The Political Economy of Conservation

Robin Burgess
LSE

Francisco Costa
FGV EPGE

Allan Hsiao
Stanford

Benjamin A. Olken
MIT

Veronica Salazar Restrepo
IIES

March 25, 2026

Abstract

Conservation is political. Forest protections safeguard our natural forest resources, but they also provoke resistance from agricultural producers who profit from clearing land. We study this tension in the Brazilian Amazon. Using detailed data on campaign donations and candidates in state and federal elections from 2002 to 2018, we show that environmental reforms led agricultural producers to sharply increase political support for pro-agricultural politicians. Donations from agricultural firm owners roughly doubled, and agricultural candidates gained 70 to 90 percent more votes. Motivated by this evidence, we develop a model of optimal regulation under the threat of political repeal. Regulation that imposes large losses on producers provokes resistance, which leads to repeal that can undo the welfare gains from conservation. We estimate the model with spatial data on land use and agricultural returns across the Brazilian Amazon from 2001 to 2019. Pigouvian taxation maximizes welfare in isolation, but it triggers certain repeal and thus achieves little in expectation. Optimal policy softens taxation and targets sites where conservation gains are large relative to producer losses. This targeted approach achieves substantial feasible welfare gains where Pigouvian taxation cannot.

1 Introduction

Agricultural production drives deforestation that accounts for one fifth of global carbon emissions. This negative externality motivates efforts to strengthen forest regulation, but such regulation faces political headwinds. Large agricultural producers derive sizeable profits from deforestation and thus face strong incentives to lobby against regulation. How can we design conservation policy that accounts for this political reality?

We tackle this question in the Brazilian Amazon. In 2004, the Brazilian government launched the Action Plan for the Prevention and Control of Deforestation in the Legal Amazon (PPCDAm), which strengthened enforcement of existing forest protections through satellite monitoring, expanded protected areas, and tougher penalties for illegal clearing. Deforestation fell sharply. We compile detailed data on individual campaign donors and candidates in state and federal elections from 2002 to 2018, linking them to firm-level records from the Brazilian tax authority to identify agricultural sector participants.

We use a difference-in-differences strategy that compares agricultural donors and candidates to their non-agricultural counterparts before and after the implementation of PPCDAm. The identifying assumption is that, absent the reform, political contributions by agricultural firm owners would have evolved in parallel with those of other firm owners. We support this assumption with pre-treatment event studies and present robustness checks using matched samples, as well as triple-difference specifications that exploit variation in enforcement intensity across municipalities and that compare the Legal Amazon to the rest of Brazil.

We find that PPCDAm led to large increases in donations from agricultural donors, particularly to candidates with ties to the agricultural sector. Donations from agricultural firm owners roughly doubled relative to pre-reform levels. Triple-difference estimates confirm that the increase was concentrated where enforcement was strongest. Switching to the candidate side, we document that agricultural candidates received 70 to 90 percent more votes in both state and federal elections, with meaningful positive effects on the probability of being elected to office. These results suggest deliberate efforts to back politicians who oppose environmental regulation. The agricultural sector met regulation with resistance.

Such resistance creates a trade-off for the regulator. We characterize this trade-off

with a simple theory framework that emphasizes the incidence of regulation and the threat of repeal. The typical theoretical approach is instead to maximize aggregate welfare, while addressing distributional impacts with lump-sum transfers. We document political resistance to regulation that is in conflict with the idea that transfers can navigate politics. We take seriously that regulation imposes large losses on agricultural producers, provoking resistance that risks undoing the welfare gains from conservation. The regulator must therefore balance the environmental benefits of strong regulation against the risk of repeal. We show that optimal regulation falls short of the Pigouvian benchmark, and that precise targeting can ease the trade-off. Regulation that targets sites with high emissions and low agricultural values achieves large welfare gains with limited producer losses, reducing the risk of repeal. By contrast, blunt regulation that burdens all producers, including those on efficient land, invites resistance without meaningful environmental benefit.

We study these forces quantitatively with a simple empirical model of agricultural land use in the Brazilian Amazon. We compile spatial data on agricultural production and suitability across the Amazon for the period from 2001 to 2019. The data reveal substantial heterogeneity in both agricultural returns and carbon stocks. We estimate a logit model of land allocation that recovers the responsiveness of the agricultural frontier to changes in profitability. Agricultural returns combine spatial variation in agricultural suitability with time variation in global commodity prices. Estimation reduces to a Poisson regression that we estimate in pooled long differences over 10-year horizons, as well as in annual panel regressions. Our baseline analysis relies on the long-difference specification, which isolates longer-term, decadal responses. By contrast, the panel regression isolates short-run responses to year-on-year fluctuations, and it delivers estimates that are substantially smaller. We use the estimated model to simulate counterfactual regulation, computing welfare, producer surplus, emissions, and the risk of political repeal under alternative policy instruments.

Our counterfactual analysis yields three main findings. First, Pigouvian taxation maximizes welfare in the absence of political resistance, but it also imposes large burdens on producers. This heavy burden prompts political resistance that can trigger repeal. In expectation, Pigouvian taxation achieves little when strong political resistance guarantees repeal. Second, optimal policy softens taxation and targets regulation toward sites where emissions reductions are large relative to producer losses. This two-dimensional targeting – on both emissions and profits – achieves substantial

feasible welfare gains where Pigouvian taxation cannot. Third, policies constrained to target emissions alone, including regional taxes and protected areas, can still effectively navigate the trade-off between welfare and repeal. These constrained policies dominate Pigouvian taxation, particularly when the value of abatement is high, even though they fall short of optimal policy. The gap between optimal and constrained policy reflects the informational challenge of targeting on profits, which are typically the private information of the firm.

This paper contributes to work at the intersection of political economy and the environment (Oates and Portney 2003, Besley and Persson 2023). A growing literature studies how political forces shape climate and environmental outcomes, including climate attitudes and public opinion (Dechezleprêtre et al. 2025, Hsiao and Kuipers 2025), environmental issues in politics (Bombardini et al. 2025, Longuet-Marx 2025), climate adaptation (Hsiao 2025), green investment (Kashner 2026), policy lobbying (Kang 2016), and regulation of the commons (Ryan and Sudarshan 2022, Berman and Hickok 2026). We bring this perspective to tropical deforestation, where the tension between economic development and environmental protection is acute and the political constraints on regulation are binding.

A large empirical literature studies tropical deforestation in Brazil and Indonesia, which house two of the world's most important tropical ecosystems.[1] In Brazil, recent work has used a range of empirical methods to study the role of policy in shaping land use (Souza-Rodrigues 2019, Assunção et al. 2024, Leite-Mariante and Restrepo 2024, Araujo et al. 2026, Domínguez-Iino 2026). Several papers evaluate the PPCDAm reform, as well as related enforcement and financial instruments (Nepstad et al. 2009, Hargrave and Kis-Katos 2013, Assunção et al. 2015, 2019, Assunção et al. 2022). In Indonesia, the expansion of palm oil cultivation has sparked forest fires and destroyed peatland (Busch et al. 2015, Balboni et al. 2025, Hsiao 2026). More broadly, agricultural expansion drives extensive deforestation at global scale (Hsiao et al. 2026). We document that regulation induces political resistance, which can potentially reverse the initial gains from forest protection.

A smaller set of papers focuses specifically on the politics of deforestation. Globally, electoral cycles shape deforestation through illegal logging, forest fires, and weakened enforcement (Burgess et al. 2012, Balboni et al. 2021, Sanford 2023, Harding et al. 2024). In Brazil, agricultural interests and political interference have eroded

[1] Recent surveys include Balboni et al. (2023) and Costa et al. (2025a).

conservation gains (Burgess et al. 2025, Costa et al. 2025b, Dahis and Katovich 2025). In this crucial setting, we quantify how politics shape the effects and optimal design of environmental regulation. We emphasize the need to account for political sustainability in evaluating policy impacts. The future of the Amazon depends on durable conservation in the face of opposition.

2 The Brazilian Amazon

We describe our institutional setting and data. Agricultural expansion drives deforestation in the Brazilian Amazon, the world's largest tropical forest. A major reform strengthened the enforcement of forest protections, slowing deforestation. But subsequent political realignment weakened these protections, raising questions about the durability of conservation policy.

2.1 Background

Brazil's Constitution and first Forest Code offer protection against deforestation in the form of the "80 percent rule." According to this rule, private landowners in the Legal Amazon must preserve at least 80 percent of their property as forested land. However, weak enforcement limited the effectiveness of these policies. By the early 2000s, deforestation rates had reached alarming levels, prompting a shift in policy (MMA 2013).

In 2004, the Brazilian government introduced the Action Plan for the Prevention and Control of Deforestation in the Legal Amazon (PPCDAm), a comprehensive initiative that integrated technology investments, regulatory enforcement, and economic incentives. A key innovation was the satellite-based DETER monitoring system, which enabled real-time detection of forest loss and strengthened enforcement by the Brazilian Environmental Agency (IBAMA) (Assunção et al. 2023). Legal reforms also reinforced conservation efforts. Deforestation on untitled land was reclassified as a felony, empowering authorities to seize assets used in illegal activities, and protected areas were expanded (Soares-Filho et al. 2010, Harding et al. 2021). These policies under PPCDAm led to a sharp decline in deforestation, which fell by 70 percent between 2006 and 2012.

As priorities shift and economic pressures build, questions remain around long-

term sustainability (Crouzeilles et al. 2017, Costa et al. 2025b). In 2012, the New Forest Code weakened regulation by granting amnesty to properties that had been illegally deforested before 2008 (Soares-Filho et al. 2014, Freitas et al. 2018). Enforcement agencies also faced budget cuts. From 2013 to 2016, IBAMA's budget fell sharply, and reductions to field staff limited monitoring capacity (CGU 2016). These shifts, coupled with broader political realignment in favor of agricultural expansion, coincided with a resurgence in deforestation and the reversal of prior conservation gains (Rochedo et al. 2018, Burgess et al. 2025).

2.2 Election and campaign data

We measure political activity in the form of campaign donations. Our primary data source is the universe of formal donations recorded in all legislative, state, and federal elections from 2002 to 2018. We focus on these elections because state and federal governments hold the greatest authority over environmental policy and enforcement. Elections occur every four years. All campaign donations must be registered and publicly disclosed by the Supreme Electoral Court (TSE), ensuring comprehensive coverage of political contributions. We observe donors, candidates, and donation values. The same TSE data record electoral outcomes: total votes and election outcomes for each candidate at the municipality and election-year level.

We leverage firm-level data from the Brazilian tax authority (*Receita Federal*) to identify agricultural donors and candidates, who are most likely to oppose environmental regulations. These data include the full registry of formal firms and their owners. We use this information to construct a panel of all donors who once owned a firm in Brazil, and we link firm owners to their campaign contributions and candidate affiliations. We classify agricultural sector donors as individuals with firms registered in the agricultural sector. We supplement these data by building a panel of candidates who received donations from firm owners, distinguishing between those funded by donors associated with the agricultural sector and those funded by donors from other sectors. Appendix A provides details on the construction of these panel datasets.

We restrict our analysis to donations made in elections within the Legal Amazon, ensuring that our estimates reflect political financing patterns relevant to environmental policymaking in the region. Table 1 presents the number of individual donors

Table 1: Individual donors and candidates

	Individuals			Donation to/from	
	Total	Agricultural firm owners	Infractors	Agricultural firm owners	Infractors
	(1)	(2)	(3)	(4)	(5)
Donors	62,987	3,509	3,564	1,488	530
Candidates	5,605	390	274	1,786	1,716

This table describes individual donors and candidates in state and federal elections in the Legal Amazon from 2002 to 2018. Columns 1 to 3 count distinct individuals: all donors or candidates (1), those who own agricultural firms (2), and those with a history of IBAMA environmental infractions (3). Columns 4 and 5 count individuals who gave or received at least one donation to or from agricultural firm owners (4) or infractors (5). The data come from the Supreme Electoral Court (TSE) for donations and from the Brazilian tax authority (*Receita Federal*) for firm ownership.

and candidates in our sample. Column 1 reports the total number of distinct individuals: 62,987 donors and 5,605 candidates. Of these, column 2 shows that 3,509 donors and 390 candidates own agricultural firms. Furthermore, column 4 reports the number of individuals who gave (or received) at least one donation to (or from) agricultural firm owners. We observe that 1,488 donors gave donations to candidates who owned agricultural firms, while 1,786 candidates received donations from donors who owned agricultural firms.

2.3 Enforcement and deforestation data

We also incorporate data on environmental violations from IBAMA, Brazil's environmental enforcement agency, to identify donors and candidates with a history of regulatory infractions. We obtain individual-level infraction and embargo records. Each record identifies the recipient using a partially masked CPF or CNPJ, along with the violation date and municipality. We link these records to donors and candidates to construct infraction histories. Columns 3 and 5 of Table 1 report the number of donors and candidates with at least one infraction. Column 3 shows that 3,564 donors and 274 candidates have a history of IBAMA environmental infractions. Column 5 shows that 530 donors gave donations to candidates with a history of IBAMA infractions, while 1,716 candidates received donations from donors with a history of IBAMA infractions.

Turning to regulatory outcomes, the Brazilian National Space Research Institute (INPE) produces annual deforestation estimates through the PRODES program. These data draw on Landsat satellite imagery to detect and measure forest loss across the Legal Amazon. We aggregate the PRODES data to the municipality level for 2000 to 2018. We use this annual deforestation data as a descriptive measure of the impacts of regulation over time, including as it evolves from 2004 to 2006 during the implementation of the PPCDAm reform. For each municipality, the reduction in deforestation during this period serves as a proxy for the extent of exposure to the reform.

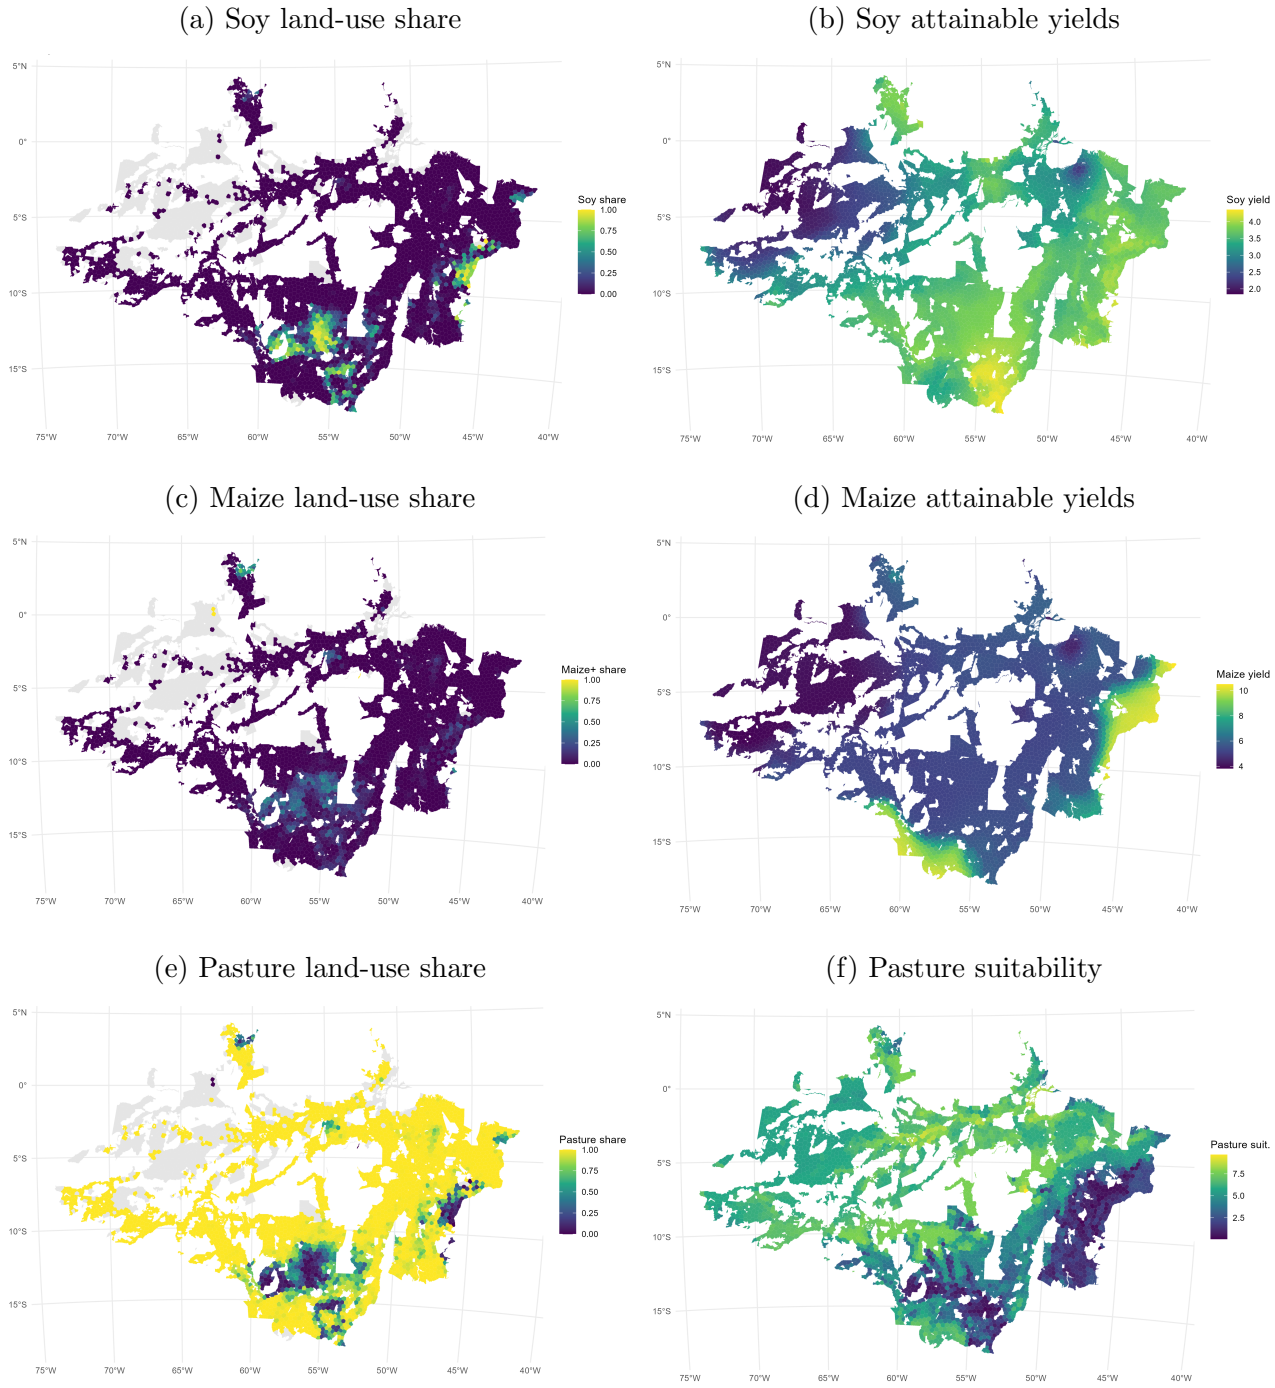
Finally, we use cloud cover data from Assunção et al. (2023) as a measure of exogenous variation in enforcement intensity. The DETER satellite system, which guides IBAMA field operations, cannot detect deforestation beneath persistent cloud cover. We use data on mean municipal cloud cover to capture this plausibly exogenous variation in monitoring capacity across municipalities. Municipalities with high cloud cover are difficult for IBAMA to monitor and thus have lower exposure to the PPCDAm reform. Municipalities with lower cloud cover are easier to monitor and thus more exposed to the reform.

