehtari. 2019. "R-squared for Bayesian regression models." *The American Statistician*.

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.

Hoffman, Matthew D., and Andrew Gelman. 2014. "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo." *Journal of Machine Learning Research*, 15(47): 1593–1623.

IBGE. 2017. "Censo Agropecuário: Tabelas 6882, 6911."

Nesterov, Yurii. 2009. "Primal-dual subgradient methods for convex problems." *Mathematical programming*, 120(1): 221–259.

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."

SEAB-PR. 2021. "Preço Médio - Recebido Pelo Agricultor: Boi Gordo, Arroz (Em Casca), Cana-de-Açúcar, Milho, Mandioca, 1990-2021."

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Souza, Carlos M., Julia Z. Shimbo, Marcos R. Rosa, Leandro L. Parente, Ane A. Alencar, Bernardo F. T. Rudorff, Heinrich Hasenack, Marcelo Matsumoto, Laerte G. Ferreira, Pedro W. M. Souza-Filho, Sergio W. de Oliveira, Washington F. Rocha, Antônio V. Fonseca, Camila B. Marques, Cesar G. Diniz, Diego Costa, Dyeden Monteiro, Eduardo R. Rosa, Eduardo Vélez-Martin, Eliseu J. Weber, Felipe E. B. Lenti, Fernando F. Paternost, Frans G. C. Pareyn, João V. Siqueira, José L. Viera, Luiz C. Ferreira Neto, Marciano M. Saraiva, Marcio H. Sales, Moises P. G. Salgado, Rodrigo Vasconcelos, Soltan Galano, Vinicius V. Mesquita, and Tasso Azevedo. 2020. "Reconstructing Three Decades of Land Use and Land Cover Changes in Brazilian Biomes with Landsat Archive and Earth Engine." *Remote Sensing*, 12(17).

Souza Jr, Carlos M, Julia Z. Shimbo, Marcos R Rosa, Leandro L Parente, Ane A. Alencar, Bernardo FT Rudorff, Heinrich Hasenack, Marcelo Matsumoto, Laerte G. Ferreira, Pedro WM Souza-Filho, et al. 2020. "Reconstructing three decades of land use and land cover changes in brazilian biomes with landsat archive and earth engine." *Remote Sensing*, 12(17): 2735.

Stan Development Team. 2023. "Stan Modeling Language Users Guide and Reference Manual." Version 2.33.

Other (e.g., Data Files)
Click here to access/download;Other (e.g., Data Files);Figure captions and alt text.pdf
This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Figure captions and alt text

Carbon prices, forest conservation and reforestation in the Brazilian Amazon

Journal of Political Economy | MS 20250778R2

Main manuscript: Figures 1–14

Figure 1

Caption: 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).

Alt text: Log-scale scatterplot of emissions per capita against GDP per capita. Most countries form an upward-sloping cluster; the Brazilian Amazon lies well above countries with similar income and above the United States, European Union, China, and India.

Figure 2

Caption: Percentage of maximum carbon density: actual vs. theoretical. The black solid curve gives the an empirically-based construction of the carbon recovery fraction as a function of the recovery time. The blue dashed curve plots $1-\exp\left(-\alpha\, t\right)$.

Alt text: Line chart of the fraction of maximum carbon density reached against age of regrowing forest, where a solid data line rises steeply and concavely from one-third at year zero to near nine-tenths by year thirty, staying above a lower, more gradual dashed theoretical line throughout.

Figure 3

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

Alt text: Two maps of 1,043 Amazon sites. Agricultural area is concentrated along the southern and eastern edges, while stored carbon is greatest in western, northern, and central areas.

Figure 4

Caption: Carbon sequestration parameters and Agricultural productivity heterogeneity

Alt text: Two maps of site-level productivity parameters. Carbon sequestration capacity is generally higher in western and central sites, while agricultural productivity is concentrated along the southern and eastern margins.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Figure 5

Caption: Agricultural area and cumulative CO_2 capture.

Alt text: Two line charts over fifty years for transfer prices of zero, fifteen, and twenty-five dollars. Agricultural area rises toward one-fifth of the biome without transfers but falls to near zero with them; cumulative capture turns negative without transfers and climbs steadily with them.

Figure 6

Caption: Evolution of transfer payments for two choices of b.

Alt text: Line chart of net transfer payments over fifty years, where both curves rise sharply to a peak near year eight to ten, then decline gradually. The higher transfer level peaks around twenty-two billion dollars, the lower around eleven billion.

Figure 7

Caption: Agricultural area changes after 30 years.

Alt text: Four Amazon maps compare agricultural shares in 2017 and after thirty years under transfers of zero, ten, and twenty-five dollars. Agriculture expands along the southeast without transfers, contracts sharply at ten dollars, and nearly disappears at twenty-five dollars.

Figure 8

Caption: Agricultural area evolution over time

Alt text: Six Amazon maps track agricultural shares from 2017 to 2047 under a fifteen-dollar transfer. Agriculture retreats rapidly from most sites during the first fifteen years and remains concentrated in a small southern and eastern fringe.

Figure 9

Caption: Ambiguity-adjusted densities for four sites. The left plots report sites with the largest entropies for θ and right plots report sites with the largest entropies for γ. The upper plots are $b=0$ and the lower plots for $b=15$.

