Optimal Environmental Targeting in the Amazon Rainforest*

Juliano Assunção, Robert McMillan, Joshua Murphy, and Eduardo Souza-Rodrigues†

March 26, 2022

Abstract

This paper sets out a data-driven approach for targeting environmental policies optimally in order to combat deforestation. We focus on the Amazon, the world's most extensive rainforest, where Brazil's federal government issued a 'Priority List' of municipalities in 2008 – a blacklist to be targeted with more intense environmental monitoring and enforcement. First, we estimate the causal impact of the Priority List on deforestation (along with other relevant treatment effects) using 'changes-in-changes' (Athey and Imbens, 2006), finding that it reduced deforestation by 43 percent and cut emissions by 49 million tons of carbon. Second, we develop a novel framework for computing targeted optimal blacklists that draws on our treatment effect estimates, assigning municipalities to a counterfactual list that minimizes total deforestation subject to realistic resource constraints. We show that the ex-post optimal list would result in carbon emissions

*We would like to thank Fernanda Brollo, Sylvain Chabe-Ferret, Francisco Costa, Clarissa Gandour, Kelsey Jack, Nick Kuminoff, Charles Manski, Ismael Mourifié, Alvin Murphy, Marcel Oestreich, Martino Pelli, Alex Pfaff, Alberto Salvo, Paul Scott, Edson Severini, Kate Sims, Aloysius Siow, and Aleksey Tetenov for many helpful comments and suggestions. Thanks also to seminar participants at Arizona State University, Brock University, the Instituto Escolhas, Iowa State University, Toulouse School of Economics, University of British Columbia, University of Ottawa, University of Toronto, the 2018 LSE-NHH Conference, the Montreal Workshop in Environmental and Resource Economics, the 2018 NBER Summer Institute, the INRA Environmental and Natural Resources Conservation Workshop in Montpellier, the 25th European Association of Environmental and Resource Economists Conference, and the 6th World Congress of Environmental and Resource Economists for their feedback. Additional comments and suggestions from the Editor and four anonymous referees helped us improve the paper considerably. Faisal Ibrahim and Brenda Prallon provided outstanding research assistance. Financial support from SSHRC and the University of Toronto Mississauga is gratefully acknowledged. Assunção would also like to thank the CNPq for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of their respective organizations. All remaining errors are our own.

†Contact information: Assunção – Departamento de Economia, PUC-Rio, Rua Marques de Sao Vicente, 225/F210, 22453-900 - Rio de Janeiro/RJ, Brazil (email: juliano@econ.puc-rio.br); McMillan – Department of Economics, University of Toronto, 150 St. George Street, Toronto, ON M5S 3G7, Canada, and NBER (email: mcmillan@chass.utoronto.ca); Murphy – Natural Resources Canada, 580 Booth Street, Ottawa, ON K1A 0E4, Canada (email: joshua.murphy4@canada.ca); Souza-Rodrigues – Department of Economics, University of Toronto, 150 St. George Street, Toronto, ON M5S 3G7, Canada (email: e.souzarodrigues@utoronto.ca).

over 10 percent lower than the actual list, amounting to savings of more than $1.29 billion (36% of the total value of the Priority List), with emissions over 23 percent lower on average than a randomly selected list. The approach we propose is relevant both for assessing targeted counterfactual policies to reduce deforestation and for quantifying the impacts of policy targeting more generally.

Keywords: Policy Targeting, Optimal Regulation, Monitoring, Deforestation, Amazon, Carbon Emissions, Changes-in-Changes, Resource Constraints, Partial Identification, Minimax Ambiguity

1 Introduction

In many developing countries, the effective implementation of environmental policies is undercut by weak institutions, as recent research has documented clearly (see, e.g., Greenstone and Jack, 2015). The resulting unregulated and often-illegal activities that prevail can cause severe duress to fragile ecosystems, leading to outcomes that are both damaging and inefficient. As a prominent instance of this phenomenon, several studies (notably by Burgess et al., 2012) have highlighted the role of illegal logging and land clearing as drivers of tropical deforestation, widely understood to be a critical contributor to global carbon emissions (see IPCC, 2013). In settings such as these where existing institutions are over-stretched, targeted monitoring and enforcement policies may be advantageous, helping to focus limited resources where they can have higher-than-average impacts.

This paper measures the causal effects of blacklist-type government regulations – a widespread form of targeting – and then explores how such targeted regulations can be optimized. It does so in the context of deforestation, focusing on the Amazon, the world's most extensive rainforest and a vitally important ecosystem, whose fundamental roles in storing carbon, conserving biodiversity, maintaining water quality and even modulating the Earth's climate are well established (Foley et al., 2005; Bonan, 2008). Deforestation in the Amazon has been a source of international concern for at least the past 30 years, spurring regulatory activity on the part of Brazil's federal government in particular. Regulations introduced in Brazil coincided with a marked slowdown in deforestation between 2004 and 2017, with the annual deforested area falling by 75 percent.[1] As other factors may have been responsible for this decline, changing commodity prices among them, policy makers are interested in understanding the efficacy of actual regulations in reducing deforestation, and how such regulations might be further refined. Yet the literature has not supplied a means to assess, in a systematic quantitative way, which policy configurations would be likely to have most impact in

[1] Annual deforestation has risen more recently, approaching rates last seen in 2008.

limiting future deforestation given relevant constraints. Filling that gap is the central task of this paper.

Our analysis is built around an important regulatory change that occurred in 2008, when Brazil's federal government issued a blacklist of 36 municipalities (out of a total of 524) with especially high levels of deforestation – the so-called 'Priority List.' The listed municipalities were to be the focus of more rigorous monitoring and stricter penalties, with the list being renewed every year subsequently.

The paper's first goal is to estimate the causal treatment effect of the Priority List on deforestation in the Brazilian Amazon. We do so over the period 2006-2010. Given the official criteria did not specify exactly how the list was chosen, we start by investigating the effective selection procedure that assigned municipalities to the Priority List. The patterns we find in the data indicate that the federal government adhered closely to a threshold rule; indeed, we are able to replicate the actual 2008 assignments with 98 percent accuracy using only the inferred rule. The rule essentially separates municipalities based on their deforestation levels but not on their trends: comparing municipalities on the list (versus not) leading up to the reform's introduction supports the assumption of common trends.

Based on that evidence, one could use a standard difference-in-differences ('DID') model to estimate the impact of the reform, noting that other commonly-used treatment effect estimation strategies – matching, regression discontinuity, and instrumental variables – are problematic in light of the threshold rule (as we explain in Section 4).[2] Instead, our approach is to implement the changes-in-changes ('CIC') model proposed by Athey and Imbens (2006) (henceforth 'A&I'), a nonlinear generalization of the DID model to the entire distribution of potential outcomes. In a policy evaluation context with pre- and post-policy periods, A&I show how the difference in the distribution functions of the untreated group before and after treatment can be combined with the distribution function of the treated group before treatment to predict the hypothetical distribution of the treated group in the post-treatment period, absent treatment. (In standard DID, we note the adjustments are to the average, not to the entire distribution function, and are implemented linearly.) Similarly, the counterfactual distribution function of the effects of treatment on the untreated can also be recovered.

Given these features, the CIC model is appealing in our context for two main reasons. First, it can accommodate the possibility that the Priority List was implemented (as is likely) on the group with potentially higher average benefits.[3] Treatment effects are allowed to be heterogeneous

[2] Concerning a regression discontinuity approach, for instance, there are few observations close to the threshold frontier, which limits the accuracy of such a strategy. In addition, a regression discontinuity design does not identify the policy treatment effect of interest in this paper.

[3] The official criteria to enter the Priority List reflect the assumption that deforestation is a persistent process: highly deforested locations in the past are expected to be more likely to be deforested in the future, so concentrating

across units – municipalities in our application – *and* across treatment and control groups, given the two counterfactual distributions can be arbitrarily different. Second, we can estimate the policy impacts on the untreated (the average treatment on the untreated, or 'ATU'), which is necessary when trying to shed light on optimal targeting (our second goal, looking ahead). In contrast, a DID strategy will only identify treatment effects on the treated (the 'ATT') in the presence of heterogeneous effects.[4]

In terms of the main treatment effect results, we find that the Priority List caused substantial reductions in the deforestation rate, cutting it by 43 percent in the short term (the period 2009–2010) relative to the case in which no program was enacted. This reduction amounted to avoided emissions of 34.6 million tons of carbon, with a social benefit of around \$2.5 billion, assuming a conservative social cost of carbon of \$20/tCO_2 (Greenstone et al., 2013; Nordhaus, 2014).[5] Further, there is evidence of heterogeneous treatment effects, with the average effect on the untreated being between 10 and 14 percent of the estimated effect on the treated.[6]

We also investigate the possibility that the Priority List generated *spillovers*, which we think of as the indirect effects of the policy on municipalities that were not treated directly. For example, farmers in untreated municipalities that were geographically close to a Priority municipality and that had experienced substantial deforestation in the past might think monitoring could increase there also, thus modifying their behavior. Accordingly, we split the untreated group in two (labelled the 'spillover' and 'control' groups), depending on whether untreated municipalities were more or less likely to react to the policy intervention, based on proximity and prior deforestation.

Estimates of the CIC model provide evidence of spillover effects, with the spillover group reducing deforestation in response to the intervention: the treatment effect for this group is smaller than the effect on the treated, but greater than the effect on the control group. Once we account for spillovers, 3,050 km^2 of deforestation were avoided in Priority municipalities directly, and 1,102 km^2 in the spillover group, totalling 4,152 km^2 of forested area preserved in 2009-2010 as a result of the program. The total avoided emissions amounted to 49 million tons of carbon, with a social benefit of approximately \$3.63 billion. A realistic assessment of the monitoring costs of the program points to a benefit-cost ratio in excess of 30, which is extraordinarily high by any usual standard.