2.4 Land-use data

To characterize optimal policy, we compile spatial data on agricultural land use and suitability in the Brazilian Amazon. The data form a spatial panel covering the period from 2008 to 2017. The unit of observation is a 0.25-degree cell, which corresponds to roughly 900 km^2 of land area. The underlying data come from Araujo et al. (2026) and include land allocations to soy, other crops, pasture, and forest from MapBiomas; attainable crop yields for soy and maize in tons per hectare; agro-climatic suitability for pasture; and distance to ports. We also observe carbon stocks, measured in tons of CO_2 per hectare of land.

Figure 1 presents the data. In the left column, we map land-use shares across agricultural activities. Soy cultivation is concentrated in a few high-yield clusters in the southeast, particularly in Mato Grosso and along the border with the Cerrado biome. Maize and other crops are concentrated in the same regions but account for a smaller share of agricultural land. Pasture dominates agricultural land use throughout the region, consistent with the extensive cattle ranching that characterizes

Figure 1: Agricultural land use and suitability (2008)



Maps of the Brazilian Amazon at the 0.25-degree cell level for 2008. The left column shows the share of agricultural land allocated to each activity. The right column shows attainable crop yields for soy (tons/ha) and maize (tons/ha), and agro-climatic suitability for pasture. Land-use data come from MapBiomas; yield and suitability data come from Araujo et al. (2026).

the Amazon frontier. In the right column, we map crop yields and pasture suitability. Soy and maize yields are highest in the southern and eastern portions of the region, particularly along the "arc of deforestation" where agricultural frontier expansion is concentrated. Pasture suitability is more widespread but still heterogeneous, with the highest values in the south and east.

3 Political Resistance

We study the political responses to environmental regulation. Stronger enforcement of forest regulations prompted increased campaign donations from agricultural producers. Politicians tied to the agricultural sector raised more funds and won more votes.

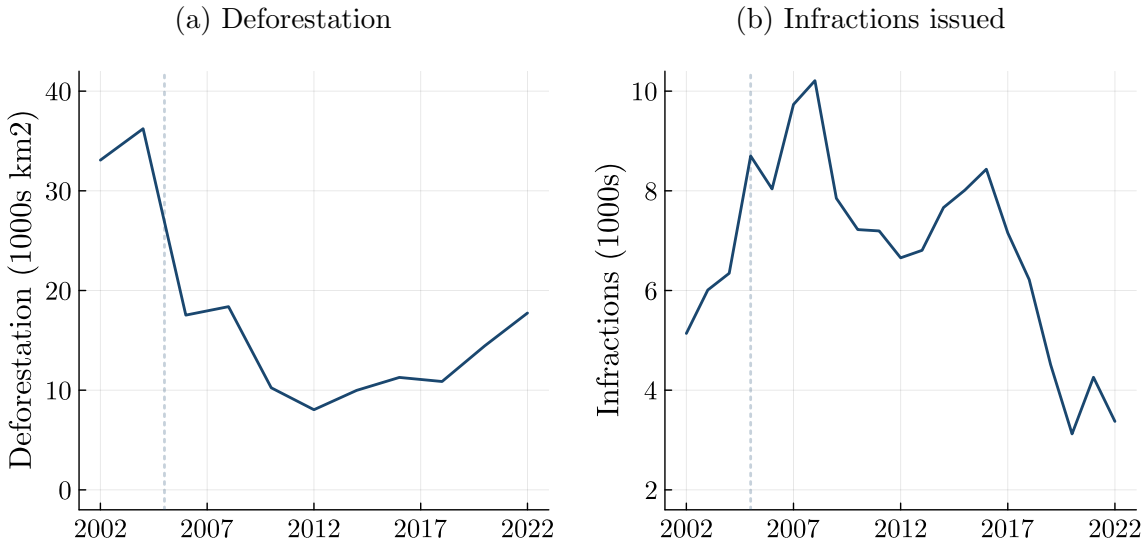
3.1 Descriptive evidence

The PPCDAm program strengthened Amazonian protections, increasing enforcement of existing policies and expanding the reach of protected areas. We see the impacts in the data. First, we validate that PPCDAm increased enforcement. Figure 2 shows time series for deforestation and environmental infractions. Deforestation falls after the start of PPCDAm implementation in 2005, while infractions increase. Each pattern is consistent with tighter enforcement of environmental regulation under PPCDAm. Furthermore, we observe an increase in deforestation and a decline in infractions after 2018, which marked a return to lower levels of enforcement.

Second, we document political resistance to PPCDAm regulation. Figure 3 plots average campaign donations for donors and candidates from 2002 to 2018 in the Legal Amazon. We include donations to candidates for both state congress (*Deputado Estadual*) and federal congress (*Deputado Federal*). We classify both donors and candidates as agricultural or non-agricultural. To do so, we identified the firms owned by donors and candidates, and we classify these firms according to their CNAE sector codes. We define agricultural firms as those connected to agriculture or food processing, which includes meatpacking, although our results are robust to excluding food processors.

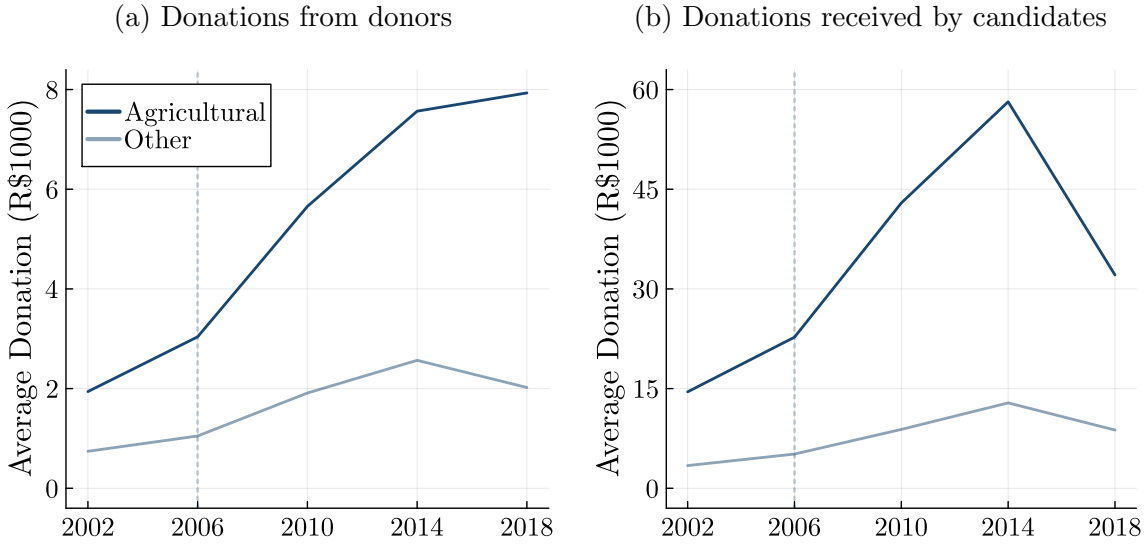
Figure 3a contrasts donations from agricultural and non-agricultural donors. We find that contributions from agricultural donors rise sharply after 2006, relative to

Figure 2: Deforestation and infractions over time



Annual totals for deforestation (left) and environmental infractions (right). The infraction year refers to when authorities formally issued the notice, which may lag the underlying violation. The data come from PRODES for deforestation and from IBAMA for infractions.

Figure 3: Campaign donations over time



Average campaign donations by election year, comparing individuals connected to the agricultural sector (navy) and those connected to other sectors (light blue). Panel (a) shows donations made by donors. Panel (b) shows donations received by candidates. The vertical lines mark 2006, the first election when PPCDAm was fully implemented. The data come from TSE for donations and from *Receita Federal* for sector classification.

those from non-agricultural donors. Figure 3b shows the same pattern on the recipient side. Candidates who own agricultural firms receive a large increase in donations after 2006, relative to candidates tied to other sectors. This increase is attenuated in 2018, perhaps reflecting campaign-finance reforms that expanded public funding and reduced reliance on private contributions.[2] Overall, these patterns indicate a disproportionate rise in the average funding of candidates with economic ties to the agricultural sector, which was the most direct target of the PPCDAm reform.

3.2 Identification

We estimate difference-in-differences specifications that formalize the visual comparisons in Figure 3. In particular, we compare agricultural and non-agricultural donations before and after the implementation of PPCDAm. The premise is that agricultural producers are most exposed to the PPCDAm reform, which imposed land use regulation that constrained the scope for agricultural expansion. Non-agricultural producers are less exposed to these constraints on agricultural expansion and thus form a control group. Changes in outcomes for non-agricultural producers therefore capture common exposure to macroeconomic, social, and political fluctuations, but not direct exposure to changes in deforestation policy.

We estimate the effect of increased environmental enforcement on political contributions with the following dynamic specification.

$$y_{imt} = \sum_{\tau} \beta_{\tau} A_i \times \mathbb{1}\{t = \tau\} + \gamma_{mt} + \varepsilon_{imt}, \tag{1}$$

where y_{imt} is the average donation made by donor i (or received by candidate i) in municipality m and election year t. The indicator $A_i = 1$ if i is an agricultural firm owner, and $\mathbb{1}\{t = \tau\}$ equals one in election year τ. The coefficients β_{τ} trace differential outcomes for agricultural versus other sectors by year, relative to 2006. We include municipality-by-election-year fixed effects γ_{mt}. Standard errors are clustered at the individual level (Abadie et al. 2023). We also report average post-2006 effects by

[2] In particular, the creation of the public electoral fund allowed parties and candidates to receive state resources for campaigns. This fund reduced reliance on private contributions, including from the agricultural sector.

pooling across election years.

$$y_{imt} = \beta\, A_i \times \mathbb{1}\{t > 2006\} + \alpha_i + \gamma_{mt} + \varepsilon_{imt}, \tag{2}$$

where $\mathbb{1}\{t > 2006\}$ is an indicator for election years after 2006. We consider alternative fixed-effects sets, including specifications without α_i as a robustness check. Our identifying assumption is that, absent PPCDAm, donations by agricultural owners would have followed parallel trends to those of other owners. We assess this assumption by examining pre-treatment coefficients in the event study.

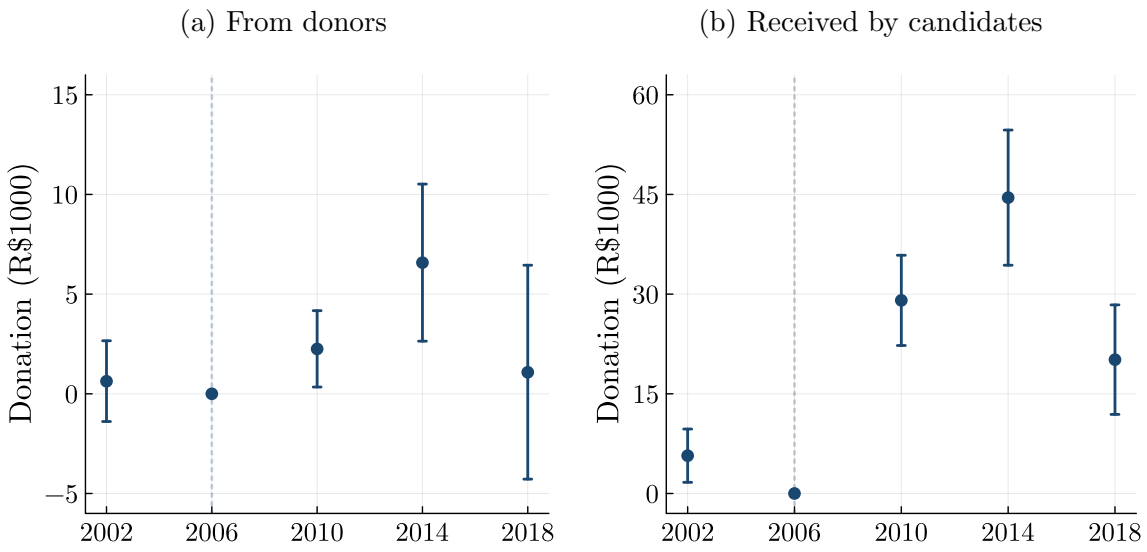
We present alternative specifications that address several potential concerns. First, agricultural firms may differ from other firms in ways correlated with donation behavior. We therefore construct a matched sample, pairing donors (or candidates) by firm characteristics: donations in 2002, firm opening year, participation in the *Simples* tax regime, ownership type (closely held vs. publicly listed), firm size (micro, small, large), and terciles of equity capital. Appendix Table A1 shows that the matched sample indeed improves the balance between agricultural and non-agricultural firms on observed characteristics. We show robustness to matched specifications, which extend equations 1 and 2 to include match-pair fixed effects.

Second, our period of study coincides with a commodity boom and increased competition from China. During this time, agriculture benefited from trade while manufacturing struggled with foreign competition (Costa et al. 2016). We show robustness to triple-difference specifications of the form

$$y_{imt} = \beta_0\, A_i \times \mathbb{1}\{t > 2006\} + \beta_1\, A_i \times \mathbb{1}\{t > 2006\} \times H_m + \alpha_i + \gamma_{mt} + \varepsilon_{imt}, \tag{3}$$

where H_m captures municipal heterogeneity. We can broaden our scope with a national sample that defines H_m as an indicator for municipalities in the Legal Amazon. The triple-difference specification then compares effects inside vs. outside the Legal Amazon. If rising agricultural incomes nationally were the main driver of our difference-in-difference effects, then our triple-difference estimates would be attenuated because we would observe similar effects outside the Amazon. We can also narrow our scope by defining H_m as cloud cover, which affects enforcement intensity by hindering satellite monitoring (Assunção et al. 2023). The triple-difference specification then compares effects in municipalities with strong vs. weak enforce-

Figure 4: Campaign donations, difference-in-differences



Event-study coefficients β_τ from equation (1) with 2006 as the reference year. Points are estimates, and bars are 95% CIs with individual-level clustering. Panel (a) shows the differential effect on average donations from agricultural firm owners. Panel (b) shows the differential effect on average donations received by candidates who are agricultural firm owners. Estimates based on the matched sample with match-pair and election-year fixed effects. See Appendix Figure A1 for results based on the complete sample.

ment intensity, thereby isolating the impacts of PPCDAm regulation on our political outcomes of interest.

3.3 Donors

Stronger regulation and enforcement brought about by PPCDAm coincide with higher political giving by agricultural donors and higher receipts for agricultural candidates. Figure 4 shows event-study coefficients from equation (1). For donors, Panel (a) shows that agricultural firm owners increased contributions after 2006 relative to owners in other sectors, with a peak in 2014. For candidates, Panel (b) shows that those with agricultural businesses received markedly more funding after 2006.

In Table 2, Panel A reports the difference-in-differences coefficients from equation (2) that capture the average increase in donations by agricultural firm owners after 2006. Columns vary the fixed effects. Column 1 includes municipality and election-year fixed effects. Column 2 replaces these with municipality-by-year fixed effects, flexibly absorbing local-by-year shocks in political contributions. Column 3

Table 2: Donor donations

	(1)	(2)	(3)	(4)
Panel A: Donations to all candidates				
$Ag \times Post$ 2006	5.70***	5.81***	4.02***	3.24***
	(0.91)	(0.97)	(0.95)	(1.19)
Effect as % of mean	237.5	242.1	167.6	135
Observations	36,670	35,195	35,195	7,100
Fixed effects	$m + y$	my	$i + my$	match + my
Panel B: Donations to agricultural candidates				
$Ag \times Post$ 2006	3.05***	3.27***	2.25***	3.36***
	(0.53)	(0.57)	(0.57)	(0.63)
Effect as % of mean	277.5	297.4	204.5	305.7
Observations	36,670	35,195	35,195	7,100
Fixed effects	$m + y$	my	$i + my$	match + my

Difference-in-differences estimates of equation (2). Outcomes are donations made by donors. *Ag* indicates agricultural firm owners; *Post* 2006 equals one for elections after 2006. Outcomes are average donations to all candidates (Panel A) and those to candidates who own agricultural firms (Panel B). Fixed effects vary by column for municipality m, election year y, individual i, municipality-year my, and match pairs. Standard errors are clustered at the individual level. Monetary values are in R$1000.

adds donor fixed effects to the municipality-by-year fixed effects, absorbing donor-specific traits that influence giving. In this specification, the differential increase in donations, β, is identified from within-donor variation over time. Column 4 turns to the matched subsample, including matched-pair and municipality-by-year fixed effects. These estimates are based on within-pair changes in donations over time. Overall, the estimates show that average donations from agricultural donors rose by roughly 130 to 242 percent relative to pre-PPCDAm averages (R$3,240 to R$5,810), all statistically significant at 1 percent.