Alt text: Four density plots compare baseline and ambiguity-adjusted productivity distributions. The agricultural-productivity densities shift substantially left for both transfer cases; the carbon-sequestration densities change little at zero transfers and shift modestly left at fifteen dollars.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy, published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Figure 10

Caption: Histogram of the years in which one of the two controls is maximal, for $b = 0$ and common $P^{ee} = 6.8$. The controls remain at zero for fifty years in 10.1% of the sites under ambiguity neutral and in 10.2% of sites under ambiguity aversion.

Alt text: Overlaid histogram of the years when deforestation or reforestation is greatest across sites. Both specifications peak near year zero, but ambiguity aversion places somewhat more mass at later positive years and less at earlier negative years.

Figure 11

Caption: Evolution of agricultural area under ambiguity neutrality and ambiguity aversion for $b = 0$ at a common business-as-usual carbon price, $P^{ee} = 6.8$.

Alt text: Line chart of agricultural land share over fifty years at a common carbon price without transfers and two aversion settings. Under ambiguity neutrality the share rises toward twenty percent; under ambiguity aversion it dips slightly and then returns near its initial level.

Figure 12

Panel labels: (a) b=0; (b) b=15

Caption: Histogram of the years in which one of the two controls is maximal, $P^{ee} = 6.8$ for ambiguity neutral and $P^{ee} = 4.8$ for ambiguity averse. For $b = 0$, 10.1% of the sites remain unchanged under ambiguity neutrality, while 13.2% do so under ambiguity aversion. For $b = 15$, 2.0% of the sites do not change under ambiguity neutrality, and 2.1% remain unchanged under ambiguity aversion.

Alt text: Two overlaid histograms compare the timing of peak land-use adjustments. With zero transfers, observations cluster near year zero with long tails; with fifteen dollars, most sites peak within the first several years and differences across ambiguity specifications are small.

Figure 13

Panel labels: (a) b=0; (b) b=15

Caption: Evolution of agricultural area under ambiguity neutrality and ambiguity aversion, using the corresponding shadow prices.

Alt text: Two line charts compare agricultural land shares under ambiguity neutrality and aversion at their respective shadow prices. With zero transfers both rise, and the ambiguity-averse path is higher; with fifteen dollars both fall rapidly and nearly overlap.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Figure 14

Panel labels: (a) b=0; (b) b=15

Caption: Evolution of agricultural area under price uncertainty for $b = 0$ and $b = 15$.

Alt text: Two line charts compare agricultural land shares with and without price-model ambiguity. For zero transfers, both paths rise and nearly overlap; for fifteen dollars, both fall rapidly toward zero with only small differences.

Online Appendix: Figures 15–22

Figure 15

Caption: Agricultural area trajectories for two different discount rates.

Alt text: Line chart of agricultural land shares under two discount rates and transfer prices of zero and twenty-five dollars. Raising the discount rate slightly increases agricultural area without transfers and has almost no effect when transfers are twenty-five dollars.

Figure 16

Panel labels: (a) distinct variances; (b) common variance

Caption: Smoothed probabilities for the two hidden state Markov chain models

Alt text: Two time-series panels plot actual cattle prices and smoothed high-state probabilities from 1995 to 2018. The distinct-variance model switches states several times; the common-variance model remains in the low state mainly from 2004 through 2008.

Figure 17

Panel labels: θ, b=0; γ, b=0; θ, b=15; γ, b=15

Caption: Relative entropy for $\xi = 1$. Sites with the highest relative entropy are circled.

Alt text: Four Amazon maps rank site-level relative entropy for agricultural and carbon-sequestration productivity under zero and fifteen-dollar transfers. High values are spatially heterogeneous; circled maximum-entropy sites lie in the south for three panels and the southeast for one.

Figure 18

Caption: Ambiguity-adjusted densities for four sites, with $\xi = 2$

Alt text: Four density plots compare baseline and ambiguity-adjusted productivity distributions under moderate ambiguity aversion. Agricultural-productivity densities shift left in both transfer cases, while carbon-sequestration densities nearly overlap the baseline.

This is not a published article. This is an accepted manuscript, without copyediting, corrections, formatting, or online data files, and before publication.
The completed version of record is expected to be published with DOI https://doi.org/10.1086/744117 in an upcoming issue of Journal of Political Economy,
published by The University of Chicago Press. Copyright 2026 The University of Chicago.

Figure 19

Caption: Ambiguity-adjusted densities for four sites, with $\xi = 0.5$

Alt text: Four density plots compare baseline and ambiguity-adjusted productivity distributions under strong ambiguity aversion. Agricultural-productivity densities shift markedly left, and the carbon-sequestration density at fifteen dollars also shifts left; adjustment is much larger than under weaker aversion.

Figure 20

Caption: Agricultural area evolution over time with ambiguity aversion, $\xi = 1$

Alt text: Six Amazon maps track agricultural shares from 2017 to 2047 under a fifteen-dollar transfer and moderate ambiguity aversion. Agriculture contracts rapidly and becomes concentrated in small southern and eastern areas.

Figure 21

Caption: Agricultural area evolution over time with ambiguity aversion, $\xi = 0.5$

Alt text: Six Amazon maps track agricultural shares from 2017 to 2047 under a fifteen-dollar transfer and stronger ambiguity aversion. Agriculture contracts rapidly but remains slightly more extensive than under moderate ambiguity aversion.

Figure 22

Caption: The densities of Bayesian R^2 for the regression equations we use to recover the γ^i's and θ^i.

Alt text: Two histograms show posterior Bayesian R-squared values. The carbon-sequestration regression is tightly centered near 0.827, while the agricultural-productivity regression is centered near 0.55 and is more dispersed.