The paper's second goal is to look beyond the actual policy and compare the Priority List with an ex-post optimal blacklist. To this end, we develop a framework for exploring the assignment of municipalities to an optimized counterfactual list based on information about treatment effects

regulatory effort in highly deforested areas may result in more substantial reductions in total deforestation.

[4]Extrapolating results from the treated group to the untreated under the assumption of homogeneous effects would bias the estimated effects on the untreated and make any ex-post policy calculations unreliable.

[5]This is likely to be a lower bound. Using the EPA's recommended current social cost of carbon estimate would more than double the estimated social benefit.

[6]Although data limitations prevent us from point-identifying the treatment on the untreated, the estimated effect on the untreated is partially identified with informatively narrow identified sets.

drawn from the first part of the analysis. A key benefit of the framework is that it allows us to investigate how knowledge of treatment effects – perhaps only partial in nature – can lead systematically to better-targeted conservation policies.

Here we suppose the federal policy maker assigns municipalities to a counterfactual list with the objective of minimizing either total deforestation (very much in line with the original goal of the Priority List itself) or total carbon emissions; a variety of other social objectives can also be accommodated. The policy maker's decision is analyzed as a treatment choice problem under ambiguity – appropriate given that some treatment effects are not point-identified – and we use the minimax criterion, assuming the policy maker chooses the ex-post list in order to achieve the best of the worst outcomes (see Manski, 2000, 2005). Further, to incorporate limited monitoring resources into the minimization problem, we consider two alternative constraints, one restricting the total area that can be monitored, and the other, the total number of municipalities on the list.[7]

Accounting for spillover effects, we show that the Priority List resulted in carbon emissions that were *at least* 12 percent higher than the ex-post optimal lists (under either constraint), while randomly selected lists of municipalities would result in emissions that were over 30 percent higher on average. The avoided emissions translate into a lower bound for the additional social value of the optimal list when compared to the Priority List of approximately \$1.29 billion over the period 2009-2010: this lower bound is approximately 36% of the total value of the Priority List.

The geographic distributions of the optimal lists reveal interesting patterns not imposed during estimation. First, the overlap between pre-existing protected areas and the area-constrained counterfactual list is much lower than that between protected areas and the original Priority List, suggesting that these two existing policies could have been made to work together more effectively.[8] Second, ignoring spillover effects, the area-constrained counterfactual list is relatively contiguous and forms a protective shield close to the deforestation frontier, which may impede the deforestation process from continuing into more pristine areas, yielding longer-term benefits. Third, when accounting for spillovers, the area-constrained optimal list becomes more geographically dispersed and less contiguous; intuitively, placing all targeted municipalities together fails to exploit the potential reduction in deforestation in adjacent locations due to spillovers.

Beyond the current application, the approach we develop is relevant for assessing counterfactual targeted policies to reduce deforestation in other contexts, based around actual policy interventions. Those interventions can be used to recover heterogeneous policy impacts, our approach then allowing researchers to trace out the quantitative implications for forest cover and carbon emissions

[7]We set the constraints at the same values as those corresponding to the Priority List, then investigate the effects of relaxing these constraints. While information about the resources effectively allocated to monitoring is very difficult to obtain, it is nevertheless reasonable to presume that the larger the area monitored or the greater the number of municipalities monitored, the higher the monitoring costs would be.

[8]A similar implication can be drawn from evidence that compares protected area policies and payments for ecological services from Mexico (see Alix-Garcia et al., 2015).

when policy makers face realistic resource constraints and knowledge of only partially identified estimates. The approach also offers a coherent framework for assessing the quantitative impacts of policy targeting more generally, as we discuss below, using credible estimates based on flexible treatment effects estimation.

The rest of the paper is organized as follows: The next section places our analysis in the context of the existing literature. Section 3 sets out relevant institutional background; Section 4 describes the data, along with descriptive evidence motivating the empirical model, which is then presented in Section 5. We provide model estimates in Section 6. Section 7 develops our counterfactual framework and presents results from the counterfactual targeting exercises, and Section 8 concludes.[9]

2 Relation to the Literature

Our paper contributes to four main areas of research. The first is a growing body of work examining the implementation of environmental policies in a developing country context: see Greenstone and Jack (2015) for a recent survey. This is a complement to the vast literature that studies such policies in a developed country context (see, e.g., the survey by Gray and Shimshack, 2011). Greenstone and Hanna (2014) argue that weak institutional arrangements in developing countries pose obstacles to effective law enforcement, showing that policies targeting improvements in air and water quality in India had varying degrees of success. In the case of climate issues, very much linked to the deforestation process analyzed in this paper, the available evidence is limited (Burke et al., 2016). Our analysis examines a widespread form of targeting and connects the causal impacts of targeting policies to the release of carbon into the atmosphere, thus contributing to both lines of research.[10]

Second, several papers examine the impact of monitoring and the role of institutions in the Amazon itself, notably Assunção et al. (2017), Assunção et al. (2019), and Burgess et al. (2019). Payments for ecological services programs have also been studied, as alternatives to command-and-control policies by Alix-Garcia et al. (2015), Jayachandran et al. (2017), Jack and Jayachandran (2018), and Simonet et al. (2019). Compared with these papers, our analysis examines the effectiveness of an optimized counterfactual policy-targeting strategy as a way of overcoming institutional and political obstacles. This type of targeted strategy can be applied in other contexts that en-

[9]The paper's Supplemental Material includes: (a) information about data sources and the construction of key variables; (b) a discussion of the Priority List selection equation; (c) evidence of possible channels linking the Priority List to deforestation; (d) a comparison of the CIC and DID models and resulting estimates; (e) how we incorporate dynamic treatment effects into our calculations; (f) an explanation of how the counterfactual optimal lists are computed in practice; and (g) several robustness exercises.

[10]Our approach also complements an earlier theoretical literature in environmental economics studying targeted regulatory strategies – see Harrington (1988) and Friesen (2003). We show how a regulator can target resources in an optimal way subject to realistic constraints using credible treatment effect estimates. Our approach allows the associated benefits to be quantified directly, on the basis of econometric evidence.

compass a substantial portion of global rainforest cover – in other parts of Amazonia, the Congo and Southeast Asia.[11]

Within the Amazon context, several recent papers have examined the effects of the Priority List: Arima et al. (2014), Cisneros et al. (2015), Andrade and Chagas (2016), Assunção and Rocha (2019), Koch et al. (2019), and Harding et al. (2021). Those studies use difference-in-differences and matching methods to obtain average treatment effects on the treated similar in magnitude to the corresponding estimates in our study. In contrast to those papers, our estimation approach using changes-in-changes allows us to recover the effects of treatment on the untreated, which we use in computing optimally targeted blacklists.

A third main strand of literature investigates the underlying causes of land use change, tropical deforestation being one important instance. This literature has considered changes in population, infrastructure, agricultural prices, political economy factors and climate-related phenomena.[12] Our results indicate that monitoring policies are important drivers of land use change and deforestation, affecting not only the municipalities that are targeted directly but also generating spillovers for neighboring areas.

Fourth, our counterfactual analysis draws on a burgeoning literature in econometrics studying treatment choice under ambiguity. Our paper is close in spirt to Manski (2000, 2006, 2010), Cassidy and Manski (2019), and Manski et al. (2021), in which decision-making under ambiguity stems from partial identification of treatment effects.[13] Few empirical applications have appeared thus far, aside from some empirical illustrations presented in existing methodological papers.[14] In an applied econometric context, our analysis is novel in that there is no study in which all the following hold simultaneously: (a) unconfoundedness assumptions fail, so that the treatment effects and the welfare objective function are partially identified; (b) the estimation of treatment effects accounts for violations of the 'Stable Unit Treatment Value Assumption' (or SUTVA); (c) the treatment choice is made under ambiguity (and also allows for spillover effects, again in violation of SUTVA); and (d) the set of admissible policies must satisfy binding capacity constraints.

---

[11] Jack and Jayachandran (2018) consider how targeting could be incorporated into the design of payments for ecological services (through manipulation of enrollment costs) to improve the cost-effectiveness of such programs. Targeting social programs in contexts other than environmental conservation is of broad interest; see, e.g., Hanna and Olken (2018) for a comprehensive discussion in the context of cash transfer programs to reduce poverty.

[12] See Stavins (1999), Pfaff (1999), Andersen et al. (2002), Mason and Platinga (2013), and Souza-Rodrigues (2019), among others.

[13] This is in contrast to another important line of research that focuses on Wald's statistical decision theory, in which ambiguity stems from sampling variation. Important work in that area includes Manski (2004), Bhattacharya and Dupas (2012), and Kitagawa and Tetenov (2018), among others – see the survey by Hirano and Porter (2020).

[14] One important empirical study is the analysis by Dehejia (2005), who examines the Greater Avenues for Independence (GAIN) program that began in California in 1986.

3 Institutional Background and Regulations

Our setting is the Brazilian Amazon, which accounts for two-thirds of the Amazon Rainforest and is itself a vast area, almost ten times the size of California. Prior to the 1960s, this area was barely occupied; access was open, and local economic activities were based largely on subsistence and the extraction of rubber and Brazil nuts (for more detail, see e.g., Souza-Rodrigues (2019)). During the 1960s and 1970s, Brazil's military dictatorship promoted the occupation of the Amazon with the explicit goals of securing national borders and developing the region, although economic recession and hyperinflation in the 1980s led government investment to be cut. In the late-1980s, ecological concerns started to shape policies in the region, with IBAMA (the Brazilian Environmental Protection Agency) being set up in 1989, acting as the national environmental police authority tasked with investigating and sanctioning environmental infractions.