In Table 2, Panel B shows effects on campaign contributions directed to candidates who are themselves agricultural firm owners. The magnitudes are even larger: agricultural donors increased donations to agricultural candidates by about 204 to 306 percent relative to the pre-PPCDAm baseline (the 2002 election). These results suggest that donors raised contributions in response to stronger environmental enforcement, seeking to expand the sector's political influence in Congress.

Table 3: Donor donations by enforcement intensity

	Heterogeneity (H) by			
	Deforestation reduction		Cloud coverage	
	(1)	(2)	(3)	(4)
$Ag \times Post$ 2006	3.59***	1.37	11.55***	9.77***
	(0.85)	(0.97)	(2.49)	(2.58)
$Ag \times Post$ 2006 $\times H$	1.10***	0.96***	-15.20***	-16.72***
	(0.34)	(0.30)	(4.49)	(4.80)
Observations	34,275	7,400	27,785	5,770
Fixed effects	my	match + my	my	match + my

Triple-difference estimates of equation (3). Outcomes are donations made by donors. H measures local enforcement salience. Columns 1–2: H is the (standardized) municipal drop in deforestation between 2004 and 2006 from PRODES. Columns 3–4: H is (standardized) cloud cover that hampers satellite monitoring from Assunção et al. (2023). Fixed effects vary by column for municipality-year my and match pairs. Standard errors are clustered at the individual level. Monetary values are in R$1000.

To test whether increased donations reflect resistance by the agricultural sector to the government's environmental policy, we examine heterogeneous effects using two measures of enforcement intensity at the municipal level. Table 3 reports estimates from the triple-difference specification in equation (3). In columns 1 and 2, we use the municipal reduction in annual deforestation from 2004 to 2006, measured using PRODES. The triple-interaction coefficients are positive and statistically significant at the 1% level, indicating larger post-2006 increases in donations in municipalities that experienced bigger declines in deforestation over this period (i.e., deforestation in 2004 minus deforestation in 2006).

In Table 3, columns 3 and 4 exploit enforcement variation arising from differential cloud cover. Assunção et al. (2023) shows that cloud cover reduces the DETER system's monitoring capacity, which identifies deforestation hotspots and guides law enforcement. The triple-interaction coefficients are negative and statistically significant at the 1% level, consistent with smaller donation increases in municipalities with higher average cloud cover (lower enforcement) relative to municipalities with less cloud cover (higher enforcement). Taken together, these exercises suggest that agricultural donors increased donations more where PPCDAm was more binding.

Finally, we estimate a triple-difference specification that compares the relative

Table 4: Donor donations in the Legal Amazon

	Sample outside Amazon			
	Rest of Brazil		Neighboring states	
	(1)	(2)	(3)	(4)
$Ag \times Post$ 2006	0.65	0.53	-0.08	-0.33
	(0.55)	(0.75)	(1.43)	(1.16)
$Ag \times Post$ 2006 $\times$ Amazon	2.70***	2.31*	3.77**	3.53**
	(0.97)	(1.24)	(1.66)	(1.58)
Observations	367,295	13,045	147,905	9,080
Fixed effects	my	match + my	my	match + my

Triple-difference estimates of equation (3). Outcomes are donations made by donors. *Amazon* indicates municipalities in the Legal Amazon. Coefficients on the triple interaction compare post-2006 changes for agricultural donors inside vs. outside the Amazon. Columns 1–2 include donors from the whole country; columns 3–4 include donors in the Legal Amazon and in neighboring regions – MA, PI, CE, BA, MG, MS, and GO. Fixed effects vary by column for municipality-year my and match pairs. Standard errors are clustered at the individual level. Monetary values are in R$1000.

increase in agricultural-sector campaign donations inside the Legal Amazon with donations from the rest of Brazil. In Table 4, columns 1 and 2 show that while donations from the agricultural sector seem to rise nationally, the point estimates outside the Amazon are small and not statistically different from zero. By contrast, the triple-interaction coefficient is about four times larger and statistically significant, implying that the increase in agricultural donations was concentrated in the Amazon. Columns 3 and 4 show similar results when we restrict the outside-Amazon comparison group to donors in neighboring regions. Altogether, the evidence indicates a post-2006 shift in campaign contributions by agricultural donors concentrated in the Amazon. Appendix Table A2 shows that these results are robust to excluding the Cerrado biome and to restricting the sample to the Cerrado only.

3.4 Candidates

Switching the unit of analysis from donors to candidates, we assess the support received over time by candidates who have ever received donations from the agricultural sector. We estimate a difference-in-differences regression analogous to equation (2), with i indexing candidates. Table 5 reports the estimates. In columns 1 and 2, the

Table 5: Candidate donations

	Total donations		Donations from: Ag. firm owners		Donations from: Infractors	
	(1)	(2)	(3)	(4)	(5)	(6)
$Ag \times Post$ 2006	32.85***	31.24***	11.08***	11.70***	5.08***	5.75***
	(2.79)	(2.53)	(1.17)	(0.94)	(0.87)	(0.88)
Effect as % of mean	176.6	168	257.8	272.2	141	159.6
Observations	15,660	6,290	15,660	6,290	15,660	6,290
Fixed effects	my	match + my	my	match + my	my	match + my

Difference-in-differences estimates of equation (2). Outcomes are donations received by candidates: total (columns 1–2), from agricultural donors (columns 3–4), and from donors with environmental infractions (columns 5–6). Fixed effects vary by column for municipality-year my and match pairs. Standard errors are clustered at the candidate level. Monetary values are in R$1000.

total campaign funds of candidates who received donations from agricultural donors increased by about 170 percent after 2006, relative to candidates financed by owners in other sectors.

Table 5 also shows changes in the composition of candidate donations, further supporting the growing political influence of the agricultural sector. Not only did the total funds received by agriculture-linked candidates increase, but donor composition shifted as well. Columns 3 and 4 show a roughly 250 percent increase in donations received by these candidates from agricultural donors. Columns 5 and 6 suggest over a 140 percent increase in donations received from donors with a history of environmental infractions. These findings indicate a transformation in both the volume and the sources of political financing.

Another way to measure political support for candidates linked to the agricultural sector is to examine electoral outcomes. Table 6 presents the effects of PPCDAm on these candidates' electoral performance, distinguishing between federal (Panel A) and state (Panel B) congressional elections.[3] In columns 1 and 2, agriculture-linked candidates receive about 70 percent more votes than candidates financed by owners in other sectors. In columns 3 and 4, the probability that agriculture-linked candidates for the federal congress are elected rises by more than 80 percent after 2006. For state

[3] Given differences in vote totals and electoral competition across state and federal races, we analyze them separately.

Table 6: Candidate votes

	Votes		Elected	
	(1)	(2)	(3)	(4)
Panel A: Federal Congress				
Ag × *Post* 2006	32.61***	38.75***	0.33***	0.41***
	(4.35)	(13.68)	(0.05)	(0.14)
Effect as % of mean	71.2	84.6	81.6	102.6
Observations	1,378	108	1,378	108
Fixed effects	*my*	match + *my*	*my*	match + *my*
Panel B: State Congress				
Ag × *Post* 2006	10.37***	5.68**	0.30***	-0.02
	(0.92)	(2.74)	(0.03)	(0.14)
Effect as % of mean	87.6	48.0	75.5	-5.5
Observations	4,995	247	4,995	247
Fixed effects	*my*	match + *my*	*my*	match + *my*

Difference-in-differences estimates of equation (2). Outcomes are thousands of votes received by candidates (columns 1–2) and an indicator for being elected (columns 3–4). Panel A considers candidates running for the Federal Congress (*Deputados Federais*) and Panel B those running for State Congress (*Deputados Estaduais*). Fixed effects vary by column for municipality-year *my* and match pairs. Standard errors are clustered at the candidate level.

contests, point estimates for election probabilities are similar in sign but become small and statistically indistinguishable from zero in the specification with matched-pair fixed effects. Taken together, these results suggest that PPCDAm increased political support for agriculture-linked candidates: they raised more campaign funds, garnered more votes, and, in federal races, were more likely to be elected.

4 Policy Trade-offs

We study a simple model of deforestation under regulation. We introduce political resistance, which can undercut regulation by leading to repeal. We discuss implications for optimal policy.

4.1 Regulation

Consider plots $i \in \mathcal{I}$. Plots contain forests, and they face a choice $x_i \in \{0, 1\}$ to develop these forests into plantations. Plantations generate private agricultural value $v_i > 0$, which includes agricultural revenue net of development and production costs. Plantation development also involves deforestation, which releases carbon emissions $e_i > 0$ at social cost $\kappa > 0$ per unit. Development is subject to regulation $\tau = \{\tau_i\}_i$, where $\tau_i \geq 0$. Private profits are private values net of regulation.

$$\pi_i(\tau_i) = (v_i - \tau_i)x_i(\tau_i) \tag{4}$$

Profit maximization yields a cutoff rule for development.

$$x_i(\tau_i) = \mathbb{1}[v_i \geq \tau_i] \tag{5}$$

All plots develop in the absence of regulation, as $v_i > 0$ implies $x_i(0) = 1$. Non-development generates zero profits.

The regulator chooses regulation τ to maximize welfare. For a given plot, welfare is agricultural profits and government revenue net of carbon emissions, which simplifies to agricultural value net of carbon emissions by equation (4). Regulation is a transfer from producers to the government.

$$\begin{aligned} w_i(\tau_i) &= \pi_i(\tau_i) + (\tau_i - \kappa e_i)x_i(\tau_i) \\ &= (v_i - \kappa e_i)x_i(\tau_i) \end{aligned} \tag{6}$$

Total welfare sums across plots.

$$W(\tau) = \sum_i w_i(\tau_i)$$

As is typical, regulation aims to force producers to internalize their social costs of deforestation.

Result 1. Pigouvian tax $\tau_i^P = \kappa e_i$ maximizes welfare $W(\tau)$ and $w_i(\tau_i)$.

Proof. By equation (6), we have $w_i(\tau_i^P) = \pi_i(\tau_i^P)$ for $\tau_i^P = \kappa e_i$. Profit maximization leads to welfare maximization. □

4.2 Resistance

We introduce political resistance, such that regulation is subject to potential repeal. We consider a simple sequential timeline. First, the regulator announces regulation. Second, agricultural producers lobby against the regulation in proportion to their regulatory burden. Stronger regulation has a higher risk of repeal. Third, the regulation is repealed with some probability. Repeal removes regulation for all producers, whether they lobby or not. Fourth, if repeal is successful, then production proceeds under zero regulation. If repeal is unsuccessful, then production proceeds under the proposed regulation.

More formally, regulation τ is repealed with probability R, which is strictly increasing in the total burden B of regulation across plots.

$$R'(B) > 0, \quad B(\tau) = \sum_i b_i(\tau_i) \tag{7}$$

For a given plot, this burden is the private profits lost under regulation.

$$\begin{aligned} b_i(\tau_i) &= \pi_i(0) - \pi_i(\tau_i) \\ &= \min\{v_i, \tau_i\} \end{aligned} \tag{8}$$

The second line follows from equations (4) and (5), by which $\pi_i(0) = v_i$ and $\pi_i(\tau_i) = \max\{v_i - \tau_i, 0\}$. If regulation halts production, then the plot loses v_i. If regulation allows but taxes production, then the plot loses τ_i.

Result 2. Repeal risk $R(B(\tau))$ is non-decreasing in τ_i.

Proof. By equation (8), $b_i(\tau_i)$ is non-decreasing in τ_i. By equations (7), it follows that R and B are non-decreasing in τ_i. □

4.3 Trade-off

The regulator chooses regulation τ to maximize feasible welfare $F(\tau)$, which accounts for repeal risk $R(\tau) = R(B(\tau))$.

$$\begin{aligned} F(\tau) &= W(\tau)[1 - R(\tau)] + W(0)R(\tau) \\ &= W(\tau) - [W(\tau) - W(0)]R(\tau) \end{aligned} \tag{9}$$

The first line expresses welfare in expectation of repeal. The second line highlights the tension between maximizing welfare $W(\tau)$ and minimizing repeal risk $R(\tau)$. Repeal leads to non-regulation and social losses, as $W(\tau) \geq W(0)$ unless $\kappa e_i < v_i$ for all plots i. The regulator must balance regulation against resistance.

Result 3. Optimal regulation $\tau_i^* \leq \tau_i^P$ maximizes feasible welfare $F(\tau)$, trading off improvements in $W(\tau)$ and $R(\tau)$.

Proof. First, $\tau_i > \tau_i^P$ does not maximize $F(\tau)$. We have $W(\tau_i) \leq W(\tau_i^P)$ by result 1 and $R(\tau_i) \geq R(\tau_i^P)$ by result 2. Thus, $F(\tau_i) \leq F(\tau_i^P)$ by equation (9). Second, $\tau_i^* \leq \tau_i^P$ faces a trade-off. On one hand, we have $W(\tau_i^*) \leq W(\tau_i^P)$ by result 1. Relative to τ_i^P, we lose in terms of $W(\tau)$. On the other hand, we have $R(\tau_i^*) \leq R(\tau_i^P)$ by result 2. We gain in terms of $R(\tau)$. □

4.4 Targeting

Without political resistance, the regulator maximizes welfare by targeting emissions with a Pigouvian tax. The standard targeting principle applies. With political resistance, the regulator maximizes *feasible* welfare with a trade-off: the Pigouvian tax maximizes welfare, but it also risks repeal. Precise targeting lessens the trade-off.

That is, the targeting principle applies along new margins. First, result 3 highlights an aggregate tension between welfare and repeal, but this tension varies across plots. Optimal policy can target plots with high emissions and low agricultural values, thereby achieving large welfare gains with limited regulatory burden. Second, result 4 shows that non-regulation is unambiguously optimal for plots with $v_i \geq \kappa e_i$, such that agricultural values exceed emission damages. The reason is that production is efficient on these plots. In allowing production to proceed without regulatory burden, we minimize the risk of global repeal, which would lead to welfare losses from inefficient plots. At the same time, precise targeting has substantial informational requirements that we will assess in our quantification exercise.

Result 4. Targeting $T(\tau_i) = \tau_i \cdot \mathbb{1}[v_i < \kappa e_i]$ improves feasible welfare $F(\tau)$.

Proof. Let $\tau_i^T = T(\tau_i)$. For plots with $v_i < \kappa e_i$, we have $\tau_i^T = \tau_i$. Targeting changes neither welfare $w_i(\tau_i)$ nor burden $b_i(\tau_i)$. For plots with $v_i \geq \kappa e_i$, we have $\tau_i^T = 0 \leq \tau_i$. It follows that $x_i(\tau_i^T) \geq x_i(\tau_i)$ by equation (5), $w_i(\tau_i^T) \geq w_i(\tau_i)$ by equation (6), and

$b_i(\tau_i^T) \leq b_i(\tau_i)$ by equation (8). Targeting raises welfare $w_i(\tau_i)$ and lessens burden $b_i(\tau_i)$. Aggregating plots, $W(\tau_i^T) \geq W(\tau_i)$, $B(\tau_i^T) \leq B(\tau_i)$, and $R(\tau_i^T) \leq R(\tau_i)$. By equation (9), $F(\tau_i^T) \geq F(\tau_i)$. □

4.5 Transfers

Our framework takes the distributional consequences of regulation seriously. The standard approach is to maximize aggregate welfare $W(\tau)$ and rely on lump-sum transfers to address distributional concerns. Under this approach, political resistance is irrelevant: losers from regulation are compensated, and the Pigouvian tax τ^P is optimal without modification. We argue that transfers encounter major difficulties in practice.

The first challenge is informational. Compensating producers requires knowing their private values v_i, which are unobserved. The regulator cannot simultaneously elicit truthful information about costs and implement efficient transfers, as producers have incentives to overstate their losses. Optimal compensation thus faces a fundamental mechanism design problem. Even well-intentioned transfer schemes will systematically over- or under-compensate individual producers.

The second challenge is fiscal. Full compensation requires transfers on the order of total forgone agricultural profits $\sum_i b_i(\tau_i)$. In the Brazilian Amazon, our empirical analysis will suggest that sum is on the order of hundreds of billions of dollars. Agricultural returns in this region reflect some of the most productive cropland and pasture in the world, and the number of affected producers spans thousands of municipalities. The scope and scale of fiscal transfers would dwarf any historical environmental budget in Brazil or elsewhere.

The third challenge is institutional. Site-specific transfers would at least partially target smallholders across remote frontier municipalities, relying on administrative capacity that is largely absent at present. Brazil has experimented with payments for ecosystem services through programs such as *Bolsa Verde* and *Bolsa Floresta*, but their coverage and payment levels are small relative to the compensation needed for the political feasibility of full Pigouvian taxation. The limited reach of these programs is itself evidence that transfers are difficult to implement at the relevant scale.

Beyond these implementation challenges, transfers face credibility and coordination problems. Transfer commitments are subject to time inconsistency: future

administrations may renege, and indeed Brazilian environmental policy has shifted dramatically across political cycles. Producers who rationally discount such promises will resist regulation despite the promise of future payments. Furthermore, most benefits of forest protection accrue globally through carbon abatement and biodiversity conservation, but there is currently no effective mechanism for large international-to-local transfers. The experience of REDD+ payments illustrates this difficulty, with modest actual payments over the course of decades.

Finally, and most directly, we observe political resistance to environmental regulation in the historical data. The effects we document in Section 3 are direct evidence that compensating transfers are not taking place in practice. Indeed, the salience of regulation in Brazilian political discourse – from debates over the Forest Code to resistance against protected area enforcement – centers on who bears the costs, and not on the hypothetical transfers that might accompany them. We therefore model the direct incidence of regulation on agricultural producers, rather than assuming that transfers can navigate the distributional consequences. The burden $b_i(\tau_i)$ in equation (8) drives political resistance, and the regulator must account for this resistance when designing policy. Our quantification takes this constraint at face value.

5 Politically Optimal Policy

We bring the theory to data with a simple empirical model, which we use to quantify the welfare and emissions impacts of counterfactual regulatory policy.

5.1 Model

Each site i has a fixed endowment of n_i hectares of land. At time t, the site allocates a_{it} hectares to agricultural use and leaves $f_{it} = n_i - a_{it}$ as forest. Development implies deforestation, as the expansion of agricultural land comes at the expense of forest. For each hectare j, the site chooses between agricultural use and forest, maximizing profits subject to idiosyncratic shocks.

$$\max \left\{ \beta v_{it} + \epsilon^a_{ijt},\, \epsilon^f_{ijt} \right\}$$

Agricultural use has value v_{it}, while forests have private value that we normalize to zero. Shocks $(\epsilon^a_{ijt}, \epsilon^f_{ijt})$ are logit-distributed and independent. Parameter $\beta > 0$ governs the responsiveness of land allocations to agricultural value. Logit shocks yield simple expressions for agricultural land and forest.

$$\frac{a_{it}}{n_i} = \frac{\exp(\beta v_{it})}{1 + \exp(\beta v_{it})}, \quad \frac{f_{it}}{n_i} = \frac{1}{1 + \exp(\beta v_{it})}, \tag{10}$$

where agricultural value v_{it} is given by log returns r_{it} net of costs c_{it}.

$$v_{it} = \log\left(\frac{r_{it}}{c_{it}}\right) \tag{11}$$

Costs c_{it} capture barriers to production, including distance to markets and the private costs of development.

5.2 Estimation

We estimate the model with data on agricultural land use and returns across the Brazilian Amazon from 2001 to 2019.

Specification

Applying the logit inversion to equation (10) and substituting equation (11), we obtain

$$\log\left(\frac{a_{it}}{f_{it}}\right) = \beta \log r_{it} - \beta \log c_{it}.$$

We observe agricultural and forested lands (a_{it}, f_{it}), and we will construct agricultural returns r_{it}. Costs c_{it} are unobserved. The specification that we take to the data is

$$\log\left(\frac{a_{it}}{f_{it}}\right) = \beta \log r_{it} + \gamma_i + \delta_t + \varepsilon_{it}, \tag{12}$$

Site and time fixed effects (γ_i, δ_t) absorb common components of unobserved costs, while error term ε_{it} captures the remaining variation. We estimate equation (12) by Poisson pseudo-maximum likelihood (PPML). In exponential form, the specification becomes $a_{it} = f_{it} \cdot \exp(\beta \log r_{it} + \gamma_i + \delta_t) \cdot \exp(\varepsilon_{it})$ with f_{it} entering as an offset. The unit of observation is a site i in a year t, and we weight sites by land area.

Our main specification estimates equation (12) in pooled long differences over a 10-year horizon. For each start year from 2001 to 2009, we compute the long difference to the corresponding endpoint 10 years later, and we stack these windows into a single dataset. We estimate a common return coefficient across windows, interacting fixed effects with window indicators to allow for window-specific intercepts. Identification thus comes from within-window, cross-site variation in baseline returns. We cluster standard errors at the level of Brazilian microregions to account for autocorrelation in time and space: sites appear in multiple windows, and nearby sites face correlated agricultural conditions. The key parameter is β, which captures the elasticity of the agriculture-to-forest ratio with respect to agricultural returns over a decade-long horizon.[4] We contrast the resulting estimates to those from a panel specification, which draws on the annual data from 2001 to 2019. The panel estimates capture short-run responses from year to year, while the long-difference estimates capture cumulative responses over a 10-year horizon.

Agricultural returns

Agricultural returns combine spatial variation in productivity with time variation in global commodity prices. We construct returns as a weighted average across agricultural activities $k \in \{\text{soy, maize, pasture}\}$.

$$r_{it} = \sum_k w_i^k y_i^k p_t^k, \tag{13}$$

where yields y_i^k are attainable crop yields for soy and maize and agro-climatic suitability for pasture, prices p_t^k are global commodity prices, and weights w_i^k aggregate across activities.[5] Returns thus vary across sites through local productivity, as well

[4] We note that parameter β differs from the elasticity of deforestation with respect to agricultural prices. First, deforestation is a flow – the change in the agricultural land stock – which is typically a small fraction of the stock itself. A small proportional change in the stock translates to a large proportional change in the flow, and so a modest stock elasticity can imply a much larger flow elasticity. Second, β understates the elasticity of agricultural lands with respect to individual crop prices. Our composite returns r_{it} aggregate across crops, and so a change in a single crop price translates less than one-for-one into composite returns. Both forces potentially compress β relative to the deforestation-price elasticity.

[5] For pasture, we convert agro-climatic suitability into a cardinal measure of productivity. Following Domínguez-Iino (2026), we use the estimated relationship

$$\text{cattle per hectare} = 0.2 + 1.2 \cdot \text{suit_pasture}.$$

as over time through global prices. We capitalize annual returns to a present value by dividing by $(1-\delta)$ for discount factor $\delta = 0.9$.

Weights w_i^k determine how local agricultural composition translates global price variation into local profitability shocks. We consider three approaches to weighting. First, we use baseline land-use shares for each site as observed in 2008. This measure directly captures local agricultural composition, but it may also partially reflect endogenous land-use choices. In particular, sites' own past land-use decisions may be correlated with unobserved determinants of future agricultural expansion. Second, we use leave-one-out shares based on aggregate shares in 2008 at the level of Brazilian microregions. We exclude sites' own shares. Third, following Fiszbein (2022), we use fitted shares from a Poisson land allocation model that predicts land-use shares from attainable yields and suitability. These fitted shares reflect expected land allocations given agro-climatic conditions.

For long differences, we re-center baseline shares to the beginning of each 10-year window, such that local agricultural composition is measured prior to the outcome change in that window. We calculate leave-one-out shares by aggregating at the microregion level. We compute fitted shares by estimating the Poisson land allocation model separately at the starting year of each 10-year window. For panel estimation, we construct all weights using baseline land allocations in 2000.

Estimates

In Table 7, Panel A presents our long-difference estimates. Higher profitability leads to expansion of agricultural land at the expense of forest, with consistent estimates across all three approaches to constructing agricultural returns. The estimated β coefficients range from 0.84 to 1.20 and are statistically significant. Column 1 uses observed baseline shares within each site, column 2 uses leave-one-out shares from the microregion, and column 3 uses fitted shares from the Poisson land allocation model. The estimates are broadly stable across specifications, with the fitted-share approach yielding the largest point estimate.

Panel B presents our panel estimates, which rely on the annual variation from 2001

We then convert stocking rates into annual output per hectare, multiplying by 0.5 tons (the average weight of a cattle head) and dividing by 3 (reflecting an average slaughter age of three years). We obtain yields live cattle weight for slaughter per hectare, which we combine with the price of live cattle to obtain pasture returns. Agricultural prices are from IPEA's farmgate price series for soy, maize, and live cattle (IPEA 2024).

Table 7: Agricultural returns and development

	Crop weights		
	Baseline shares	Leave-one-out shares	Fitted shares
	(1)	(2)	(3)
Panel A: Long differences			
Log agricultural returns	0.84***	1.02***	1.20***
	(0.18)	(0.34)	(0.28)
Observations	37,274	37,274	37,274
Fixed effects	$iw + tw$	$iw + tw$	$iw + tw$
Panel B: Panel estimates			
Log agricultural returns	0.46***	0.60***	0.67***
	(0.10)	(0.20)	(0.16)
Observations	42,620	42,620	42,620
Fixed effects	$i + t$	$i + t$	$i + t$

Panel A presents PPML estimates of equation (12) in pooled 10-year long differences with data from 2001 to 2019, while Panel B uses the full annual panel variation with two-way fixed effects. The dependent variable is agricultural land area with log forest area as an offset. Each column uses a different approach to constructing the weights in equation (13): observed baseline shares, leave-one-out shares from the microregion, and fitted shares from a Poisson land allocation model. Panel A includes site-by-window and year-by-window fixed effects (iw and tw), while Panel B includes site and year fixed effects (i and t). Standard errors are clustered at the microregion level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

to 2019. The estimates range from 0.46 to 0.67 and remain similar across weights. Fitted shares again yield the largest estimates. Moreover, the panel estimates are substantially smaller than the long-difference estimates. This gap is consistent with gradual adjustment: year-to-year profitability changes shift the agricultural frontier slowly, and so short-run responses understate the long-run elasticities. The long-difference estimates, which capture cumulative responses over a decade, are thus the appropriate inputs for our counterfactual analysis.

We validate our main results in Appendix B. First, Table B1 estimates the model at the municipality level, rather than the 0.25-degree cells of our main specification. The long-difference estimates lie in a range from 1.15 to 1.73, which is larger than our main estimates, while the panel estimates range from 0.66 to 0.97. The pattern

of larger long-run than short-run responses holds at this spatial scale, and the estimates remain similar across weights. Second, Table B2 shows that our estimates are robust to constructing agricultural returns with three- and five-year moving average prices, rather than with annual prices. Third, Figure B1 finds that our long-difference estimates are increasing in the length of the horizon, as is consistent with gradual adjustment of the agricultural frontier.[6] Fourth, Figure B2 shows that our long-difference estimates are largely stable for start years ranging from 2001 to 2009. Our main specification pools across these start years. Fifth, Figure B3 assesses pre-trends. A potential concern is that sites with high agricultural returns may be on differential land-use trajectories. Testing for pre-trends in long differences is difficult because we are limited by our sample that begins in 2001. We instead present the event study for our panel specification, which indeed shows no strong pre-trends.[7]

5.3 Counterfactual policy

The estimated model allows us to evaluate counterfactual regulatory policy. We evaluate counterfactuals for the year 2019, which is the endpoint of our sample, and we suppress the time subscript in what follows.

Regulation

Regulation reduces agricultural value, discouraging agricultural development and deforestation. We specify regulation as an ad valorem tax $\tau_i \in [0, 1]$ on agricultural returns. By equation (11), agricultural value under tax τ_i becomes

$$v_i(\tau_i) = \log\left(\frac{(1-\tau_i)r_i}{c_i}\right).$$

[6] Shortening the long-difference horizon does not lead to convergence to the panel estimate. These specifications differ in their use of the data and their fixed effects. The long-difference specification pools over start years and includes site-by-window and year-by-window fixed effects, while the panel specification uses all annual observations with site and year fixed effects.

[7] We estimate coefficients that are small and not statistically significant at leads of two to four years. We obtain a positive, significant coefficient at a lead of one year, which is arguably consistent with the time required to convert cleared land into productive agriculture. That is, land deforested in one year is typically not harvested until the following growing season. Thus, by nature, agricultural expansion must be in anticipation of returns realized one year later.

We normalize baseline regulation to zero, evaluating taxation in excess of the status quo. Baseline regulation is subsumed in costs c_i, which we hold fixed in counterfactuals. Appendix B formalizes this normalization and shows that it interprets baseline regulation as non-revenue-generating. Indeed, Brazilian policy has centered on the 80 percent rule and protected areas, rather than revenue-generating Pigouvian taxes. By equation (10), we have

$$\frac{a_i(\tau_i)}{f_i(\tau_i)} = \exp\big(\beta v_i(\tau_i)\big). \tag{14}$$

We write $a_i = a_i(0)$ and $f_i = f_i(0)$ as shorthand for the baseline values observed in the data. It follows directly that

$$\log\left(\frac{a_i(\tau_i)}{f_i(\tau_i)}\right) - \log\left(\frac{a_i}{f_i}\right) = \beta \log(1-\tau_i). \tag{15}$$

Returns r_i and costs c_i cancel, and total land $n_i = a_i(\tau_i) + f_i(\tau_i)$ is fixed. Thus, given estimated parameter β, observed (a_i, f_i, n_i), and any counterfactual tax τ_i, we can compute counterfactual agricultural land.

$$a_i(\tau_i) = n_i \left(\frac{\frac{a_i}{f_i}(1-\tau_i)^\beta}{1 + \frac{a_i}{f_i}(1-\tau_i)^\beta}\right)$$

Welfare

Our welfare metrics are as in Section 4. Producer surplus is the logit inclusive value, which we monetize and aggregate across hectares.

$$\pi_i(\tau_i) = \frac{n_i r_i}{\beta} \log\left(1 + \frac{a_i}{f_i}(1-\tau_i)^\beta\right), \tag{16}$$

where we apply equations (14) and (15) to the log-sum formula. Government revenue derives from tax τ_i on agricultural lands, while emissions derive from carbon stocks e_i on the same land. Total producer surplus, government revenue, and emissions are

$$\Pi(\tau) = \sum_i \pi_i(\tau_i), \quad G(\tau) = \sum_i \tau_i r_i a_i(\tau_i), \quad E(\tau) = \sum_i e_i a_i(\tau_i).$$

We evaluate welfare in differences relative to the status quo, defining notation $\Delta X(\tau) = X(\tau) - X(0)$. By equation (8), the burden of regulation is

$$B(\tau) = \Pi(0) - \Pi(\tau) = -\Delta\Pi(\tau). \tag{17}$$

Total welfare is the sum of producer surplus and government revenue, net of emission damages at social cost κ per ton of CO_2. In differences,

$$\Delta W(\tau) = \Delta\Pi(\tau) + \Delta G(\tau) - \kappa\Delta E(\tau), \tag{18}$$

where $\Delta G(\tau) = G(\tau)$ because we normalize baseline regulation to zero. Consumer surplus is zero under our assumption of perfectly elastic demand for agricultural goods.

5.4 Optimal policy

We study optimal policy that maximizes feasible welfare gains. By equation (9), these gains are

$$\Delta F(\tau) = \Delta W(\tau)[1 - R(B(\tau))], \tag{19}$$

where $B(\tau)$ and $\Delta W(\tau)$ are given by equations (17) and (18). To compute welfare, we must take a stance on the social costs κ of emissions. To compute feasible welfare and optimal policy, we must take a further stance on function $R(B)$. This function governs how strongly the probability R of repeal responds to the burden B of regulation. We specify it as follows.

$$R(B(\tau); \lambda) = \min\left\{\frac{B(\tau)}{\lambda}, 1\right\} \tag{20}$$

Repeal becomes certain when burden B reaches level $\lambda > 0$, such that $R = 1$ if $B \geq \lambda$. The risk of repeal decreases linearly below this level, such that $R = \frac{B}{\lambda}$ if $B < \lambda$. We choose \$500 billion as a large, conservative baseline value of λ. This figure is roughly one third of average annual GDP in Brazil during the study period. Lower values of λ will increase the likelihood of repeal, and higher values will decrease the likelihood.

We consider regulation of several forms. It is convenient to express taxes in units of value v_i by defining transformed tax $\tilde{\tau}_i = -\log(1-\tau_i)$, which enters the value function linearly. Linearity yields clean expressions for optimal policy, but the transformation also calls for careful welfare accounting that we detail in Appendix B. As a benchmark,

we quantify the impacts of Pigouvian tax $\tilde{\tau}^P$, which maximizes welfare gains ΔW.

$$\tilde{\tau}_i^P(\kappa) = \frac{\kappa e_i}{r_i} \tag{21}$$

for social costs $\kappa > 0$. High values of κ correspond to large emissions damages and strong taxation. The Pigouvian rate equals emission damages per dollar of agricultural returns, and so it varies across sites through carbon stocks e_i and returns r_i. By result 3, political resistance gives rise to trade-offs for feasible welfare, with optimal regulation generally falling short of full Pigouvian taxation. We illustrate with scaled Pigouvian tax $\tilde{\tau}^S$, which is softer than full Pigouvian taxation for scale $\sigma \in (0, 1)$.

$$\tilde{\tau}_i^S(\kappa, \sigma) = \sigma \tilde{\tau}_i^P(\kappa)$$

By result 4, we can achieve gains by further targeting on agricultural profits π_i, which we compute using the estimated model. We define efficiency index η_i as the ratio of a site's Pigouvian tax to its burden $b_i = \pi_i(0) - \pi_i(\tilde{\tau}_i^P)$ of regulation.

$$\eta_i = \frac{\tilde{\tau}_i^P}{b_i(\tilde{\tau}_i^P)}$$

Sites with high η_i are efficient to tax: they internalize large emissions externalities and contribute large tax revenues, each relative to their surplus losses and the resulting risk of repeal. Let $\eta_{(1)} \leq \eta_{(2)} \leq \cdots \leq \eta_{(N)}$ denote the order statistics. Optimal taxation $\tilde{\tau}^*$ excludes the least efficient sites with a cutoff ρ.

$$\tilde{\tau}_i^*(\kappa, \sigma, \rho) = \tilde{\tau}_i^S(\kappa, \sigma) \cdot \mathbb{1}\{\eta_i > \eta_{(\rho)}\}$$

Cutoff ρ excludes the ρ sites with the lowest efficiency, and then scale σ sets the tax on the remaining sites.

In simulating these policies, we choose parameters (σ, ρ) to maximize feasible welfare gains ΔF. For scaled Pigouvian tax $\tilde{\tau}^S(\kappa, \sigma)$, we have

$$\hat{\tau}_i^S(\kappa, \lambda) = \tilde{\tau}_i^S\big(\kappa, \hat{\sigma}(\kappa, \lambda)\big), \quad \hat{\sigma} = \arg\max_{\sigma} \big\{\Delta F\big(\tilde{\tau}^S(\kappa, \sigma); \kappa, \lambda\big)\big\}. \tag{22}$$

For optimal tax $\tilde{\tau}^*(\kappa, \sigma, \rho)$, we have

$$\hat{\tau}_i^*(\kappa, \lambda) = \tilde{\tau}_i^*\big(\kappa, \hat{\sigma}(\kappa, \lambda), \hat{\rho}(\kappa, \lambda)\big), \quad \hat{\sigma}, \hat{\rho} = \arg\max_{\sigma, \rho} \big\{\Delta F\big(\tilde{\tau}^*(\kappa, \sigma, \rho); \kappa, \lambda\big)\big\}. \tag{23}$$

We take as given the social costs κ of emissions and the level λ of repeal risk. We present counterfactual results across a range of values for (κ, λ).