In terms of the legal environment, around half the Amazon is protected – either indigenous land or conservation units such as national parks, extractive reserves, and areas of ecological interest. Deforestation in those areas is subject to strict regulations. The rest of the Amazon comprises undesignated public land where no deforestation is allowed, or private land, which accounts for approximately 20 percent of the total area (according to the Agricultural Census of 2006), where deforestation has to follow the rules of the so-called Forest Code. This code states that farms in the Amazon must preserve 80 percent of their area in the form of native vegetation, among other requirements. While deforestation on private land can be legal if it is both authorized and accords with the Forest Code, empirical evidence suggests that compliance with the code is limited (Borner et al., 2014; Rajão et al., 2020), and that most deforestation in the Amazon is illegal.

Regarding environmental monitoring, IBAMA's operations in the Amazon up to the mid-2000s were based largely on information collected and processed by IBAMA's headquarters and regional offices; land and air patrols used at the time proved limited in their effectiveness given the sheer extent of the area covered and risks posed to law enforcers. The adoption of satellite-based monitoring from the mid-2000s onward improved patrolling capabilities considerably. New monitoring procedures were set out in the first stage of the Action Plan for the Prevention and Control of Deforestation in the Legal Amazon (PPCDAm), launched in 2004.[15] Central to PPCDAm law enforcement was the use of high-frequency remote sensing technology in the form of a satellite-based system, DETER, developed by the Brazilian Institute for Space Research (INPE). This greatly increased the capacity to monitor forest-clearing activities in the Amazon, allowing land use images to be processed on a frequent basis, detecting areas experiencing a loss of forest cover, and in turn triggering DETER deforestation alerts for the attention of law enforcers. Since its introduction,

[15] The PPCDAm also led to the expansion of protected areas, mostly during its first phase (spanning 2004–2007), before the first municipalities were assigned to the Priority List in 2008.

this satellite-based system has had a significant impact in slowing deforestation, with Assunção et al. (2017) estimating deforestation would have been over 3.6 times higher in the absence of the system.

The Priority List. In 2008, the government launched the second phase of the PPCDAm, creating a blacklist to better target regulatory effort in order to combat illegal deforestation. Any Amazon municipality could be added to what became known as the 'Municípios Prioritários' List (for convenience, the 'Priority List'). Municipality-level selection criteria for this list were based on (1) total deforested area, (2) total deforested area over the past three years, and (3) the increase in the deforestation rate in at least three of the past five years; the exact rules have to be inferred as they are not in the public domain.[16] Municipalities placed on the Priority List were then subject to more intense environmental monitoring and law enforcement, as well as a raft of costly regulations (MMA, 2009).[17]

The Ministry of the Environment's Ordinance 28, issued in January 2008, listed 36 municipalities making up the initial Priority List – 7 percent of the total number of municipalities in the Brazilian Amazon. The original list was expanded to include an additional seven municipalities in 2009. Six more were placed on the list in 2011, followed by a further two in 2012. By that stage, just six municipalities had been removed from the list (one in 2010, a second in 2011, and four in 2012), and a further eight municipalities were then added in 2017. In total, 59 municipalities were eventually placed on the Priority List between 2008 and 2017, while 467 municipalities were never placed on it during the same interval.

4 Data and Descriptive Evidence

To study the impacts of the targeting policy, we have assembled a municipality-year panel data set that combines information about Priority status, land use, and other possible determinants of deforestation. Our analysis focuses on the time period 2006 to 2010, with the pre-treatment covering 2006 and 2007, and the post-treatment years being 2009 and 2010.[18] The official list of Priority municipalities comes from the Ministry of the Environment, the treatment group comprising municipalities that entered and remained on the list from 2008 to 2010 inclusive. This gives a total of

[16] The legal basis for targeting certain municipalities was set out in the Presidential Decree 6,321 in December 2007. Exiting the Priority List depended on reducing deforestation in a significant way and having at least 80 percent of the municipal private area registered in the Rural Environmental Registry system (Arima et al., 2014).

[17] These included more stringent conditions applying to the approval of subsidized credit contracts, and the requirement to develop local plans for sustainable production (see Maia et al., 2011). Private land titles were also revised in a bid to identify fraudulent documentation and illegal occupancy, and licensing requirements were made stricter for rural establishments.

[18] Comparing deforestation before the first phase of the PPCDAm in 2004 and after the implementation of the Priority List would capture the combined effects of both regulatory changes.

35 municipalities, as one exited in 2010. The untreated group consists of the set of municipalities that did not enter the list before 2010.[19]

Our municipality-year panel includes annual measures of the incremental deforestation, cumulative deforestation, and forested area remaining in each municipality; these are all drawn from the Brazilian government's satellite-based forest monitoring program, PRODES. Factors affecting deforestation in our dataset, other than being on the Priority List, include rainfall and temperature (generated by Matsuura and Willmott, 2012), protected areas (from the National Register of Conservation Units, maintained by the Brazilian Ministry of the Environment), prices of beef and crops (from the State Secretariat for Agriculture and Food Supply), 'local' gross domestic product (from the Brazilian Institute of Geography and Statistics, IBGE), soil quality (as measured by FAO-GAEZ agricultural suitability indices), and distance to the nearest port (from the 2006 National Highway Plan from the Brazilian Ministry of Transportation). We have also assembled data on crop area and the number of cattle (based on the Municipal Crop Survey and the Municipal Livestock Survey, produced by IBGE), deforestation alerts (from the Real-Time System for the Detection of Deforestation, DETER), fines issued (from IBAMA), and measures of the aboveground carbon stock (calculated by Baccini et al., 2012). We dropped observations with minimal remaining available forested area (less than 6 km^2) – small municipalities mostly located at the extreme eastern and southeastern regions of the Amazon Biome, which are not especially relevant for policies focused on preventing deforestation. In total, we have a balanced panel of 490 (out of a possible 524) municipalities within the Amazon Biome. (For further information about data sources and the construction of key variables, see Section A of the Supplemental Material.)

Around 20 percent of the Brazilian Amazon had been deforested to date – over 700,000 square kilometres, an enormous area (larger than Texas, for example). Cleared areas are used mainly for agriculture: approximately two-thirds of the deforested area comprises pasture, and around 8 percent is used for crops – see Almeida et al. (2016).[20] The evolution of deforestation is shown in Figure 1, revealing two pronounced downward steps coinciding with the start of the main phases of the PPCDAm, in 2004 and 2008 (indicated by the two vertical lines), with the deforestation rate then stabilizing over the next five years. In total, annual deforestation declined by approximately 75 percent over the period 2004-17.[21] The figure also presents initial evidence relating deforestation

[19] We focus on the *initial* list for two reasons. First, the calculation of the optimal list requires estimates of treatment effects conditional on covariates for individual units, yet combining groups of units that received treatment at different times in the estimation procedure would yield a weighted sum of different estimated average treatment effects (see, e.g., Goodman-Bacon, 2018), biasing our analyses of the counterfactual optimal list. Second, while in principle it might be possible to estimate causal effects that varied according to the year municipalities were placed on (or taken off) the list, few municipalities entered or exited from 2009 on, so there is not much that can be said with any accuracy about policy impacts in those cases.

[20] Almeida et al. (2016) also show that 20 percent of the cleared area currently takes the form of secondary vegetation. The remaining areas consist of mining, urban areas, 'unobserved' (i.e., areas whose land usage cannot be interpreted due to cloud cover or smoke from recent forest burning), and 'other.'

[21] Total incremental deforestation by year is shown in Table H1 in the Supplemental Material, together with other

to possible contributing factors. Trends in the international prices of soybeans and beef (measured on the rightmost vertical axis) suggest a positive correlation between deforestation and prices prior to 2008; this accords with the fact that most of the deforested area in the Brazilian Amazon is used for pasture (grazing cattle being reared mainly for beef) and crops (primarily soybeans and corn). After 2008, the correlation appears to be much weaker, consistent with the notion that the Priority List helped to preserve the rainforest even when international prices were rising.

The location of the municipalities on the Priority List within the Amazon is displayed in panel (a) of Figure 4; panel (b) shows the Priority List together with protected areas. Those municipalities are found mostly in the Amazon's southern and eastern regions ("the Arc of Deforestation"). Figure 2 shows where incremental deforestation occurred in 2006 and 2010 (the first and the last years used in our econometric analysis), together with Priority municipalities overlaid. It draws attention to the geographically persistent nature of the deforestation process, similar patterns holding in other years of the sample period. The geographic persistence suggests that a targeted policy might be effective, concentrating monitoring and enforcement in locations where deforestation was more likely to occur.

Selection onto the Priority List. The Priority status of a municipality depends on three official selection criteria (as noted): (1) the total amount of forested land cleared in municipality m from its inception up to and including year $t-1$ (which we label Z^1_{mt-1}); (2) the amount of forested land cleared in municipality m in the three-year period ending in year $t-1$ (Z^2_{mt-1}); and (3) an indicator showing whether municipality m experienced year-on-year growth in new deforestation at least three times in the five-year period ending with year $t-1$ (Z^3_{mt-1}). The first two criteria relate to longer-run and more recent deforestation, respectively, while the third captures whether deforestation accelerated in recent years.

Under the assumption that these variables fully determine Priority status, the selection equation can be written, generally, as:

$$G_{mt} = g\left(Z^1_{mt-1}, Z^2_{mt-1}, Z^3_{mt-1}\right), \tag{1}$$

where the indicator $G_{mt} \in \{0,1\}$ captures whether municipality m is on the Priority List in year t or not. As the precise rules determining selection are not stated publicly, we seek to infer them by exploring the extent to which Priority status is determined by the vector summarizing the three criteria, $Z_{mt-1} \equiv \left(Z^1_{mt-1}, Z^2_{mt-1}, Z^3_{mt-1}\right)$.