5.5 Results

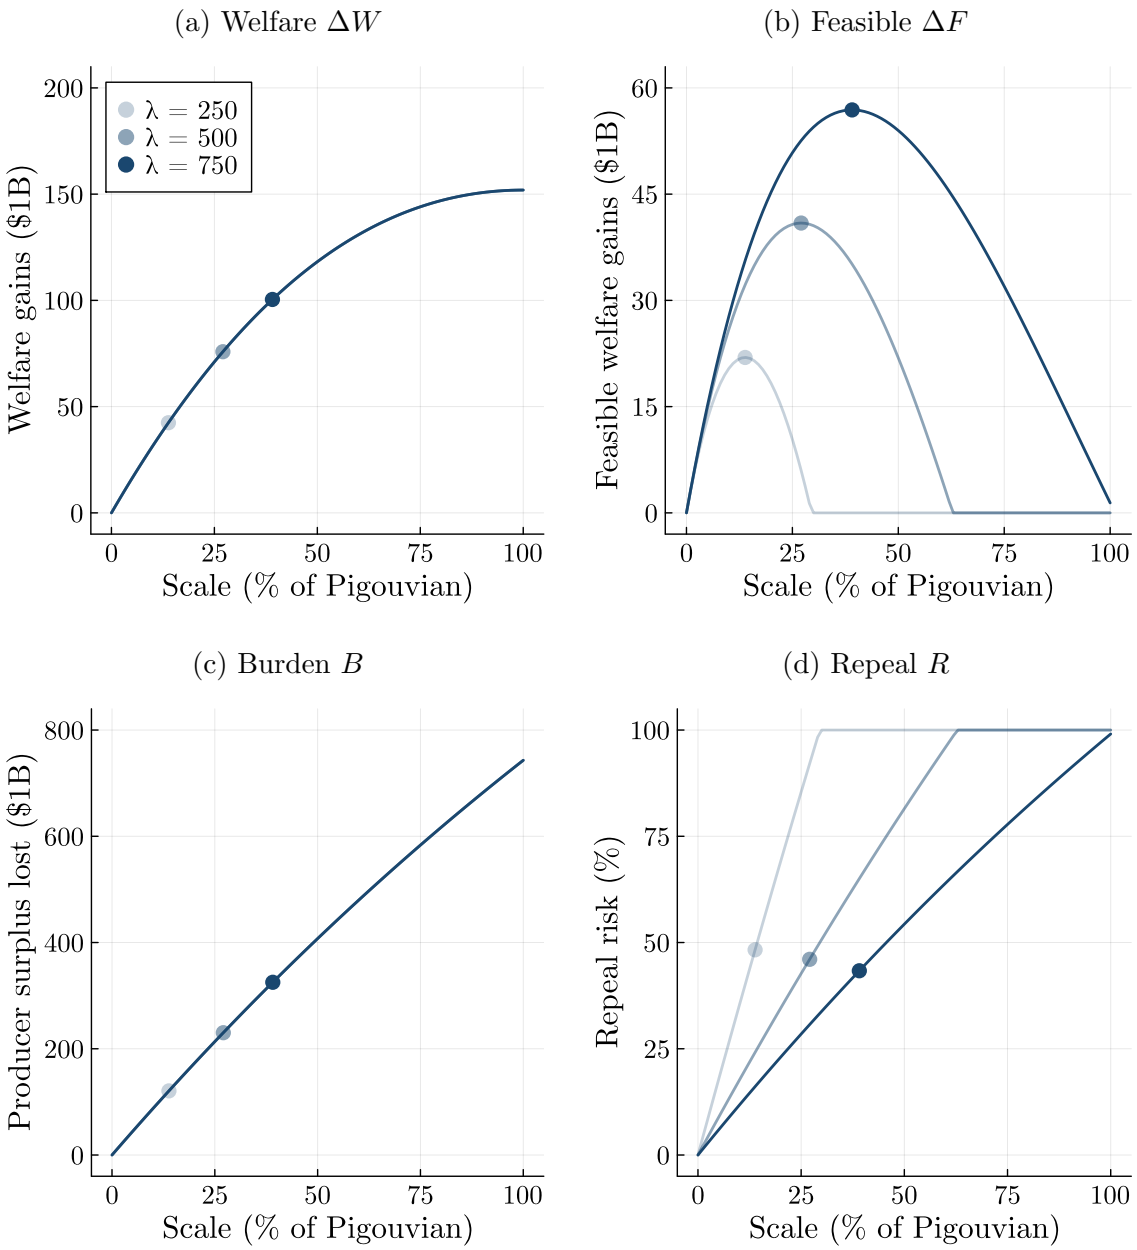
Figure 5 begins by illustrating the trade-off between welfare gains and the risk of repeal. We vary the scale σ of Pigouvian taxation, applying regulation to all sites by fixing $\rho = 0$. We set social costs at $\kappa = \$50$ per ton of CO_2, and we plot outcomes for repeal parameter $\lambda = \$250$, $\$500$, and $\$750$ billion. The y-axes show welfare gains, feasible welfare gains, the burden of regulation, and the probability of repeal. Dots mark the optimal scale parameter $\hat{\sigma}$, which varies with λ.

Focusing on the dots, Figures 5a and 5b show that both welfare and feasible welfare gains are increasing in λ. Higher λ implies a lower risk of repeal, which allows for stronger regulation and larger welfare gains. Indeed, Figure 5d shows that the risk of repeal is declining in λ. This decline is modest because lower repeal risk also encourages stronger regulation, which partially offsets the initial decrease in repeal risk. Figure 5c shows the consequence for producers: stronger regulation increases the burden of regulation, which we compute as producer surplus losses.

Focusing on the curves, we choose scale σ optimally to maximize feasible welfare gains. We see this maximization in Figure 5b, where the dots mark the peak of each curve. Consistent with result 3, the optimal scale falls short of full Pigouvian taxation at $\sigma = 100$ percent. Welfare gains are increasing in σ in Figure 5a, as implied by result 1. The curves overlap because welfare and burden depend only on the tax, and not on λ. But the burden of regulation is still increasing in σ, as we see in Figure 5c. The same holds for the risk of repeal, as we see in Figure 5d. The optimal scale $\hat{\sigma}$ balances these forces, trading off larger welfare gains against elevated repeal risk when considering strong regulation. Appendix Figure B4 illustrates the same trade-off by instead varying cutoff ρ, applying full Pigouvian taxation by fixing $\sigma = 100$ percent.

Figure 6 compares optimal and Pigouvian policies. The solid navy lines capture the impacts of optimal policy, as described in equation (23), while the dotted navy

Figure 5: Policy trade-offs



Counterfactual impacts of scaled Pigouvian taxation τ^S as a function of scale σ. We set cutoff $\rho = 0$ and thus apply Pigouvian taxation across all sites. Dots mark the optimal scale $\hat{\sigma}$ for each value of λ, which governs the sensitivity of repeal risk to producer burden. Lines correspond to λ values of \$250, \$500, and \$750 billion. Social costs are set at $\kappa = \$50$ per ton CO_2. Appendix Figure B4 instead varies cutoff ρ for fixed $\sigma = 100\%$.

lines correspond to full Pigouvian taxation, as in equation (21). Pigouvian regulation targets emissions alone and with full force, while optimal policy further targets producer profits and is sensitive to the impacts of regulation on repeal. The x-axes capture the value of abatement, which is equivalent to the social cost κ of emissions. We plot results for κ values that range from \$0 to \$100 per ton of CO_2, taking \$500 billion as our baseline value of λ.

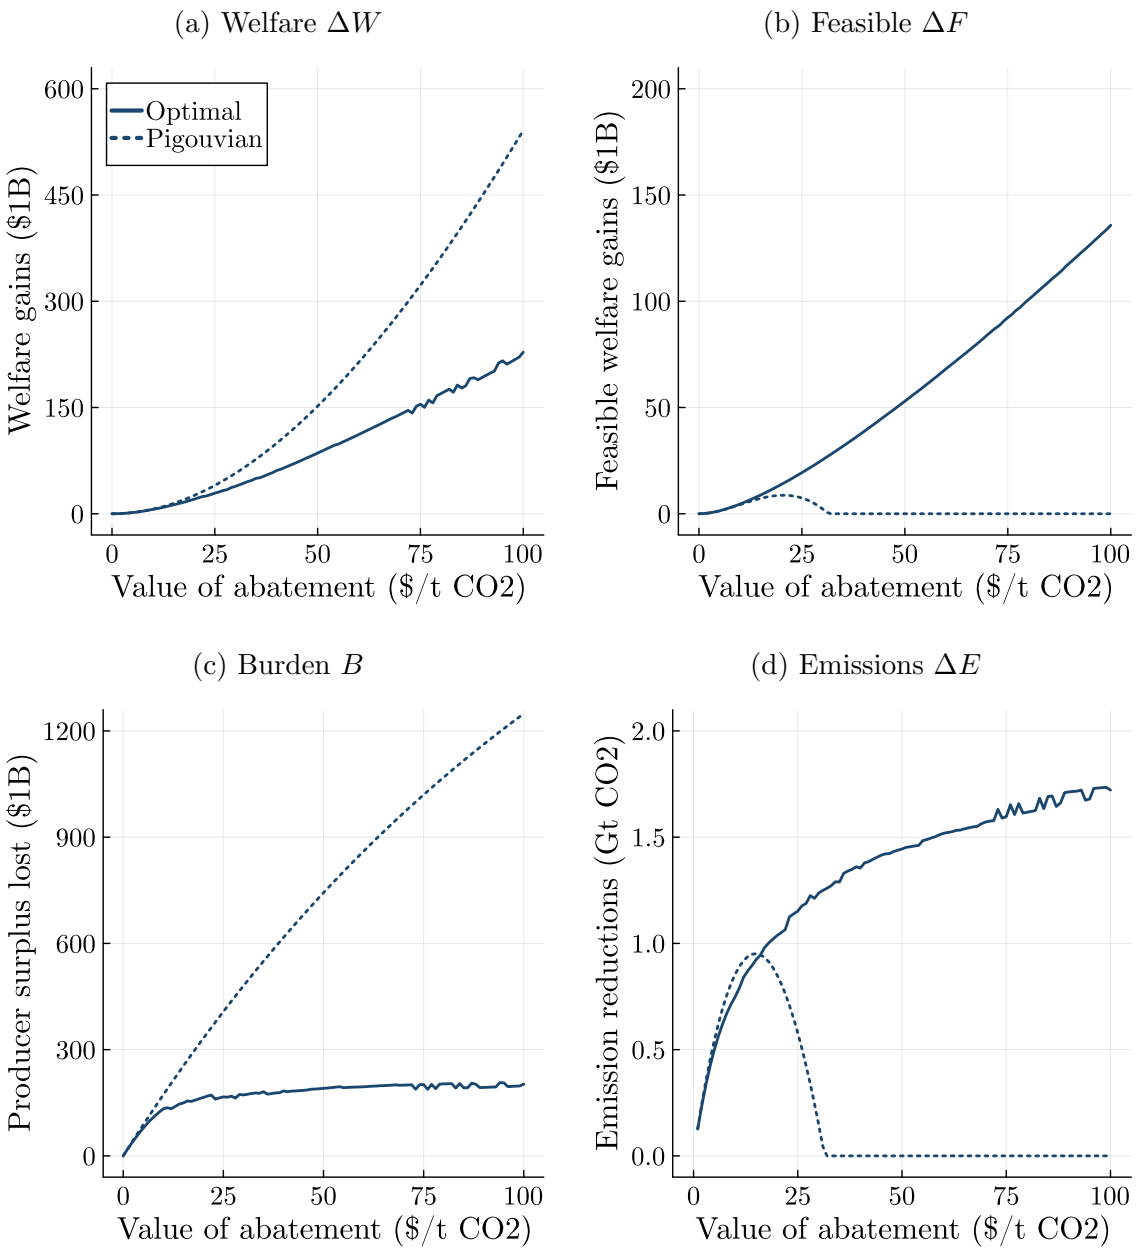
Large values of κ raise the potential welfare gains from regulation and thus encourage strong regulation. Consistent with the trade-off of result 3, strong regulation increases both welfare gains and the burden of regulation in Figures 6a and 6c. This pattern holds for both optimal and Pigouvian policies independently, as each curve is increasing from left to right.[8] The same trade-off applies to the difference between optimal and Pigouvian policies. That is, for all values of κ, Pigouvian policy achieves larger welfare gains than optimal policy, but Pigouvian policy also imposes a greater burden on producers.

Consistent with result 4, optimal policy is substantially better at navigating this trade-off. Figure 6b shows that optimal policy achieves higher feasible welfare gains than Pigouvian taxation for all κ, with the gap widening at higher social costs. At $\kappa = \$50$ per ton of CO_2, the optimal policy generates feasible welfare gains of over \$50 billion. By contrast, Pigouvian taxation imposes a large burden of \$750 billion on producers, resulting in certain repeal and zero feasible welfare gains. Moreover, the problem of repeal arises early on. We encounter certain repeal as low as $\kappa = \$30$ per ton of CO_2, which arguably represents a modest social cost of carbon and thus a modest Pigouvian tax. Similarly, Figure 6d shows that optimal policy leads to roughly equivalent or higher emissions reductions in expectation, as strong Pigouvian taxation fails to survive political resistance and the risk of repeal.

At the same time, the precise targeting of optimal policy relies on information about producer profits. Our estimated empirical model allows us to calculate these profits counterfactually, but these values are typically the private information of the firm. Regulators may instead only have access to measures of carbon stocks through

[8] The solid curves for optimal policy exhibit a small amount of non-monotonicity. This jaggedness arises from two sources. The first is the discrete nature of cutoff ρ, which involves the inclusion and exclusion of individual sites. The second is the discrete nature of ranking by efficiency index η_i, which depends on κ through Pigouvian tax $\tilde{\tau}_i^P(\kappa)$. Changes in κ lead to changes in η_i and thus changes in order statistics $\eta_{(\rho)}$. In both cases, the effects of discreteness are minimal because we have a large number of sites.

Figure 6: Optimal policy



Counterfactual impacts of optimal policy $\hat{\tau}^*$ (solid) and Pigouvian tax τ^P (dotted) as a function of social costs κ. We set $\lambda = \$500$ billion. Optimal policy jointly chooses scale σ and cutoff ρ to maximize feasible welfare gains ΔF, while Pigouvian taxation sets $\sigma = 100\%$ and $\rho = 0$.

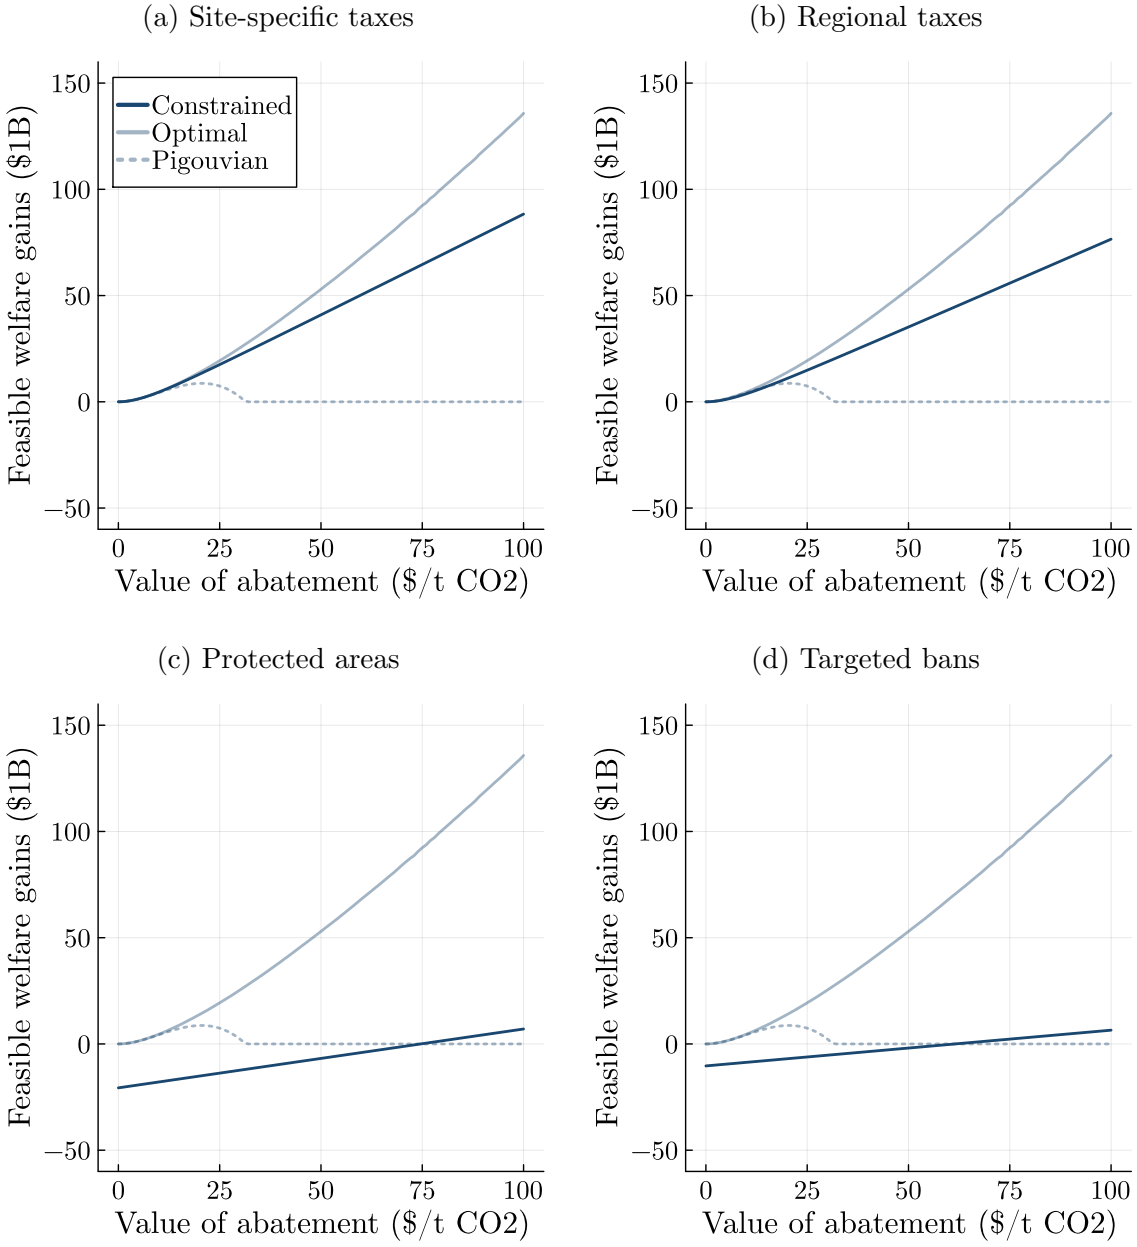
satellite accounting of aboveground biomass and forest density.