Figure 3 plots all combinations of Z^1_{mt-1} and Z^2_{mt-1} for a given value Z^3_{mt-1}: the scatterplot in panel (a) holds $Z^3_{mt-1} = 0$, while the scatterplot in panel (b) holds $Z^3_{mt-1} = 1$. In both panels,

information – the number of fines issued, the annual expansion of protected areas, and the number of municipalities added to the Priority List.

municipality-year observations with $G_{mt} = 0$ (i.e., not on the list) and $G_{mt} = 1$ (on the list) are marked with crosses and dots, respectively. The two panels indicate that regulators adhered closely to a threshold rule involving the first and second criteria: both Z^1_{m2007} and Z^2_{m2007} had to cross pre-determined thresholds in order for municipality m to qualify for the Priority List, while the third criterion (whether deforestation accelerated in recent years) is not important. Specifically, the thresholds drawn in both panels of Figure 3 are 2,137 km^2 for Z^1_{mt-1} and 222 km^2 for Z^2_{mt-1}. Using only these inferred thresholds, we are able to replicate the actual 2008 assignments with 98 percent accuracy – evidence that a strict threshold selection rule was followed in practice.[22]

One important consequence of this threshold rule being followed is that additional factors such as local political influence are unlikely to lead to manipulation close to the relevant thresholds that determine the Priority List's initial composition. This is perhaps surprising, given evidence that corruption is an important and widespread problem in Brazil, with well-documented consequences for deforestation – see, for example, Cisneros et al. (2013). We also find that the fraction of municipalities in which the mayor is affiliated with the political coalition of the Brazilian President is the same among Priority and non-Priority municipalities (approximately 40 percent in each group), suggesting the policy was not used as punishment against political enemies at the local level.[23]

The empirical form taken by the selection function $g(\cdot)$ also has implications for the viability of several widely-used identification strategies. Because there is very little overlap in the data among Priority and non-Priority groups given Z_{mt-1}, selection-on-observables techniques (e.g., propensity score matching) are problematic in this context. The applicability of regression discontinuity (RD) designs also runs up against the fact that there are few observations close to the threshold frontier (in addition to which an RD does not identify the policy treatment effect of interest in this paper). Further, while the criteria variables in Z_{mt-1} might seem to be natural instruments for Priority status, they are invalid when the unobservables affecting deforestation decisions are serially correlated. These considerations will help motivate the estimation approach presented in the next section.

Spillovers. In the main analysis, we consider the possibility that the Priority List generated spillovers – that is, indirect impacts on municipalities that were not treated directly by the Priority List. Such spillovers could work in at least two distinct ways. On the one hand, by concentrating monitoring in areas where a disproportionate amount of deforestation occurred (so-called 'hot

[22] These thresholds were estimated using a classification tree algorithm; see Section B of the Supplement for details.

[23] We are grateful to Fernanda Brollo for generously sharing the political coalition data. Of note, Pailler (2018) finds that, on average, deforestation rates increase in election years when an incumbent mayor runs for re-election (in our setting, 2008 was a mayoral election year), and finds no significant changes to deforestation in the years leading up to, or following, the election year. This suggests that re-election incentives did not affect deforestation during the pre-treatment period (2006-2007) nor during the post-treatment years (2009-2010).

spots'), the intervention might simply shift deforestation rather than reduce it, depending on the costs involved – the problem of 'leakage' in the regulation literature. On the other hand, farmers in untreated municipalities might deforest less if they expected the intervention to increase monitoring in non-targeted locations, reasonable given that regulators could benefit from spatial economies of scale in monitoring and also might place highly deforested municipalities on the blacklist in the near future.

As a precursor to investigating whether spillovers are present, we split the untreated group in two according to whether such municipalities are more or less likely to react to the policy intervention. Specifically, we consider two plausible conditions for designating 'spillover' municipalities: whether a municipality (i) shares a border with a treated municipality, and (ii) it had high levels of deforestation historically. We define the second condition on the basis of the threshold criteria that were (implicitly) adopted by the Brazilian government, shown in Figure 3.[24] The group of untreated municipalities satisfying both conditions – being a neighbor of a Priority municipality and having high levels of past deforestation – is then designated the 'spillover' group.[25]

Comparing Treated versus Untreated Municipalities. Next, we compare Priority and non-Priority municipalities. We first consider a 'baseline' cross section from 2007, right before the policy's introduction, grouping all untreated municipalities together.

Table 1 shows the summary statistics. It makes clear that treated and untreated groups differ in important ways, as expected. In Priority municipalities, incremental deforestation and total historical deforestation – the first selection criterion considered by the Ministry of the Environment when assigning Priority status – are higher. Priority municipalities are also larger and have higher local agricultural GDP, higher carbon stocks per hectare, and are subject to more stringent policy measures.

When we break the untreated municipalities further into 'Spillover' and 'Control' groups, the summary statistics in the table confirm that the spillover group falls between the treated and control groups in virtually all instances. In terms of the evolution of deforestation, the profiles across the three groups are similar, especially after 2005, and present no signs of any anticipation effects. (See Figure H6 in the Supplemental Material, which presents (a) deforestation in levels and (b) the log odds ratio of deforestation shares across the three groups – the outcome variable in our empirical framework, as explained below in Section 5.) To summarize: while deforestation slowed down for all three groups after 2008, the slowdown among Priority municipalities is more prominent

[24] That is, a municipality is deemed to have 'high levels of historical deforestation' if Z^1_{mt-1} and Z^2_{mt-1} exceed 70 percent of the thresholds – that is, whether $Z^1_{mt-1} \geq 0.7 \times 2,137$ km^2 and $Z^2_{mt-1} \geq 0.7 \times 222$ km^2.

[25] The empirical results presented in Sections 6 and 7 are robust to diffferent definitions of how close past deforestation is to the threshold criteria (see the Supplemental Material, Section G). Municipalities with deforestation levels near the selection thresholds but without a treated neighbor might also react to the Priority List in anticipation of stricter monitoring, yet only 4 municipalities satisfy this condition, so we cannot split the untreated group further.

than for spillover units which, in turn, is more pronounced than for control units. The evidence supports the view that the selection rule effectively separated municipalities based on their levels of deforestation shares but not on their trends: that is, the third selection criterion, Z^3_{mt-1}, does not help predict Priority status (as already noted).

5 Empirical Framework

In this section, we present an empirical framework for studying targeted environmental regulations in the Amazon. This is based on the changes-in-changes (CIC) model proposed by A&I, which provides a nonlinear generalization of the difference-in-differences (DID) model to the entire distribution of outcomes.

Our empirical approach is shaped by certain data constraints. We focus on municipal-level deforestation and treat that as a function of the regulatory environment (among other factors – e.g., commodity prices and local climatic conditions), given we do not observe the land use decisions of individual farmers but instead have land-use panel data at the municipal level. On the policing front, we have only limited information about the intensity of monitoring, so we use a binary measure of treatment – assignment to the Priority List – and follow a treatment effects approach, noting that modeling the micro-level decisions of farmers and regulators is not feasible. The data constraints notwithstanding, our approach enables us to recover causal treatment effects based on aggregate data and credible policy variation.

5.1 The Changes-in-Changes Model

We make use of standard potential outcomes notation, with capital letters denoting random variables and lower-case letters denoting corresponding realized values. Each municipality m belongs to group $G_m \in \{0, 1\}$, where group 0 is the control group and group 1 is the treatment group – extensions to more than two groups are straightforward. Let A_{mt} denote the total forested area in municipality m at the beginning of year t, and let D_{mt} be the amount of deforestation occurring in m during that year. The share of newly deforested area Y_{mt} is the ratio of D_{mt} to A_{mt}. We use superscript $j \in \{0, 1\}$ to indicate the potential outcome that arises under the policy regime j. The observed share of deforestation for municipality m at time t can then be written:

$$Y_{mt} = (1 - G_m) \times Y^0_{mt} + G_m \times Y^1_{mt} \, .$$

The potential share of deforestation is given by the specification:

$$Y^j_{mt} = h^j \left(X_{mt}, U_{mt}, t \right),$$

for $j \in \{0,1\}$. X_{mt} is a municipality-level vector of observed factors, including prices and agro-climatic conditions (see Section 4). U_{mt} is a municipality-level unobservable term that can incorporate municipality fixed effects (reflecting permanent differences across m in terms of, say, unmeasured soil quality, climatic conditions, or topography) in addition to time-varying unobservables; for instance, we allow for, but are not restricted to, a decomposition of the form $U_{mt} = \alpha_m + \eta_{mt}$. The function h^j depends on t, which allows for flexible time trends. In terms of the impact of the policy, one might expect $h^1(x,u,t) \le h^0(x,u,t)$ for any (x,u,t), given that the Priority List increases monitoring and enforcement intensity. Because we do not restrict the way in which the functions h^j are affected by treatment status j, municipalities at different stages of the deforestation process may respond differently to the policy intervention, giving rise to heterogeneous treatment effects.

We impose four assumptions on the model. Following A&I, we make

Assumption 1 Monotonicity: *The functions $h^j(x,u,t)$, for $j \in \{0,1\}$, are strictly increasing in u.*

Imposing strict monotonicity of h^j on the unobservable u restricts all unobservable factors influencing deforestation shares to be captured by a single scalar; while it involves a loss of generality, the same restriction applies to the standard DID model, and is common in empirical work more generally.[26] In our context, it implies that a location more suitable for agricultural activities (and so more prone to deforestation) is associated with a higher value of u (i.e., has a higher rank).

Noting the assumption does limit treatment effect heterogeneity to some degree – multiple random coefficients are ruled out, for example – we are still able to capture the following relevant features of our setting by combining the monotonicity assumption with a flexible function h^j. First, the impact of the blacklist policy can vary depending on the rank of the municipality, given the model does not restrict the way the policy indicator j and the unobservable u interact. Thus, conditional on observables, locations more prone to deforestation (i.e., those at higher ranks) may require more intense monitoring in order to reduce the amount of deforestation there, while better preserved locations (i.e., those at lower ranks) may be more sensitive to a given level of monitoring and enforcement. Second, the model allows for interactions between the time indicator t and the unobservables. Such interactions are reasonable in our setting because unobservable conversion costs may increase and/or land quality may decrease over time as deforestation in a municipality progresses – if, for example, farmers opt first to deforest in locations with lower conversion costs or higher land quality (as is plausible). Third, because h^j can change flexibly over time, the intervention may have dynamic impacts. For example, it might take time for potential deforesters to update their beliefs about the probability of being caught and fined.[27]

[26] The distinction between *weak* and *strict* monotonicity is innocuous in the CIC model when the dependent variable is continuous, as is the case here. (See the discussion in A&I's footnote 14 on page 439.)

[27] We note that Assumption 1 imposes a 'rank invariance' condition (discussed more fully in Section D of the

Assumption 2 Time Invariance Within Groups: *Conditional on each group G and on observables X, the unobservable U has an identical distribution over time.*

This assumption requires any unobservable differences between Priority and non-Priority municipalities to be stable over time. This is a key condition for the CIC model, playing a role similar to the parallel trends assumption in the standard DID model: in order to construct counterfactual predictions based on the observable distributions, some form of stability over time is necessary. A test analogous to the DID pre-treatment parallel trends test is available in the case of CIC – a placebo test of whether the actual distribution of deforestation shares equals the counterfactual distribution if treatment were assigned falsely in a different year; we present results from this placebo test in the following section.

While restrictive, Assumption 2 is less demanding than it might appear, for three reasons. First, the realizations of U_{mt} may vary over time, and can be serially correlated – due, for instance, to the presence of fixed effects – although they must be drawn from the same distribution.[28] Second, the distribution of unobservables does not have to be the same across treatment and control groups: rather, treatment effects can be heterogeneous across municipalities *and* across groups G. The assumption thus allows for policy interventions targeted at a group with potentially higher average benefits.[29] Third, Assumption 2 (in combination with Assumption 1) is less restrictive than the parallel trends assumption underlying the standard DID estimator: while the CIC model allows group and time effects to differ across individuals with different (observed and unobserved) characteristics, the standard DID model implicitly imposes constant group and time effects. (A comparison of the CIC and the DID models is provided in Section D of the Supplemental Material.) Of note, Assumption 2 allows the distribution of X_{mt} to vary by group and over time. The groups do not need to be balanced in terms of their observable characteristics, nor do we need to reweight and balance them, in order to estimate treatment effects.

Identification. In discussing identification, we adapt key results from A&I to our setting. Let $F_{Y^j_{gt}}$ denote the conditional distribution function of the potential share of deforestation Y^j_{mt} given $G = g$ and $X = x$; we omit the conditioning variables X to simplify the notation. Let the inverse distribution be given by $F^{-1}_{Y^j_{gt}}(q)$ for any quantile $q \in [0, 1]$. (When it is clear from the context, we

Supplemental Material). Rank invariance preserves the intuitive notion that, conditional on observables, a relatively highly deforested location in the data remains a relatively highly deforested location under alternative counterfactual policies – specifically, it preserves the rank u. While all the CIC results extend to a more appealing 'rank similarity' condition by incorporating two types of unobservables, one for each policy regime (allowing a relatively highly deforested municipality to be *more likely but not guaranteed* to remain a relatively highly deforested municipality under alternative policies), we maintain Assumption 1 (and our single-unobservable notation) for ease of exposition.

[28] This implies a rank similarity condition over time, as we discuss in Section D of the Supplement.

[29] Consistent with there being systematic unobservable across-group differences, higher unobservables lead to both higher levels of new deforestation and a higher probability of being placed on the Priority List through past deforestation.

also use the shorthand Y^j_{gmt} to denote the potential outcome variable for a municipality m in group g.) The following notation will be useful: the support of Y_{mt} is denoted by $\mathbb{Y}_t$, the support of Y_{mt} conditional on $G = g$, by $\mathbb{Y}_{gt}$, and the support of Y^j_{mt} given $G = g$, by $\mathbb{Y}^j_{gt}$.

For expositional ease, consider two consecutive periods t and $t+1$, before and after treatment. Athey and Imbens (2006, Theorem 3.1 and Corollary 3.1) show that under Assumptions 1 and 2, the counterfactual distribution of Y^0_{1mt+1} (i.e., the distribution of the treated group $g = 1$ in the absence of the policy intervention, $j = 0$, at $t+1$) is identified on $\mathbb{Y}_{0t+1}$ (i.e., the support of the control group at $t+1$) and is given by

$$F_{Y^0_{1t+1}}(y) = F_{Y_{1t}}\left(F^{-1}_{Y_{0t}}\left(F_{Y_{0t+1}}(y)\right)\right), \tag{2}$$

where $y \in \mathbb{Y}_{0t+1}$. In words, the counterfactual distribution $F_{Y^0_{1t+1}}$ can be calculated based on the distribution of three *observables*: the distribution of deforestation shares for the same group but prior to treatment ($F_{Y_{1t}}$), and the distributions of the share of deforestation for the control group both before and after treatment ($F_{Y_{0t}}$ and $F_{Y_{0t+1}}$). Note that the distribution of Y for the treated group under the treatment at $t+1$ (i.e., after treatment) is trivially identified: $F_{Y^1_{1t+1}} = F_{Y_{1t+1}}$. By comparing the observed $F_{Y_{1t+1}}$ with the counterfactual $F_{Y^0_{1t+1}}$, we can obtain various treatment effects on the treated (average effects and quantile effects, for example).[30]

When the dataset covers one time period before treatment, the model is just-identified. With more than one pre-treatment time period, there is more than one way to identify $F_{Y^0_{1t+1}}$, and so the model becomes overidentified. Of note, $F_{Y^0_{1t+1}}$ is identified only on $\mathbb{Y}_{0t+1}$; outside this support, $F_{Y^0_{1t+1}}$ is not identified. A similar expression to (2) holds for the *untreated* group under the same assumptions (Athey and Imbens, 2006, Theorem 3.2):

$$F_{Y^1_{0t+1}}(y) = F_{Y_{0t}}\left(F^{-1}_{Y_{1t}}\left(F_{Y_{1t+1}}(y)\right)\right), \tag{3}$$

where $y \in \mathbb{Y}_{1t+1}$. Thus equation (3) provides information about treatment effects on the untreated. As before, the counterfactual distribution for the untreated $F_{Y^1_{0t+1}}$ is not identified outside the support of the treated group, Y_{1mt+1}.

Support Conditions and Partial Identification. As just mentioned, the counterfactual distribution of (say) the treated group, $F_{Y^0_{1t+1}}$, is identified only on the support of the control group $\mathbb{Y}_{0t+1}$, and it is not identified outside this support. Consequently, when the support $\mathbb{Y}_{0t+1}$ is not comprehensive enough, i.e., when $\mathbb{Y}_{0t+1} \subset \mathbb{Y}^0_{1t+1}$, then $F_{Y^0_{1t+1}}$ is identified on the subset $\mathbb{Y}_{0t+1}$, but it is not identified at the tails of $\mathbb{Y}^0_{1t+1}$ (assuming all portions of the support are connected). In

[30]See Figure 1, page 442, in A&I for a clear exposition. We provide intuition for (2) in the context of our study in the Supplemental Material, Section D.

this case, we cannot point-identify the average treatment effect on the treated; a similar problem can arise for the untreated group.

Still, it is possible to bound counterfactual distributions (and average treatment effects) in a spirit similar to Manski (2003). By putting all remaining probability mass outside $\mathbb{Y}_{0t+1}$ at the left and right end points of $\mathbb{Y}^0_{1t+1}$, we obtain the lower and upper bounds for $F_{Y^0_{1t+1}}$, denoted by $F^L_{Y^0_{1t+1}}$ and $F^U_{Y^0_{1t+1}}$, respectively. (Once again, the same holds for the untreated group.) To implement this solution, we need prior information relating to the counterfactual support $\mathbb{Y}^0_{1t+1}$. Assumption 3 provides such prior information, as implemented previously in the empirical literature (see, e.g., Ginther, 2000; Lee, 2009).

Assumption 3 Support: *Assume* $\mathbb{Y}^j_{gt} = \mathbb{Y}_{gt}$ *for* $j, g = 0, 1$, and for any t.

Assumption 3 implies that while the policy intervention may affect the distribution of deforestation shares, it does not affect the support of the distribution. Note that, because deforestation shares always lie between zero and one, we could take $\mathbb{Y}^j_{gt} = [0, 1]$ for all j, g, and t, making the assumption innocuous. In practice, the amount of deforestation in a given municipality in a given year is typically small because of high conversion costs, implying that observed deforestation shares are mostly closer to zero than to one in the data. (The maximum deforestation share in a municipality-year in the data is 0.19.) Constructing bounds based on the $\mathbb{Y}^j_{gt} = [0, 1]$ support would then result in unnecessarily wide bounds for treatment effects in practice. We therefore use the observed supports in the data $\mathbb{Y}_{gt}$, making Assumption 3 relevant in practice.[31]

Under Assumption 3, we cannot point-identify the counterfactual distributions of both treated and untreated groups simultaneously when $\mathbb{Y}_{1t+1} \neq \mathbb{Y}_{0t+1}$. For instance, if $\mathbb{Y}_{1t+1} \subset \mathbb{Y}_{0t+1}$, we can point-identify the counterfactual distribution for the treated group $F_{Y^0_{1t+1}}$ (identified on the larger set $\mathbb{Y}_{0t+1}$), but not the control group, $F_{Y^1_{0t+1}}$ (identified only on the smaller set $\mathbb{Y}_{1t+1}$). In this case, we point-identify the average treatment on the treated (ATT), but we can only partially identify the average treatment on the untreated (ATU).