Figure 7 shows the impacts of policy that is constrained to targeting with this information on emissions alone. We plot feasible welfare impacts in navy blue for several candidate policies. As benchmarks, the light-blue lines present the impacts of the optimal and Pigouvian policies, as shown in Figure 6. In Figure 7a, we consider site-specific taxes that might be implemented through licensing fees paired with site inspections. This policy is equivalent to scaled Pigouvian tax $\tilde{\tau}^S$, as described in equation 22. It is more flexible than Pigouvian tax $\tilde{\tau}^P$ in that it optimizes over the scale σ of Pigouvian taxation in order to navigate the trade-off between welfare and repeal. At the same time, it is less flexible than optimal tax $\tilde{\tau}^*$ in that it does not optimize over the cutoff ρ, which would require information on producer profits. This constrained policy achieves smaller feasible welfare gains than optimal policy, but it also avoids the ineffectiveness of Pigouvian taxation at high levels.

We then turn to policies more similar to those used in practice. In Figure 7b, we again implement scaled Pigouvian tax $\tilde{\tau}^S$, but we tax based on average emissions at the mesoregion level rather than site-specific emissions. This regional tax is finer than full Pigouvian $\tilde{\tau}^P$ in that it optimizes over scale σ, but it is also coarser in that it does not target site-specific emissions. Full Pigouvian taxation dominates at low levels of taxation because of its site specificity, but the constrained tax dominates at higher levels because of its scale flexibility.

In Figures 7c and 7d, we assess bans on production. We simulate protected areas by banning development in the 5 percent of microregions with the highest average carbon density, and we simulate targeted bans by banning development in the five percent of sites with the highest site-specific carbon density. Both policies fall far short of optimal policy because they do not effectively navigate the risk of repeal. While bans achieve meaningful emissions reductions by targeting high-emission sites, they do so with an infinite tax that imposes maximal producer losses. Bans also fall short of Pigouvian policy when the value of abatement is low. At the extreme, when $\kappa = 0$, taxation is distortionary. The Pigouvian tax falls to zero, while bans continue to prohibit development on banned sites. At the same time, bans dominate Pigouvian policy when the value of abatement is high. In this case, bans achieve large welfare gains because the value of emissions reductions is large. They also avoid certain repeal because they affect no more than 5 percent of microregions or sites. Pigouvian policy instead imposes large taxes on all sites, to the point of certain repeal.

Figure 7: Constrained policy



Feasible welfare gains ΔF under constrained, emissions-only policies (navy) as a function of social costs κ. Light-blue lines show Pigouvian and optimal benchmarks from Figure 6. We set $\lambda = \$500$ billion. These policies optimize over scale σ but not cutoff ρ. Optimizing over σ requires only information on carbon stocks, while optimizing over ρ requires further information on producer profits. Site-specific taxes reflect carbon density at the site level. Regional taxes depend instead on the average carbon density across sites within a Brazilian mesoregion. Protected areas prohibit development of sites in the 5% of Brazilian microregions with the highest regional carbon density. Targeted bans prohibit development of the 5% of sites with the highest site-specific carbon density.

6 Conclusion

Environmental regulation can generate political resistance that undermines its own durability. We study this dynamic in the Brazilian Amazon, where conservation policy provoked a political backlash from agricultural producers. Using rich political data, we find that producers responded to environmental reforms by increasing support for pro-agricultural politicians who opposed the reforms. Motivated by this evidence, we develop a model of optimal regulation under the threat of political repeal. We estimate the model with data on land allocation and agricultural returns across the Brazilian Amazon from 2001 to 2019.

Our counterfactual analysis shows that accounting for the risk of repeal fundamentally changes optimal policy. Pigouvian taxation maximizes welfare in isolation but imposes large burdens on producers, triggering certain repeal and thus achieving little in expectation. Optimal policy softens taxation and further targets regulation toward sites where emissions reductions are large relative to producer losses. This targeted approach achieves substantial feasible welfare gains where Pigouvian taxation cannot. We emphasize that the design of conservation policy should account for these political constraints.

References

Abadie, Alberto, Susan Athey, Guido Imbens, and Jeffrey Wooldridge. When should you adjust standard errors for clustering? *Quarterly Journal of Economics*, 138 (1):1–35, 2023.

Araujo, Rafael, Francisco Costa, and Marcelo Sant'Anna. Efficient conservation of the Brazilian Amazon: Estimates from a dynamic model. *Review of Economic Studies*, 93(1):72–105, 2026.

Assunção, Juliano, Clarissa Gandour, and Rudi Rocha. Deforestation slowdown in the Legal Amazon: Prices or policies? *Environment and Development Economics*, 20(6):697–722, 2015.

Assunção, Juliano, Clarissa Gandour, Romero Rocha, and Rudi Rocha. The effect of rural credit on deforestation: Evidence from the Brazilian Amazon. *Economic Journal*, 130(626):290–330, 2019.

Assunção, Juliano, Robert McMillan, Joshua Murphy, and Eduardo Souza-Rodrigues. Optimal environmental targeting in the Amazon Rainforest. *Review of Economic Studies*, 90(4):1608–1641, 2022.

Assunção, Juliano, Clarissa Gandour, and Romero Rocha. DETER-ing deforestation in the Amazon: Environmental monitoring and law enforcement. *American Economic Journal: Applied Economics*, 15(2):125–56, 2023.

Assunção, Juliano, Lars Peter Hansen, Todd Munson, and José Scheinkman. Carbon prices and forest preservation over space and time in the Brazilian Amazon. 2024.

Balboni, Clare, Robin Burgess, Anton Heil, Jonathan Old, and Benjamin Olken. Cycles of fire? Politics and forest burning in Indonesia. *AEA Papers and Proceedings*, 111:415–419, 2021.

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

Balboni, Clare, Robin Burgess, and Benjamin Olken. The origins and control of forest fires in the tropics. *Review of Economic Studies*, 2025.

Berman, Aaron and Nathaniel Hickok. Environmental regulation with irreversible investments: Evidence from High Plains aquifer depletion. 2026.

Besley, Timothy and Torsten Persson. The political economics of green transitions. *Quarterly Journal of Economics*, 138(3):1863–1906, 2023.

Bombardini, Matilde, Frederico Finan, Nicolas Longuet-Marx, Suresh Naidu, and Francesco Trebbi. Climate politics in the United States. 2025.

Burgess, Robin, Matthew Hansen, Benjamin Olken, Peter Potapov, and Stefanie Sieber. The political economy of deforestation in the tropics. *Quarterly Journal of Economics*, 127(4):1707–1754, 2012.

Burgess, Robin, Francisco Costa, and Benjamin Olken. National borders and the conservation of nature. 2025.

Busch, Jonah, Kalifi Ferretti-Gallon, Jens Engelmann, Max Wright, Kemen Austin, Fred Stolle, Svetlana Turubanova, Peter Potapov, Belinda Margono, Matthew

Hansen, and Alessandro Baccini. Reductions in emissions from deforestation from Indonesia's moratorium on new oil palm, timber, and logging concessions. *Proceedings of the National Academy of Sciences*, 112(5):1328–1333, 2015.

CGU, Controladoria Geral da União. Relatório de Avaliação da Execução de Programas de Governo N. 69: Ações Relativas á Fiscalização Ambiental sob Responsabilidade do IBAMA. *Brasília: CGU*, 2016.

Costa, Francisco, Jason Garred, and João Paulo Pessoa. Winners and losers from a commodities-for-manufactures trade boom. *Journal of International Economics*, 102:50–69, 2016.

Costa, Francisco, Allan Hsiao, Heitor Pellegrina, and Eduardo Souza-Rodrigues. Deforestation. *VoxDevLit*, 2025a.

Costa, Francisco, Dimitri Szerman, and Juliano Assunção. The environmental costs of political interference: Evidence from power plants in the Amazon. *Journal of Public Economics*, 242:105314, 2025b.

Crouzeilles, Renato, Rafael Feltran-Barbieri, Mariana Ferreira, and Bernardo Strassburg. Hard times for the Brazilian environment. *Nature Ecology & Evolution*, 1(9): 1213, 2017.

Dahis, Ricardo and Erik Katovich. Seeds of influence: The political economy of Brazil's soy boom. 2025.

Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva. Fighting climate change: International attitudes toward climate policies. *American Economic Review*, 115(4):1258–1300, 2025.

Domínguez-Iino, Tomás. Efficiency and redistribution in environmental policy: An equilibrium analysis of agricultural supply chains. *Journal of Political Economy*, 134(4), 2026.

Fiszbein, Martin. Agricultural diversity, structural change, and long-run development: Evidence from the United States. *American Economic Journal: Macroeconomics*, 14(2):1–43, 2022.

Freitas, Flavio, Gerd Sparovek, Göran Berndes, U. Martin Persson, Oskar Englund, Alberto Barretto, and Ulla Mörtberg. Potential increase of legal deforestation in Brazilian Amazon after forest act revision. *Nature Sustainability*, 1(11):670, 2018.

Harding, Robin, Mounu Prem, Nelson Ruiz, and David Vargas. Buying a blind eye: Campaign donations, regulatory enforcement, and deforestation. *American Political Science Review*, 118(2):635–653, 2024.

Harding, Torfinn, Julika Herzberg, and Karlygash Kuralbayeva. Commodity prices and robust environmental regulation: Evidence from deforestation in Brazil. *Journal of Environmental Economics and Management*, 108:102452, 2021.

Hargrave, Jorge and Krisztina Kis-Katos. Economic causes of deforestation in the Brazilian Amazon: A panel data analysis for the 2000s. *Environmental and Resource Economics*, 54:471–494, 2013.

Hsiao, Allan. Sea level rise and urban adaptation in Jakarta. 2025.

Hsiao, Allan. Coordination and commitment in international climate action: Evidence from palm oil. *Econometrica*, 94(1):1–33, 2026.

Hsiao, Allan and Nicholas Kuipers. Climate crisis and policy inaction in Indonesia. *American Journal of Political Science*, 2025.

Hsiao, Allan, Jacob Moscona, Benjamin Olken, and Karthik Sastry. Climate change, deforestation, and the expansion of the global agricultural frontier. 2026.

Instituto de Pesquisa Econômica Aplicada. IPEA data: Agricultural farmgate prices, 2024. https://ipeadata.gov.br.

Kang, Karam. Policy influence and private returns from lobbying in the energy sector. *Review of Economic Studies*, 83(1):269–305, 2016.

Kashner, Zane. Building with externalities: Local governments and wind farms. 2026.

Leite-Mariante, Gabriel and Veronica Salazar Restrepo. Does conservation work in general equilibrium? 2024.

Longuet-Marx, Nicolas. Party lines or voter preferences? Explaining political realignment. 2025.

MMA, Ministério do Meio Ambiente. *Plano de Ação para Prevenção e Controle do Desmatamento na Amazônia Legal (PPCDAm) 3a Fase (2012-2015) pelo uso sustentável e conservação da Floresta*. Brasília, 2013.

Nepstad, Daniel, Britaldo Soares-Filho, Frank Merry, André Lima, Paulo Moutinho, John Carter, Maria Bowman, Andrea Cattaneo, Hermann Rodrigues, Stephan Schwartzman, David McGrath, Claudia Stickler, Ruben Lubowski, Pedro Piris-Cabezas, Sergio Rivero, Ane Alencar, Oriana Almeida, and Osvaldo Stella. The end of deforestation in the Brazilian Amazon. *Science*, 326(5958):1350–1351, 2009.

Oates, Wallace and Paul Portney. The political economy of environmental policy. In *Handbook of Environmental Economics*, volume 1, pages 325–354. Elsevier, 2003.

Rochedo, Pedro, Britaldo Soares-Filho, Roberto Schaeffer, Eduardo Viola, Alexandre Szklo, André Lucena, Alexandre Koberle, Juliana Leroy Davis, Raoni Rajão, and Regis Rathmann. The threat of political bargaining to climate mitigation in Brazil. *Nature Climate Change*, 8(8):695, 2018.

Ryan, Nicholas and Anant Sudarshan. Rationing the commons. *Journal of Political Economy*, 130(1):210–257, 2022.

Sanford, Luke. Democratization, elections, and public goods: The evidence from deforestation. *American Journal of Political Science*, 67(3):748–763, 2023.

Soares-Filho, Britaldo, Paulo Moutinho, Daniel Nepstad, Anthony Anderson, Hermann Rodrigues, Ricardo Garcia, Laura Dietzsch, Frank Merry, Maria Bowman, Letícia Hissa, Rafaella Silvestrini, and Cláudio Maretti. Role of Brazilian Amazon protected areas in climate change mitigation. *Proceedings of the National Academy of Sciences*, 107(24):10821–10826, 2010.

Soares-Filho, Britaldo, Raoni Rajão, Marcia Macedo, Arnaldo Carneiro, William Costa, Michael Coe, Hermann Rodrigues, and Ane Alencar. Cracking Brazil's Forest Code. *Science*, 344(6182):363–364, 2014.

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

APPENDIX

A Political Resistance

A.1 Data construction

We detail the construction of the donor and candidate panel data, which we use to characterize political resistance in Section 3.

Campaign finance and electoral records

We obtain campaign finance records from the Brazilian Supreme Electoral Court (TSE), which publicly discloses all registered donations for every election since 2002. For each donation, the records identify the donor by CPF (individuals) or CNPJ (firms), the candidate by CPF, the donation amount, the office contested, and the state of the election. We use elections from 2002 to 2018, covering five state and federal legislative cycles (2002, 2006, 2010, 2014, and 2018), as well as the four intervening municipal cycles (2004, 2008, 2012, and 2016).

We process the raw files year by year, standardizing variable names and cleaning CPF and CNPJ identifiers. We strip leading zeros, remove null placeholders, and standardize name strings. We retain individual donors (11-digit CPFs) separately from corporate donors (14-digit CNPJs), and we produce a deduplicated list of all distinct individual donors and candidates appearing in any election from 2002 to 2018. Corporate donors are tracked separately and excluded from the main panel.

We also use the TSE records to construct electoral outcomes. For each candidate, we compute total votes received at the municipality level, collapsing across polling zones. We retain the candidate's declared occupation code to flag self-reported agricultural occupations as a supplementary classification.

Firm registry

We link individual donors and candidates to the firm registry of the *Receita Federal do Brasil* (RF), which we extract from the CNPJ database as of 2019. The registry records, for each establishment, the CNPJ identifier, primary CNAE sector code, state and municipality, legal form, firm size (*porte*), registration in the *Simples* regime,

and equity capital. We also use the individual owners' file (*sócios* and *representantes legais*), which identifies each owner by a partially masked CPF and a standardized name.

We classify a firm as agricultural if its CNAE code falls below 1,100,000, which covers primary agriculture, livestock, fishing, and food and beverage processing including meatpacking. For each individual in the owners file, we aggregate characteristics across all firms the individual owns. We assign the modal sector and municipality, and we code a dummy variable for each individual that measures whether they own at least one agricultural firm. Equity capital is categorized into terciles, and firm size is recorded as micro, small, or large.

Name matching

The RF database masks 6 of the 11 CPF digits, preventing direct identifier matching between TSE and RF records. We instead link individuals using a probabilistic name-matching procedure implemented in Python. For each donor and candidate in the TSE, we compare the standardized name against all records in the RF owners file using a fuzzy string algorithm. We use a high similarity score threshold of 0.9 for a match, and we validate candidates using the partial CPF digits where available. For each matched individual, we assign firm characteristics from the RF record: sector classification, municipality, and infraction history. Table A1 reports characteristics of the matched sample.

Environmental infractions

We use publicly available infraction and embargo records from IBAMA. Each record identifies the recipient by name, a partially masked CPF or CNPJ, date, and municipality. We create three merge keys: (1) masked CPF and name, (2) CNPJ (for corporate entities), and (3) name alone. For each individual, we record the year of the first infraction and the year of the first embargo, which represents a more severe enforcement action involving an asset seizure order.

We link infraction records to TSE donors and candidates using all three strategies in sequence, applying each in turn and filling unmatched records with the next strategy. We classify a donor or candidate as an *infractor* if any strategy yields a match, and as an *embargoee* if they appear in the embargo subfile.

Panel construction

We construct separate balanced panels for donors and candidates. For donors, we focus on the five federal and state legislative election years: 2002, 2006, 2010, 2014, and 2018. Starting from the deduplicated individual donor list, we link each donor to the firm registry and infraction records using the name-matching and merging procedures described above. For each donor-year cell, we sum donations made to all candidates, to agricultural candidates, and to candidates with infraction histories. We winsorize donation values at the 99th percentile within year, donor type, and state, where our goal is to limit the influence of extreme values. We then expand to a balanced panel with one observation per donor-election-year pair, filling in zeros for years in which a donor made no recorded donation. Time-invariant characteristics, including sector, municipality, and infraction status, are imputed within each donor identifier to fill gaps created by missing elections.

For candidates, we apply an analogous procedure using all election years from 2002 to 2022 at two-year intervals, totalling 11 periods. For each candidate-year cell, we merge donations received from each donor type (CPF or CNPJ), as well as total votes and election outcomes from the electoral returns. We fill in zeros for non-participation years, and we impute time-invariant characteristics as we do for donors.

We restrict our main analysis to donors and candidates linked to Legal Amazon municipalities. We do so using an official IBGE crosswalk of municipalities to the Legal Amazon perimeter, as defined by Brazilian federal law. We drop observations with missing, zero-padded, or otherwise invalid CPF values. After these restrictions, the final sample contains 62,987 individual donors and 5,605 candidates (Table 1).

A.2 Validation

We include additional validation checks for Section 3.

Matching

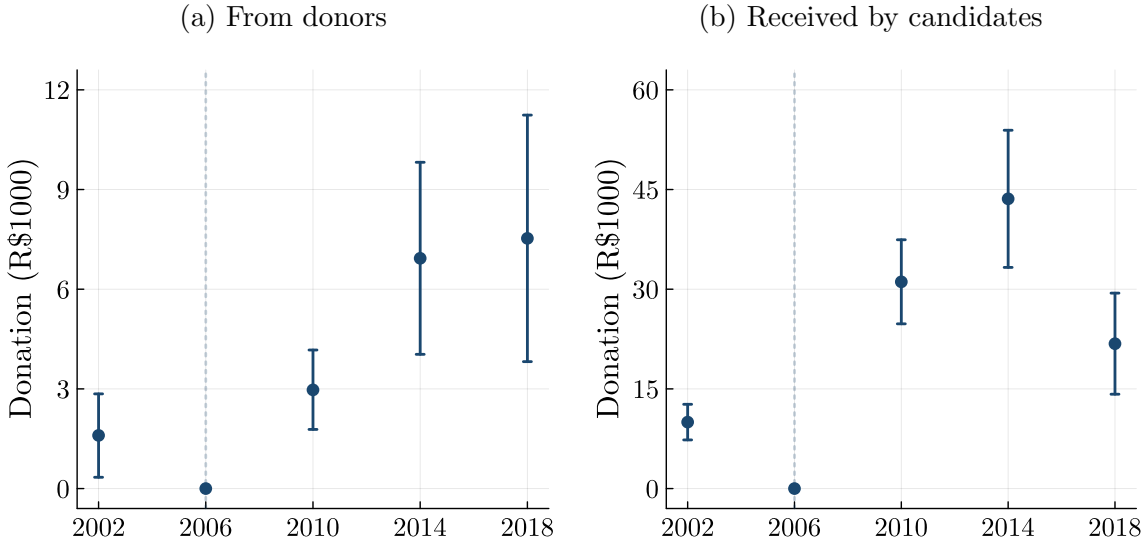
Table A1 assesses covariate balance between agricultural and non-agricultural firms in the full and matched samples. Figure A1 reports difference-in-differences event studies using the complete, unmatched sample, complementing the matched estimates in the main text.

Table A1: Balance of pre-treatment characteristics, agricultural vs. other firms

	Full sample				Matched sample			
	Agri.	Other	Diff.	t-stat	Agri.	Other	Diff.	t-stat
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Donors								
Observations	810	6,524			810	810		
Donation in 2002	2.06	0.52	1.54	2.74	2.06	1.61	0.45	0.68
Firm opening year	2001	2003	-1.75	-4.04	2001	2001	0.18	0.32
Simples tax regime	0.81	0.64	0.17	11.4	0.81	0.80	0.01	0.44
Firm size: micro	0.40	0.54	-0.14	-7.91	0.40	0.40	0	0.10
Firm size: small	0.08	0.09	-0.01	-0.98	0.08	0.10	-0.02	-1.21
LTDA	0.82	0.65	0.17	11.8	0.82	0.83	-0.01	-0.59
Capital: 1st tercile	0.61	0.55	0.07	3.78	0.61	0.58	0.04	1.52
Capital: 3rd tercile	0.34	0.36	-0.02	-1.15	0.34	0.37	-0.03	-1.35
Panel B. Candidates								
Observations	629	2,409			629	629		
Donation in 2002	14.5	3.36	11.2	7.62	14.5	8.82	5.69	3.37
Firm opening year	1999	2002	-2.61	-5.40	1999	1999	0.22	0.34
Simples tax regime	0.84	0.76	0.08	4.85	0.84	0.84	0.01	0.23
Firm size: micro	0.47	0.47	0	0.11	0.47	0.48	-0.01	-0.45
Firm size: small	0.08	0.08	0	0.03	0.08	0.08	-0.01	-0.52
LTDA	0.72	0.60	0.13	6.16	0.72	0.77	-0.05	-1.88
Capital: 1st tercile	0.66	0.62	0.03	1.57	0.66	0.64	0.02	0.71
Capital: 3rd tercile	0.28	0.31	-0.03	-1.60	0.28	0.29	-0.02	-0.62

This table reports balance on pre-treatment firm characteristics for agricultural and other firms. Donations in 2002 are in R$1000. We show the mean difference between agriculture minus other, as well as the t-statistic of the mean equality test. Means for binary variables are proportions. The matched sample pairs agricultural and non-agricultural firms on the characteristics listed. We report two-sided t-tests in columns 4 and 8.

Figure A1: Total donations, difference-in-differences



Event-study coefficients β_τ from (1) with 2006 as the reference year. Points are estimates, and bars are 95% CIs with individual-level clustering. Panel (a) shows the differential effect on average donations from agricultural firm owners. Panel (b) shows the differential effect on average donations received by candidates who are agricultural firm owners. Regressions include municipality-by-election-year.

The Cerrado

Table A2 examines robustness of the national triple-difference to the inclusion and exclusion of the Cerrado biome. A potential concern is that this biome, where soybean production is concentrated, may disproportionately drive the results. To assess this, we report two additional specifications. Columns 1 and 2 exclude municipalities in the Cerrado entirely, comparing agricultural donations inside the Legal Amazon to the rest of the country. The triple-interaction coefficient remains positive and similar in magnitude, indicating that the Amazon-specific increase in agricultural donations does not depend on the Cerrado. Columns 3 and 4 restrict the sample to the Cerrado only, comparing municipalities inside versus outside the Legal Amazon boundary. The point estimates are again consistent with the main findings. Both sets of matched specifications lose statistical significance due to the smaller samples, but the pattern of results is robust across these splits.

Table A2: Donor donations by biome

	No Cerrado		Cerrado only	
	(1)	(2)	(3)	(4)
$Ag \times Post$ 2006	0.34	1.66	1.78***	-0.15
	(0.69)	(1.73)	(0.55)	(3.14)
$Ag \times Post$ 2006 $\times$ Amazon	2.55**	2.10	2.57*	0.72
	(1.23)	(1.49)	(1.42)	(3.05)
Observations	304,175	6,710	63,120	2,760
Fixed effects	my	match + my	my	match + my

Triple-difference estimates of equation (3). Outcomes are donations made by donors. *Amazon* indicates municipalities in the Legal Amazon. Columns 1–2 exclude municipalities in the Cerrado biome; columns 3–4 include only municipalities in the Cerrado. Fixed effects vary by column for municipality-year my and match pairs. Standard errors are clustered at the individual level. Monetary values are in R$1000.

B Politically Optimal Policy

B.1 Fitted shares

Our third approach to constructing the weights w_i^k in equation (13) uses fitted shares from a land allocation model, following Fiszbein (2022). The idea is to predict each site's agricultural composition from agro-climatic fundamentals, purging the influence of idiosyncratic local factors that may be correlated with unobserved determinants of future development.

Let $k \in \{\text{soy, maize, pasture}\}$ denote agricultural activities. For each activity k, we estimate a Poisson regression of land area on local productivity measures in the baseline year:

$$a_i^k \sim \text{Poisson}(\mu_i^k), \quad \mu_i^k = a_i \exp\left(\theta_1^k y_i^{\text{maize}} + \theta_2^k y_i^{\text{soy}} + \theta_3^k y_i^{\text{pasture}}\right),$$

where a_i^k is the land area allocated to activity k in site i, a_i is total agricultural land, and y_i^k are attainable crop yields for soy and maize and agro-climatic suitability for pasture. The term $\log(a_i)$ enters as an offset, ensuring that predicted allocations scale with total agricultural area. Using the fitted values $\hat{\mu}_i^k$, we compute the predicted share of activity k in site i as

$$\hat{w}_i^k = \frac{\hat{\mu}_i^k}{\sum_\ell \hat{\mu}_i^\ell}.$$

These fitted shares reflect the expected land allocation across activities given local agro-climatic conditions. They are used in place of observed shares when constructing agricultural returns r_{it} in equation (13).

B.2 Validation

We validate the main estimates of Table 7. Table B1 estimates the model at the municipality level rather than the 0.25-degree cell level, presenting both long-difference and panel estimates. Table B2 estimates the model with moving average prices rather than contemporaneous yearly prices. Figure B1 shows how the pooled long-difference estimates vary by horizon length, ranging from 1 to 15 years. Figure B2 shows the long-difference estimates for each start year with a 10-year horizon. Figure B3 reports panel estimates with leads and lags, which allow us to assess pre-trends.

Table B1: Agricultural returns and development at the municipality level

	Crop weights		
	Baseline shares	Leave-one-out shares	Fitted shares
	(1)	(2)	(3)
Panel A: Long differences			
Log agricultural returns	1.15***	1.24***	1.72***
	(0.31)	(0.39)	(0.43)
Observations	7,782	7,782	7,782
Fixed effects	$iw + tw$	$iw + tw$	$iw + tw$
Panel B: Panel estimates			
Log agricultural returns	0.66***	0.70***	0.97***
	(0.17)	(0.21)	(0.24)
Observations	8,680	8,680	8,680
Fixed effects	$i + t$	$i + t$	$i + t$

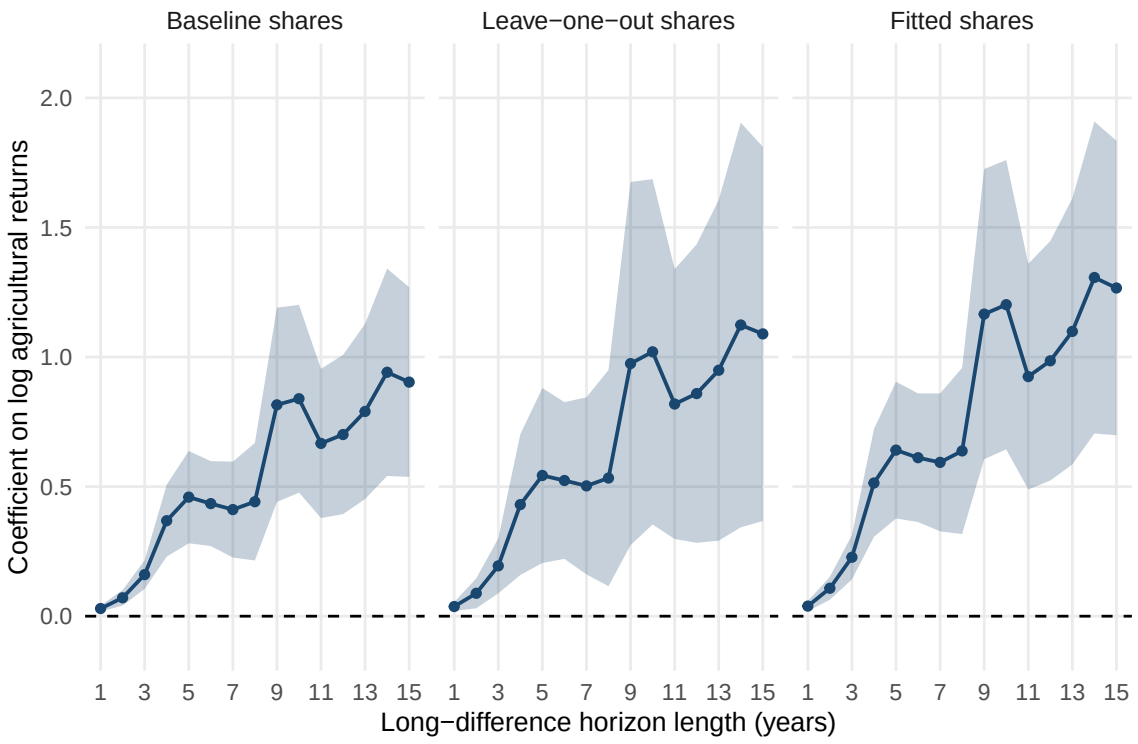
PPML estimates at the municipality level. Panel A presents estimates of equation (12) in pooled 10-year long differences with data from 2001 to 2019, while Panel B uses the full annual panel variation with two-way fixed effects. The dependent variable is agricultural land area with log forest area as an offset. Each column uses a different approach to constructing the weights in equation (13): observed baseline shares, leave-one-out shares from the microregion, and fitted shares from a Poisson land allocation model. Panel A includes site-by-window and year-by-window fixed effects (iw and tw), while Panel B includes site and year fixed effects (i and t). Standard errors are clustered at the microregion level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B2: Agricultural returns and development with moving average prices

	Crop weights		
	Baseline shares	Leave-one-out shares	Fitted shares
	(1)	(2)	(3)
Panel A: Yearly prices			
Log agricultural returns	0.84***	1.02***	1.20***
	(0.18)	(0.34)	(0.28)
Observations	37,274	37,274	37,274
Fixed effects	$iw + tw$	$iw + tw$	$iw + tw$
Panel B: 3-year moving-average prices			
Log agricultural returns	0.92***	1.12***	1.29***
	(0.20)	(0.38)	(0.32)
Observations	28,966	28,966	28,966
Fixed effects	$iw + tw$	$iw + tw$	$iw + tw$
Panel C: 5-year moving-average prices			
Log agricultural returns	0.92***	1.11***	1.28***
	(0.19)	(0.38)	(0.31)
Observations	20,690	20,690	20,690
Fixed effects	$iw + tw$	$iw + tw$	$iw + tw$

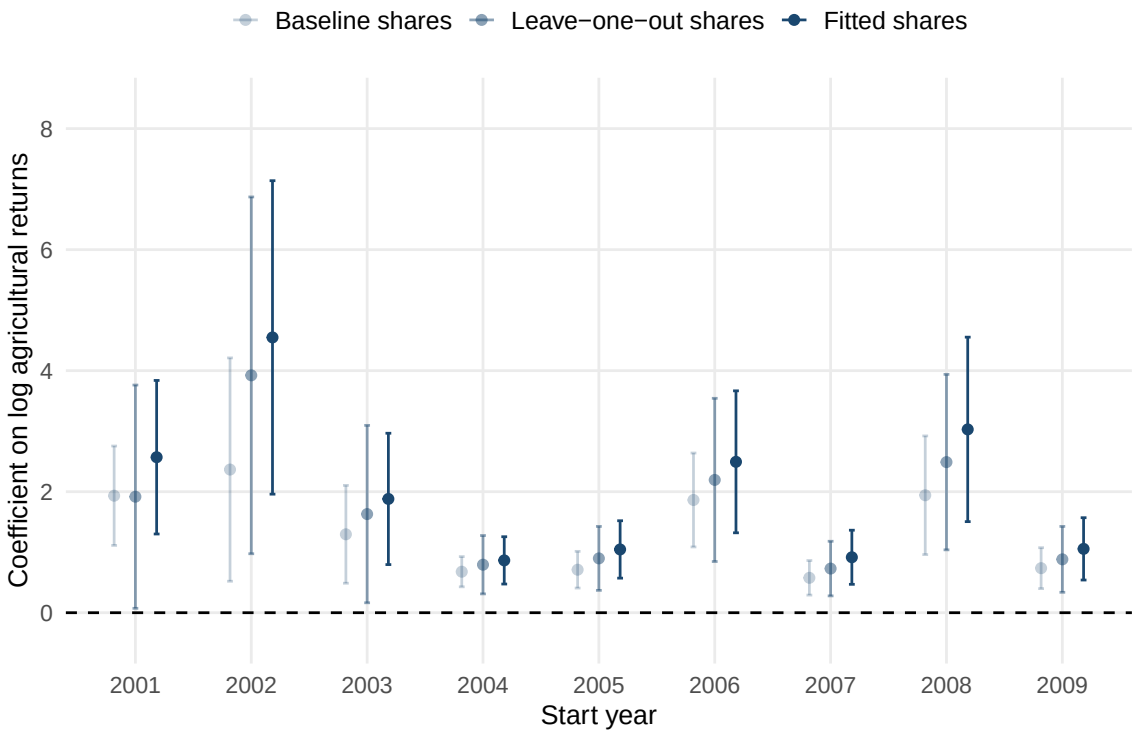
PPML estimates with moving average prices. Panel A presents estimates of equation (12) in pooled 10-year long differences with data from 2001 to 2019, using contemporaneous yearly prices. It matches our main estimates in Table 7. Panel B uses 3-year centered moving average prices, and Panel C uses 5-year centered moving average prices. The dependent variable is agricultural land area with log forest area as an offset. Each column uses a different approach to constructing the weights in equation (13): observed baseline shares, leave-one-out shares from the microregion, and fitted shares from a Poisson land allocation model. All panels include site-by-window and year-by-window fixed effects (iw and tw). Standard errors are clustered at the microregion level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure B1: Agricultural returns and development by long-difference horizon