Semiparametric Specification. Although the CIC model can be estimated completely nonparametrically (Athey and Imbens, 2006; Melly and Santangelo, 2015), we adopt a semiparametric specification because of data limitations. The simplest and most parsimonious procedure is to partial-out the covariates X_{mt} and apply the CIC model to the residuals, as A&I suggest.

[31] This procedure leads to worst-case bounds because it does not incorporate possible additional restrictions, such as continuity or smoothness on the counterfactual distributions. In order to minimize the impact of outliers, we follow the literature and trim observations below the 3rd and above the 97th percentiles (Ginther, 2000; Lee, 2009). The empirical results are robust to the trimming – for example, dropping observations below and above the percentiles $[2.5, 97.5]$ and $[3.5, 96.5]$. See Section G of the Supplemental Material. Note that if the policy intervention reduces deforestation shares by also shifting the supports to the left – i.e., by reducing the maximum deforestation possible – our bounds become more conservative (i.e., wider than the true bounds).

To that end, we adopt the logit model, as is common in the empirical land use literature (Stavins, 1999; Pfaff, 1999; Souza-Rodrigues, 2019). This model is appealing in our context. It can be motivated in terms of there being a continuum of farmers who make binary choices (to deforest or not), aggregated up to the municipality level; one can then trace the share of deforestation back to underlying individual decisions – helpful in interpreting the empirical results. From a measurement perspective, the model does not predict negative deforestation (in contrast to a linear model). This is important in our setup because the estimated ex-post optimal list depends crucially on having reasonable predictions for counterfactual deforestation, yet a linear model predicts negative deforestation for over 16 percent of these observations, which may lead to biased ATT estimates and produce misleading results when constructing the counterfactual optimal list. In addition, the logit specification allows for heterogeneous effects of covariates X_{mt} on deforestation in a parsimonious way. This is useful when selecting the ex-post list: if heterogeneous effects were restricted to depend only on unobservables, the ex-post list would only select all municipalities in the group with the highest average impact of treatment. Finally, the logit model has a convenient functional form that makes it easy to partial-out the covariates in order to estimate the CIC model. A fully nonparametric model would require estimating all conditional distribution functions, given X, in equations (2) and (3) nonparametrically, which is not practical.[32]

Assumption 4 Semiparametric Model: *The potential share of newly deforested area, Y_{mt}^{j}, for $j \in \{0,1\}$, in a municipality m at t is given by*

$$Y_{mt}^{j} = \frac{\exp\left[X_{mt}'\beta + V_{mt}^{j}\right]}{1 + \exp\left[X_{mt}'\beta + V_{mt}^{j}\right]}, \tag{4}$$

where (i) $V_{mt}^{j} = v^{j}\left(U_{mt}, t\right)$, with the functions $v^{j}\left(u, t\right)$ satisfying Assumption 1 (i.e., strict monotonicity on u); (ii) the unobservable U_{mt} is independent of X_{mt} given G_m; and (iii) V_{mt}^{j} satisfy the support condition in Assumption 3.

Assumptions 4(i) and 4(iii) apply the CIC model to the residuals V_{mt}. Assumption 4(ii) is the typical extension of the zero correlation assumption from linear to nonlinear models. (Note that this is a semiparametric model because it leaves both the function $v^{j}\left(U_{mt}, t\right)$ and the distribution of U_{mt} unspecified.) By regressing the log odds ratio of the share of deforestation on covariates, we can identify and estimate the coefficients β; then we can back out the residuals V_{mt}, and apply the CIC model to them.[33]

[32] An alternative solution, proposed by Kottelenberg and Lehrer (2017), is to reweight the observations based on propensity scores. Although appealing, this solution is of limited use in the current context because of the lack of common support on propensity scores induced by the selection rule (see Section 4).

[33] Specifically, as A&I note, let I_{mt} be a vector of dummy variables indicating group status (control versus treat-

5.2 Average Treatment Effects

We calculate the average treatment effects as follows: start with the logistic function $\varphi(x, v) = \exp(x'\beta + v) / (1 + \exp(x'\beta + v))$. From (4), the potential share of new deforestation is given by $Y_{mt}^j = \varphi(X_{mt}, V_{mt}^j)$. Expected deforestation under intervention j, D_{mt}^j, conditional on observables (X_{mt} and A_{mt}) and on the group $G = g$, is given by

$$E\left[D_{mt}^j | X_{mt}, A_{mt}, G_m = g\right] = \int \left[\varphi(X_{mt}, v) \times A_{mt}\right] dF_{V_{gt}^j}(v), \tag{5}$$

where the distribution $F_{V_{gt}^j}$ is either observed (using the residuals of the log odds ratio regression) or is obtained from the CIC model (i.e., from either (2) or (3) applied to the residuals V_{mt}). Given (5), average treatment effects are defined in the standard way. When the counterfactual distribution of the residuals is not point-identified, we bound the conditional expectation as

$$\begin{aligned} & \int \left[\varphi(X_{mt}, v) \times A_{mt}\right] dF_{V_{gt}^j}^L(v) \\ \leq \; & E\left[D_{mt}^j | X_{mt}, A_{mt}, G_m = g\right] \\ \leq \; & \int \left[\varphi(X_{mt}, v) \times A_{mt}\right] dF_{V_{gt}^j}^U(v). \end{aligned} \tag{6}$$

Bounds on average treatment effects follow naturally from (6).

In turn, to measure the carbon emissions that result from the deforestation process, we use the equality $E_{mt}^j = D_{mt}^j \times CS_m$, where E_{mt}^j are the potential carbon emissions under policy j, and CS_m is the average difference in carbon stock, comparing forested and deforested areas within municipality m. (For simplicity, we ignore carbon decay and assume all carbon stock is immediately released into the atmosphere once a plot of land is deforested.)

6 Empirical Results

In this section, we present the estimated average treatment effects, along with CIC specification tests. The following covariates are included in the partialing-out regression: lagged rainfall and rainfall squared, lagged temperature, the share of protected areas, lagged price of beef and price of crops, lagged municipality GDP, state dummies, agricultural suitability (as measured by FAO-GAEZ agricultural potentials for soy and for corn), distance to the nearest port, predetermined cropland area (as of 2001), and predetermined number of cattle (as of 2001). The supplement provides further details, including a discussion of possible mechanisms and the relationship between

ment) interacted with time dummies. In the first stage, we estimate the regression $\log\left(\frac{Y_{mt}}{1-Y_{mt}}\right) = X_{mt}'\beta + I_{mt}'\gamma + \nu_{mt}$, then construct the residuals with the group-time effects left in: $\log\left(\frac{Y_{mt}}{1-Y_{mt}}\right) - X_{mt}'\widehat{\beta} = I_{mt}'\widehat{\gamma} + \widehat{\nu}_{mt}$.

our CIC estimates and those from a standard DID model.[34]

Treatment Effects. The estimates, first abstracting from spillovers, are reported in Table 2. In the top panel, we present the estimated average effects on deforestation, as explained in Section 5. The columns in the panel give the estimated ATT, ATU, and ATE, respectively, which are shown separately (in the rows) for 2009 and 2010. The results here are based on the average of the baseline years, 2006 and 2007: results for alternative baselines are shown in Section G of the supplement. The bottom panel of the table reports the estimated total cumulative treatment effects in those two years (summed over all municipalities), in terms of deforestation and carbon emissions. That panel also reports the value of the total emissions avoided, assuming a social cost of carbon of \$20/tCO_2 (Greenstone et al., 2013; Nordhaus, 2014). The numbers in square brackets are lower and upper bound estimates for the partially identified sets, and the numbers in parentheses are 95 percent confidence intervals.[35] We compute the confidence intervals for the parameters of interest (rather than for the identified sets) based on the approach developed by Imbens and Manski (2004); for the point-identified parameters, we use the standard procedure.[36]

Considering the treated group first, all the estimates are statistically significant. The estimated ATT is approximately -24 km^2 for 2009, and -58 km^2 for 2010. Greater estimated impacts in the later year are consistent with farmers updating their beliefs about the stringency of the new policy regime. According to the estimates, the treated group would have deforested a total of 6,916 km^2 in the period 2009-10 in the absence of treatment, which is 71 percent higher than the amount of deforestation observed in the data, implying that the Priority List led to the preservation of 2,886 km^2 of forested area. This translates into avoided emissions of 34.6 million tons of carbon in the same period. In turn, the estimated social benefit of the program in terms of avoided emissions is approximately \$2.54 billion.

[34] In Section C of the supplement, we present the regression that partials out the covariates X_{mt}. We also include the estimated factual and counterfactual distributions of the residuals $F_{V^j_{gt}}$, and provide suggestive evidence regarding possible mechanisms that link the Priority List to deforestation.

[35] Following A&I, the statistical uncertainty is modelled here using a repeated sampling framework over the distribution of the unobservables U_{mt}. This can be thought of as conventional 'super-population-based' (or sampling-based) inference, in contrast with recent work by Abadie et al. (2020) and others considering finite-population design-based uncertainty, where all population units are observed in the data and all the randomness comes only from the treatment assignment. Although we observe all municipalities in the Brazilian Amazon, the selection procedure placing municipalities onto the blacklist means completely random treatment assignments are not credible here. Thus we adopt the usual super-population-based statistical inference, viewing the associated uncertainty as an approximation (perhaps conservative in nature) to finite-population-based statistical uncertainty.

[36] In practice, our confidence intervals are computed based on the standard i.i.d. nonparametric bootstrap, where the i.i.d. resampling occurs in the cross-sectional dimension. Given that the ATU and ATE are both partially identified (as we discuss below), we used 500 bootstrap replications to compute standard errors for the lower- and upper-bound estimators, and plug them into the confidence interval formula (see Imbens and Manski (2004) equations (6) and (7) on page 1850). For the point-identified ATT, we followed a similar procedure using the same bootstrap replications to compute the standard error for the estimator, then plugging it into the standard confidence interval formula, warranted given that A&I have shown that the ATT estimator is asymptotically normally distributed.

Treatment effects on the untreated are not point-identified, but the identified sets are highly informative, the effects being statistically significant. The estimated effects in 2009 range from -3.7 km^2 to -4 km^2, and increase to between -6 km^2 and -6.8 km^2 in 2010. The difference between the estimated ATT and ATU provides evidence of heterogeneous treatment effects, suggesting that the government did select municipalities with potentially higher average impacts. We note that such results could not be obtained using a DID strategy, given that it only identifies effects on the treated. Further, extrapolating results from the treated group to the untreated under the assumption of homogeneous effects would (in light of these heterogeneous estimates) bias up the estimated effects on the untreated.

Table 3 presents the estimated treatment effects, now incorporating potential spillovers. We find that the ATTs are greater than when potential spillover effects are ignored. This is attributable to the fact that lower average reductions in deforestation after treatment now arise in the control group, given that it does not include those municipalities more likely to respond to the policy intervention. The estimates indicate that the Priority List avoided the clearing of 3,050 km^2 of forested area and emissions of 36.5 million tons of carbon during 2009-10. The estimated ATUs are in line with the estimates obtained when we assumed away any spillovers effects.

We denote by 'ATS' the average treatment effect of including a spillover municipality on the Priority List. The estimated ATSs are statistically significant. While they are only partially identified, the estimated sets are again highly informative: the effects are approximately -11.5 km^2 for 2009, and increase to between -16 km^2 and -20 km^2 for 2010. Similar to the other groups, impacts are greater during the second year of the program. The magnitudes of the ATSs fall between the estimated ATT and ATU, providing further evidence of heterogeneous effects.

In terms of total effects, the direct impact of the program avoided 3,050 km^2 of deforestation in Priority municipalities in 2009-10 (as noted), and the indirect impact in the same period discouraged the clearing of 1,102 km^2 of forested area in spillover municipalities – approximately 26% of the total impact.[37] Thus, farmers in untreated municipalities both geographically close to a Priority municipality and that experienced large areas of rainforest being cleared in the past reacted to the policy by *reducing* deforestation. They might have believed that monitoring could increase there, because of the higher risk of being placed on the list in the near future. Leakage or general equilibrium effects would tend to work in the opposite direction, but they are likely to be small. Any leakage would be limited in the short run by the costs of shifting deforestation to other areas. Equilibrium price effects are also likely to be small because the prices of the main products in the Amazon (meat, soybeans, and corn) are determined on international markets, important inputs like

[37] Note that the ATS differs conceptually from an estimate of the reaction of farmers in spillover municipalities to the existence of the policy intervention itself. The latter is given by the indirect effects shown at the bottom of Table 3. Of note, Andrade and Chagas (2016) obtained qualitatively similar results for the (indirect) spillover effects.

fertilizers are mostly imported, and the vast majority of Brazilian beef, soy, and corn production takes place in biomes other than the Amazon. In addition, migration costs would likely limit short-run equilibrium impacts in local labor markets.[38]

We estimate a total of 4,152 km^2 of forested area preserved by the policy. In turn, the program avoided 49 million tons of carbon emissions, with a social benefit of approximately $3.63 billion: these are our preferred summary estimates of the policy impacts. Compared to the combined budget allocated to INPE's and IBAMA's monitoring and enforcement activities for 2009 and 2010, a total of around $117 million, the estimates indicate a benefit-cost ratio of approximately 31 for the program, which is unusually high. Based on that ratio, further investments in monitoring and enforcement would be highly worthwhile.[39]

The estimates just presented call to mind the issue of *possible mechanisms* causing the reductions in deforestation. In Section C of the Supplemental Material, we present evidence relating to these. There, we focus on observables, investigating the effects of treatment status on a proxy for monitoring (the number of alerts given out by INPE), a proxy for enforcement intensity (the evolution of the number of fines issued by IBAMA), the total volume of rural credit concessions, and the share of protected areas. As we show, the suggestive evidence points to an increase in enforcement intensity in Priority municipalities, with policing in untreated units remaining relatively stable. This is consistent with an improvement in state capacity to implement environmental regulations, rather than a simple reallocation of fixed resources following the Priority List's introduction.

Testing. We now discuss the results of three tests applied to the CIC model. Taking these in turn, we first assess whether the actual distribution of deforestation shares equals the counterfactual distribution when imposing the policy intervention (falsely) one year early, in 2007 – this serves as a placebo test, the CIC analog to the DID pre-treatment parallel trend test. Second, we test whether the Priority List affects the entire distribution of outcomes, similar to the placebo test but using the correct timing of the intervention. Third, we test whether the counterfactual distribution is everywhere below the actual distribution, as would be the case if the absence of treatment resulted in more deforestation everywhere – a stochastic dominance test.

Table 4 presents the results of these tests. We apply each test to both the log odds ratio of

[38]Indirect equilibrium price effects could lead to more deforestation in control (non-spillover untreated) municipalities as well, tending to bias the magnitudes of our estimated treatment effects upward. Our sense is that any bias would be small for the same reasons listed in the main text.

[39]The sum of 2009 and 2010 budgets for INPE's Amazon satellite monitoring program is approximately $3.95 million. IBAMA's two monitoring and enforcement programs related to deforestation are the "Prevention and Combat of Deforestation, Fires and Forest Fires," and the "Environmental Policy Management" programs. In 2009-10, their combined budget was around $113 million, approximately 10% of IBAMA's total budget. This combined budget provides an overestimate of IBAMA's monitoring costs related to the Priority List, as it includes expenditures on other monitoring activities and in areas other than the Amazon biome. Source: `https://www.siop.gov.br/modulo/login/index.html#/`.

deforestation shares not conditioning on covariates and to the residuals (V_{mt}) after partialling-out the covariates, as explained in Section 5. In all cases, the p-values correspond to both the Kolmogorov-Smirnov and the Cramer-von Mises statistics. For the two outcomes and associated test statistics, we fail to reject the null of 'no impact' when the policy intervention is wrongly imposed in 2007 – i.e., the placebo test passes. In contrast, we reject the null of no impact when the policy intervention is set correctly in 2008, and we find strong evidence in favor of stochastic dominance, as one might expect given the estimated treatment effects discussed above.

7 Optimal Policy Targeting

In this section, we develop a counterfactual framework for targeting regulations optimally based on the estimated treatment effects in the previous section. Then we present and discuss the results from various counterfactual targeting exercises.

7.1 Policy Targeting Framework

Suppose a policy maker wishes to assign municipalities to the Priority List in order to minimize total deforestation (or total emissions), and that she has information about the conditional average treatment effects estimated above, along with the covariates we have used. For expositional ease, denote the expected deforestation under intervention j by $\mathcal{D}^j_{mt} \equiv E[D^j_{mt}|X_{mt}, A_{mt}, G_m]$, for $j \in \{0,1\}$. Also, denote the counterfactual assignment rule in period t by $\phi_t = (\phi_{1t}, ...\phi_{Mt})$, which maps municipalities $m = 1, ..., M$ to the Priority List treatment; this can be either deterministic $\phi_{mt} \in \{0,1\}$ or probabilistic $\phi_{mt} \in [0,1]$. For a given time period (and ignoring spillovers for now), the policy maker solves the problem

$$\min_{\phi_t \in [0,1]^M} \sum_{m=1}^{M} \left[\phi_{mt}\, \mathcal{D}^1_{mt} + (1-\phi_{mt})\ \mathcal{D}^0_{mt}\right]. \tag{7}$$

The minimum deforestation is achieved (trivially) by a singleton rule that allocates m to the treatment when $\mathcal{D}^1_{mt} \leq \mathcal{D}^0_{mt}$. When the equality holds, any random allocation is optimal.

The minimization problem in (7) abstracts from two important considerations. The first involves resource constraints. The original Priority List sought to direct limited resources where they were expected to have the greatest impact. Given that information about the actual resources allocated to policing is difficult to obtain, we incorporate limited monitoring resource constraints into the policy maker's minimization problem in two alternative ways. One constraint limits the total area $\overline{S}$ that can be monitored under the Priority List (given the plausible notion that the costs of monitoring and punishing illegal deforestation increase with the *total area* covered by the policy).

We write this as:

$$\sum_{m=1}^{M} s_m \times \phi_{mt} \leq \overline{S}, \tag{8}$$

where s_m is the area of municipality m. The alternative constraint applies to the *total number* of municipalities $\overline{M}$ that can be placed on the list:

$$\sum_{m=1}^{M} \phi_{mt} \leq \overline{M}. \tag{9}$$

This constraint is reasonable when monitoring costs are primarily a function of the *number* of districts that inspectors must visit.[40] Once the constraints are taken into account, the estimated magnitudes of the treatment effects for *all* municipalities influence the assignment of *each* municipality m to the list.