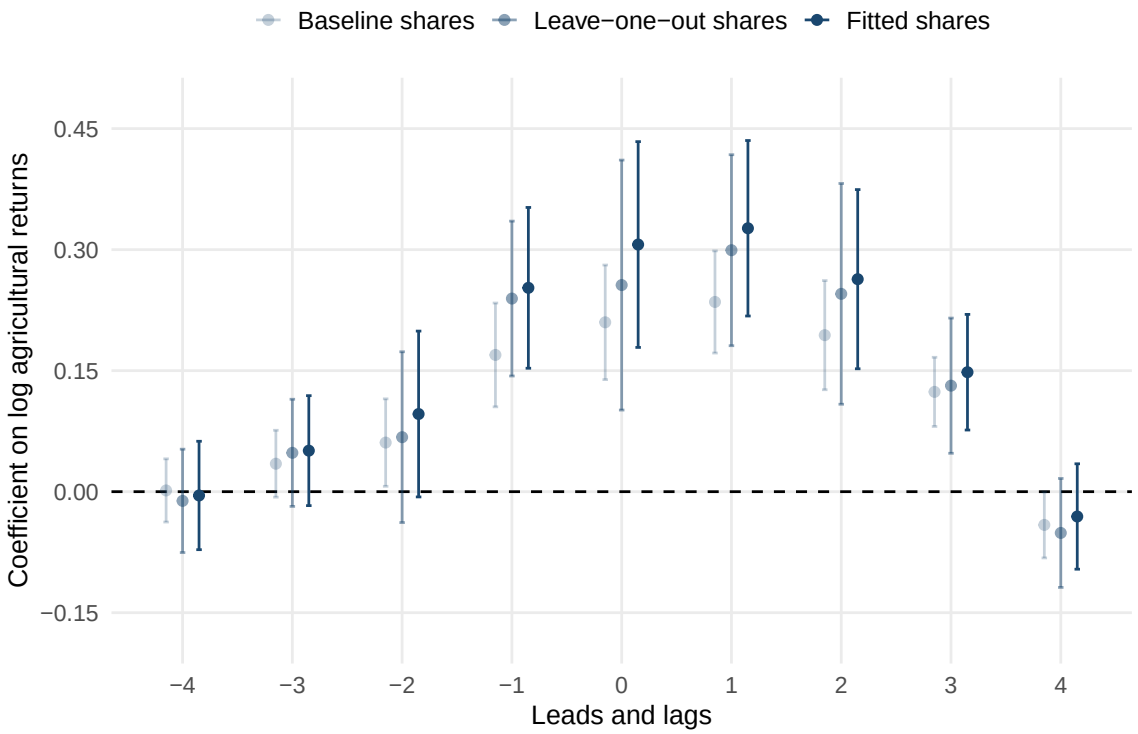
PPML estimates with varying horizon lengths, ranging from 1 to 15 years. For each horizon length, we pool over all available start years. Each panel uses a different approach to constructing the weights in equation (13): observed baseline shares, leave-one-out shares from the microregion, and fitted shares from a Poisson land allocation model. The dependent variable is agricultural land area with log forest area as an offset. All specifications include site-by-window and year-by-window fixed effects (iw and tw). Standard errors are clustered at the microregion level. The shading marks 95% confidence intervals.

Figure B2: Agricultural returns and development by start year



PPML estimates with varying start years, ranging from 2001 to 2009. Each point corresponds to one 10-year long-difference specification. Each set of estimates uses a different approach to constructing the weights in equation (13): observed baseline shares, leave-one-out shares from the microregion, and fitted shares from a Poisson land allocation model. The dependent variable is agricultural land area with log forest area as an offset. All specifications include site and year fixed effects (i and t). Standard errors are clustered at the microregion level. Vertical bars mark 95% confidence intervals.

Figure B3: Agricultural returns and development with leads and lags



PPML estimates with leads and lags, ranging from 1 to 4 years. Each set of estimates corresponds to one panel specification with two-way fixed effects (i and t), using a different approach to constructing the weights in equation (13): observed baseline shares, leave-one-out shares from the microregion, and fitted shares from a Poisson land allocation model. The dependent variable is agricultural land area with log forest area as an offset. Standard errors are clustered at the microregion level. Vertical bars mark 95% confidence intervals.

B.3 Baseline regulation

The main text normalizes baseline regulation to zero and subsumes it in costs. We show that counterfactual allocations and welfare remain well-defined under this normalization. In particular, suppose site i faces baseline regulation $\tau_i^0 \in [0,1]$ in the form of an ad valorem tax on agricultural returns, in addition to non-regulatory costs c_i^0. Agricultural value under baseline regulation is

$$v_i = \log\left(\frac{(1-\tau_i^0)r_i}{c_i^0}\right).$$

We recover equation (11) by defining $c_i = c_i^0/(1-\tau_i^0)$, such that $v_i = \log(\frac{r_i}{c_i})$ with costs c_i absorbing both non-regulatory costs c_i^0 and baseline regulation τ_i^0.

Under a further, counterfactual ad valorem tax τ_i on gross returns r_i, equation (15) continues to hold under either representation. That is, we have $\log\left(\frac{a_i(\tau_i)}{f_i(\tau_i)}\right) - \log\left(\frac{a_i}{f_i}\right) = \beta\log(1-\tau_i)$, where cost c_i cancels whether or not it contains baseline regulation τ_i^0. For fixed total land $n_i = a_i(\tau_i) + f_i(\tau_i)$, we can solve for counterfactual $a_i(\tau_i)$.

$$a_i(\tau_i) = n_i\left(\frac{\frac{a_i}{f_i}(1-\tau_i)^\beta}{1+\frac{a_i}{f_i}(1-\tau_i)^\beta}\right)$$

This counterfactual allocation depends only on estimated parameter β, observed (a_i, f_i, n_i), and counterfactual tax τ_i. Baseline regulation τ_i^0 does not enter.

Turning to welfare, counterfactual producer surplus is still given by equation (16). We rewrite by applying equation (10) and $n_i = a_i(\tau_i) + f_i(\tau_i)$ to the logit inclusive value.

$$\pi_i(\tau_i) = \frac{n_i r_i}{\beta}\log\left(1 + \frac{a_i(\tau_i)}{n_i - a_i(\tau_i)}\right).$$

Emissions are $E(\tau) = \sum_i e_i a_i(\tau_i)$. Baseline regulation τ_i^0 does not enter in either case. Both depend on $a_i(\tau_i)$, which does not depend on baseline regulation as shown above. Finally, counterfactual government revenue derives from levying counterfactual tax τ_i on gross returns r_i. Assuming that baseline regulation τ_i^0 does not generate revenue, we have that baseline $G(0) = 0$ and thus $\Delta G(\tau) = G(\tau) - G(0)$. We can compute $G(\tau) = \sum_i \tau_i r_i a_i(\tau_i)$ without isolating τ_i^0.

B.4 Constrained optimization

The optimal policy in Section 5 is parameterized by a scale σ and a cutoff ρ, which together determine the tax on each site. A natural concern is whether this two-parameter family is too restrictive – that is, whether an unconstrained, site-specific tax vector could achieve substantially higher welfare. We assess this concern by comparing the constrained $(\hat{\sigma}, \hat{\rho})$ solution to the unconstrained solution that maximizes over the full, site-specific vector of taxes. We compute the unconstrained solution with the L-BFGS algorithm, which uses finite-difference gradients to explore the solution space. We use our constrained solution as the starting point, and we run the algorithm for 200 iterations.

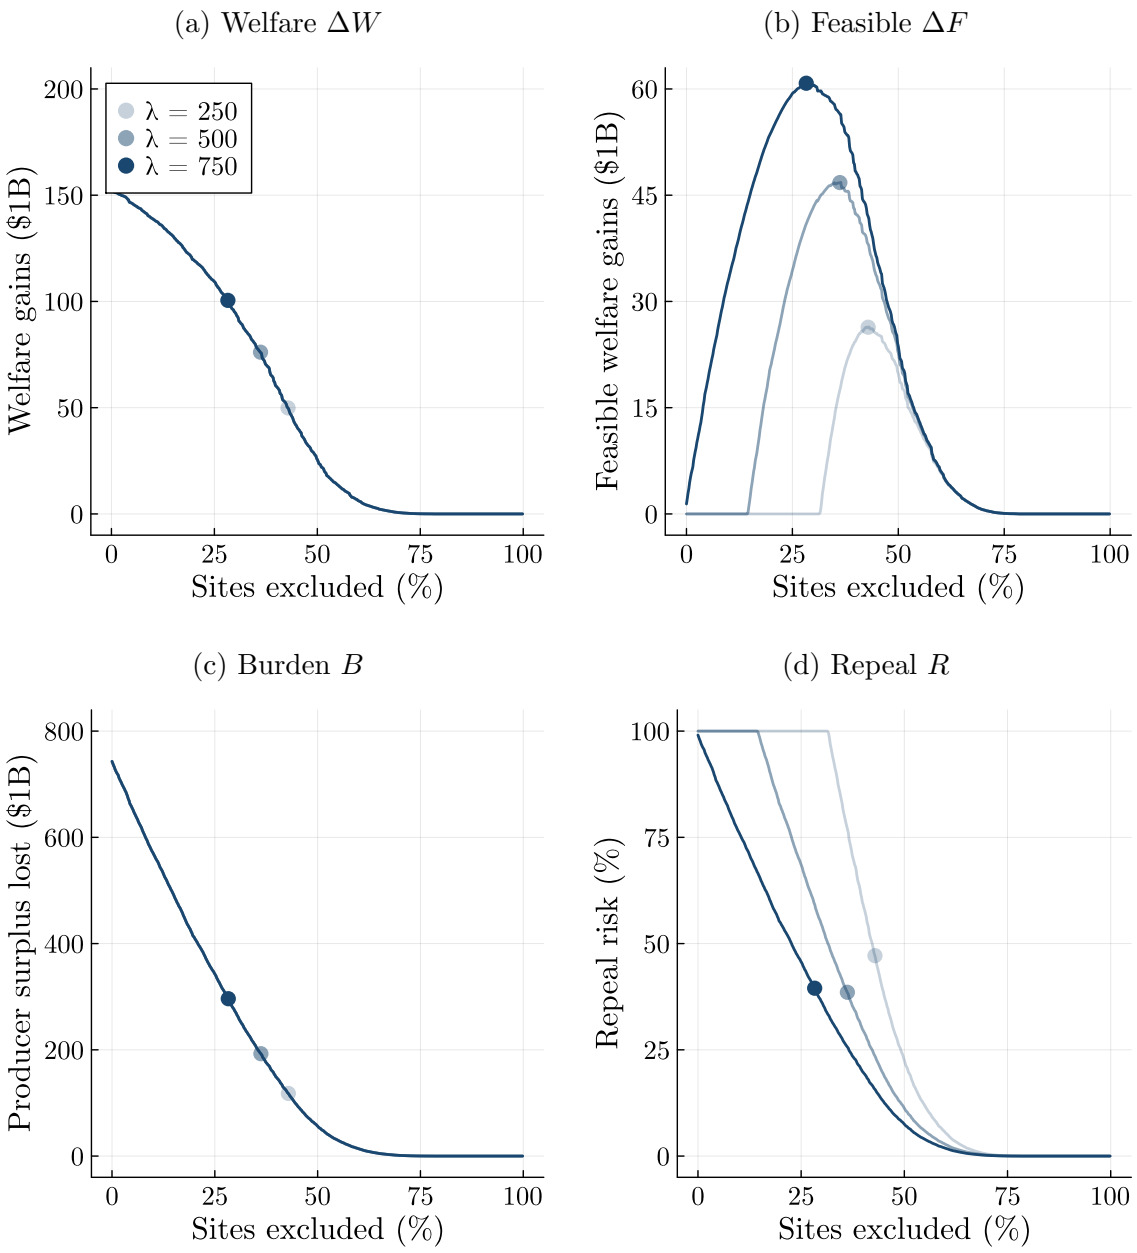
Table B3 presents the results. The first three columns validate that our analytic Pigouvian tax indeed maximizes welfare, as the theory predicts. The last three columns show that the feasible welfare gap between constrained and unconstrained policies is roughly 7 percent across values of κ. L-BFGS is not guaranteed to have converged, as the full tax vector includes more than 2,500 dimensions and thus sits within a very large solution space. Thus, while the gaps that we measure are lower bounds, they suggest that the constrained policy does well in capturing the vast majority of feasible welfare gains. In this sense, our (σ, ρ) parameterization provides a tractable and reasonable approximation to the unrestricted optimum.

Table B3: Constrained vs. unconstrained optimal policy

	Pigouvian (ΔW, \$1B)			Optimal (ΔF, \$1B)		
κ	Analytic	L-BFGS	Gap (%)	$(\hat{\sigma}, \hat{\rho})$	L-BFGS	Gap (%)
25	40.1	40.1	0.0	19.3	20.8	6.8
50	152	152	0.0	53.0	57.4	7.7
75	323	323	0.0	92.3	99.7	7.4
100	541	541	0.0	136	146	7.1

This table compares welfare and feasible welfare gains under constrained and unconstrained policies. Welfare gains are in \$1B. Social costs are set at κ = \$25 to \$100 per ton of CO_2, as shown in the first column. We set $\lambda = 500$ throughout. We solve for unconstrained policies with the L-BFGS algorithm, which uses finite-difference gradients over 200 iterations, starting from the constrained solution.

Figure B4: Solving for optimal policy, cutoff ρ



Counterfactual impacts of Pigouvian taxation at full scale ($\sigma = 100\%$) as a function of the cutoff ρ, which excludes a fraction of sites ranked by their tax-to-surplus-cost ratio. Dots mark the optimal cutoff $\hat{\rho}$ for each value of λ, which governs the sensitivity of repeal risk to producer burden. Lines correspond to λ values of \$250, \$500, and \$750 billion. Social costs are set at $\kappa = \$50$ per ton of CO_2. This figure complements Figure 5, which instead varies the scale σ of Pigouvian taxation, fixing the cutoff at $\rho = 0$.

B.5 Welfare accounting

The main text presents welfare in terms of ad valorem taxes $\tau_i \in [0, 1]$. In implementation, it is convenient to work instead in value units (v-units) with tax

$$\tilde{\tau}_i = -\log(1 - \tau_i) \geq 0,$$

such that ad valorem tax $\tau_i = 1 - \exp(-\tilde{\tau}_i)$. This transformation is useful because v-unit tax $\tilde{\tau}_i$ enters the value function linearly. By equation (11), we have

$$v_i(\tau_i) = v_i + \log(1 - \tau_i) = v_i - \tilde{\tau}_i = v_i(\tilde{\tau}_i).$$

In the logit model, the ratio of agricultural to forested land becomes $\frac{a_i}{f_i}\exp(-\beta\tilde{\tau}_i)$, and the expressions for counterfactual land, producer surplus, and welfare follow from substituting $\tau_i = 1 - \exp(-\tilde{\tau}_i)$ into equations (15) to (18). The main text monetizes producer surplus in equation (16) by multiplying by r_i. Here, we instead have

$$\pi_i(\tilde{\tau}_i) = \frac{n_i}{\beta}\log\left(1 + \frac{a_i}{f_i}\exp(-\beta\tilde{\tau}_i)\right),$$

which gives total $\Pi(\tilde{\tau}) = \sum_i \pi_i(\tilde{\tau}_i)$ in v-units. Government revenue per hectare of agricultural land is v-unit tax $\tilde{\tau}_i$, and so total $G(\tilde{\tau}) = \sum_i \tilde{\tau}_i a_i(\tilde{\tau}_i)$ is also in v-units. Emissions in v-units are $E(\tilde{\tau}) = \sum_i \frac{\kappa e_i}{r_i} a_i(\tilde{\tau}_i)$.[9]

These components deliver a Pigouvian tax of simple form. Summing over the above expressions, we have that welfare in v-units is

$$W(\tilde{\tau}) = \sum_i \left[\pi_i(\tilde{\tau}_i) + \tilde{\tau}_i a_i(\tilde{\tau}_i) - \frac{\kappa e_i}{r_i} a_i(\tilde{\tau}_i)\right].$$

By the expression for $\pi_i(\tilde{\tau}_i)$, we have $\pi_i'(\tilde{\tau}_i) = -a_i(\tilde{\tau}_i)$. Taking the first order condition,

$$W'(\tilde{\tau}_i) = -a_i(\tilde{\tau}_i) + a_i(\tilde{\tau}_i) + \tilde{\tau}_i a_i'(\tilde{\tau}_i) - \frac{\kappa e_i}{r_i} a_i'(\tilde{\tau}_i) = \left(\tilde{\tau}_i - \frac{\kappa e_i}{r_i}\right) a_i'(\tilde{\tau}_i) = 0.$$

We thus obtain Pigouvian tax $\tilde{\tau}_i^P = \frac{\kappa e_i}{r_i}$.

[9] First, we compute emissions in physical units by multiplying carbon stocks e_i, measured in tons of CO_2 per hectare, by agricultural development a_i, measured in hectares. Second, we convert to damages in monetary units by multiplying by social costs of carbon κ, measured in dollars per ton of CO_2. Third, we convert to v-units by dividing by site returns r_i, given log utility.

We largely work in v-units, including when optimizing over welfare measures. The key advantage of v-units is a transfer property. Because $v_i(\tilde{\tau}_i)$ is linear in $\tilde{\tau}_i$, a marginal tax increase reduces producer surplus by a_i v-units for a given site, and it increases government revenue by the same a_i v-units. Taxation is thus a transfer in v-units, and indeed the surplus and revenue terms cancel in the first-order condition above. This transfer property does not hold in dollars. Because ad valorem tax $\tau_i = 1-\exp(-\tilde{\tau}_i)$ is concave in $\tilde{\tau}_i$, the government's marginal dollar revenue $(1-\tau_i)r_i a_i$ falls short of the producer's marginal dollar loss $r_i a_i$ at positive tax rates. As a result, the Pigouvian tax does not take a closed form in dollars. We therefore optimize in v-units throughout.

We convert to dollars in two places. First, because we set parameter λ in dollars, we convert total burden B to dollars when computing the probability $R(B;\lambda)$ of repeal. Second, to facilitate interpretation, we convert welfare gains to dollars when presenting our results. In both cases, we multiply aggregate v-unit quantities by aggregate $\bar{r}$, which provides a simple approximation to multiplication by site-specific r_i. We compute $\bar{r}$ as the area-weighted mean of agricultural returns across sites.