The second consideration relates to partial identification: when the support conditions are violated, we can only partially identify counterfactual expected deforestation. This means that an ex-post policy evaluation must be analyzed as a treatment choice problem under ambiguity (Manski, 2005). Here we consider the minimax criterion, assuming the policy maker chooses the blacklist in order to minimize total deforestation in the worst-case scenario.

Formally, define the vector $\mathcal{D}_{mt} = (\mathcal{D}^0_{mt}, \mathcal{D}^1_{mt})$. The set of admissible values for $\mathcal{D}_{mt}$ is $\mathbb{D}_{mt} = \mathbb{D}^0_{mt} \times \mathbb{D}^1_{mt}$, with $\mathbb{D}^j_{mt} = [\underline{\mathcal{D}}^j_{mt}, \overline{\mathcal{D}}^j_{mt}]$ for $j \in \{0, 1\}$, where $\underline{\mathcal{D}}^j_{mt}$ and $\overline{\mathcal{D}}^j_{mt}$ are the lower and upper bounds given by the terms on the left-hand side and the right-hand side of inequality (6), respectively. Next, define the tuple $\mathcal{D}_t \equiv (\mathcal{D}_{1t}, ..., \mathcal{D}_{Mt})$, and the product set $\mathbb{D}_t \equiv \prod_{m=1,...,M} \mathbb{D}_{mt}$. The policy maker's problem under the minimax criterion is

$$\min_{\phi_t \in [0,1]^M} \max_{\mathcal{D}_t \in \mathbb{D}_t} \sum_{m=1}^{M} \left[\phi_{mt} \, \mathcal{D}^1_{mt} + (1 - \phi_{mt}) \; \mathcal{D}^0_{mt}\right], \tag{10}$$

subject either to the 'total area' constraint (8), or to the 'total number of municipalities' constraint (9). The minimization problem (10) subject to either constraint simplifies to a linear programming problem that is straightforward to solve numerically (see Section F of the supplement). In the

[40] Ideally, we would have precise information about the expected monitoring costs for each municipality m in each time period t, both without and with the treatment. Then we could replace the constraints (8) and (9) with the restriction

$$\sum_{m=1}^{M} \left[\phi_{mt} E\left[MC^1_{mt} | X_{mt}, G_m\right] + (1 - \phi_{mt}) \, E\left[MC^0_{mt} | X_{mt}, G_m\right]\right] \leq K_t,$$

where MC^j_{mt} are the expected monitoring costs, and K_t is the government's budget constraint. (We note that our framework can accommodate other objective functions, for example augmenting monitoring costs with the social costs of deforesting land.) Such an approach is not feasible here, however. While we know IBAMA's and INPE's total budgets, in practice we do not have information about the true budget constraint for the Priority List K_t; nor do we know how much of the total is allocated to monitoring nor how monitoring costs are distributed across municipalities.

empirical exercise, when using constraint (8), we set $\overline{S}$ equal to the total area occupied by the municipalities that were effectively put on the list in 2008 (i.e., the treated group). Similarly, when using constraint (9), we set $\overline{M} = 35$, which is the number of municipalities in the treated group. Doing so allows us to assess how close the observed Priority List was to the ex-post optimal assignment. We also investigate how the results change as we relax the constraints.

In the presence of spillover effects, the objective function is non-linear and non-differentiable in ϕ, so that we cannot solve the minimax problem using standard methods. Instead, to find the global minimum, we use a stochastic search algorithm – more precisely, a genetic algorithm that allows for integer optimization in high-dimensional constrained minimization problems. (See Section F of the supplement for a detailed description of the computation of the ex-post optimal list in this case.) Of note, we do not select a list that changes over time as this complicates the problem substantially, given the combinatorics involved. We therefore include the total of expected deforestation during 2009 and 2010 for each municipality in the objective function (10).[41]

7.2 Policy Targeting Results

We now present results absent potential spillovers, then show how targeted policies are affected once spillover effects are taken into account.

'No Spillovers' Case. The top panel of Table 5 compares the original Priority List with the ex-post optimal list obtained by solving the relevant constrained minimizations, absent spillovers: the left panel considers the total area $\overline{S}$ that can be monitored as the constraint (see equation (8)), while the right panel fixes the number of municipalities $\overline{M}$ (equation (9)).

Overall, the proportion of municipalities that appear on both lists is high: 83 percent when the constraint involves the total area, and 93.5 percent when the constraint is a maximum number of municipalities. When the policy maker is constrained to 'police' a pre-specified overall area, she can reduce deforestation in the worst-case scenario by replacing ten large municipalities on the Priority List with 73 municipalities that are smaller in size but that would help reduce total deforestation. In contrast, when the restriction applies to the number of municipalities, the policy maker would do better – as expected – by replacing small municipalities (comprising almost half of the Priority List) by municipalities that are larger in size (the total area covered by this list being 42 percent larger than the original).

Figure 4 presents the geographic distribution of municipalities on the various lists. Recall that the top left panel presents the actual Priority List and the top right panel shows the Priority

[41] The amount of expected deforestation for each m in (10) is the sum of the expected deforestation in 2009 (calculated based on equation (5)), and the expected deforestation in the following year, taking into account the counterfactual remaining forested area from the previous year, as explained in the supplement's Section E – see equation (E3).

List together with protected areas (composed of conservation units and indigenous reserves). The bottom left panel then shows the optimal list when the constraint is the total area covered, and the bottom right panel, the counterfactual list when the constraint is the number of municipalities: the bottom two panels also display protected areas, for reference.

Two interesting patterns emerge – features that were not imposed during the course of the estimation strategy. First, the overlap between the protected areas and the area-constrained counterfactual list is much smaller than the overlap between the protected areas and the original Priority List; indeed, the former comprises an area approximately half the latter. This suggests that these two policies could be combined together in an even more productive way than in the actual implementation by the Brazilian government. Second, more specifically, the geographic distribution of the area-constrained counterfactual list traces out a protective shield close to the deforestation frontier; that frontier, the "Arc of Deforestation," is located along the southeastern edge of the Amazon Biome. In the current context, the Priority List may therefore serve to work alongside the protected areas in preventing the deforestation process from continuing into more pristine regions.

To shed some light on which observable factors might be more important in determining whether a municipality is placed on the optimal list or not, we estimate simple reduced-form regressions, regressing the optimal list indicator on the covariates X, the 'criteria' variables Z, and the Priority status indicator G. Based on that exercise, the two most important factors predicting the optimal list status are the share of protected areas (consistent with our analysis of the geographic distribution of the optimal list – see Figure 4), and Priority status itself, which is unsurprising, given the overlap between the two lists presented in Table 5.[42]

Next, we seek to quantify the *consequences* of optimally targeted policies. We do so by comparing the maximum possible expected deforestation and the carbon emissions achieved under the optimal list with the corresponding outcomes under the Priority List, along with another benchmark: a list composed of municipalities that are selected randomly.[43] The top panel of Table 6 presents the results. Compared to the area-constrained optimal list, the Priority List results in around 6 percent more deforestation and 5 percent higher carbon emissions in 2009–2010. The estimated avoided emissions translate into an additional social value of the optimal list (i.e., in excess of the value of the Priority List) of at least \$622 million for that two-year span alone.

We find that the ex-post optimal list fixing the number of municipalities performs slightly better than the area-constrained optimal list. But since it covers a much larger area, monitoring costs are likely to be significantly higher in the former case. In comparison, randomly selecting

[42] As an aside, we note that these regressions are meant to be suggestive only: they ignore the fact that the assignment of a municipality onto the optimal list depends on the characteristics of *all* municipalities in the presence of capacity constraints. (Results are available upon request.) For an evaluation of the impacts of protected areas on deforestation, see Pfaff et al. (2015).

[43] We simulated 1000 random lists with $\overline{M} = 35$ and computed the average resulting counterfactual deforestation and emissions.

35 municipalities onto the list would result, on average, in 22 percent more deforestation and 26 percent higher emissions than the number-constrained optimal list.

To investigate further how the two main policies in the Brazilian Amazon – the Priority List and protected areas (PAs) – interact with each other, we calculate the optimal blacklist when fixing the share of PAs in every municipality to be the same. Focusing on the area-constrained lists, we find that when setting the share of PAs in all municipalities to equal the sample average, the carbon emissions of a restricted 'fixed-PA' optimal list are around 66% higher than the optimal list, providing further evidence underlining the importance of coordinating both policies to reduce deforestation. This is also substantially higher than the emissions under the Priority List, suggesting that the policy makers did incorporate (to some extent, even if implicitly) the geographic distribution of protected areas when implementing the Priority List. When we increase the share of PAs in all municipalities to the sample average plus one sample standard deviation, the emissions of the 'fixed-PA' optimal list are just 14% greater than the optimal list – a reasonable result given that a much larger area would have been protected, leaving therefore less room for the optimal list to reduce emissions further.

Next, we are interested to see how much the minimax solution for carbon emissions is affected by relaxing the constraints we have been imposing. Figure 5 presents the level of emissions at the optimum for the total area constraint. The vertical line shows the maximum $\overline{S}$ that corresponds to the area covered by on the Priority List. The horizontal line corresponds to the amount of carbon emissions estimated directly from the data for 2009-10.

As the figure makes clear, the minimax carbon emissions decrease rapidly when a small area is covered by the optimal list and level off for large $\overline{S}$ eventually, indicating that the benefits of including additional municipalities on the list decrease with $\overline{S}$. Because monitoring costs are likely to increase with $\overline{S}$, concentrating efforts on a strategically-selected subregion of the Amazon emerges as a suitable policy. Furthermore, the minimum area needed for the optimal list to generate the same amount of emissions as the original Priority List is approximately 550 thousand km^2, which is 72 percent of the area covered by the original list (763 thousand km^2); this corresponds to the point in the figure where the minimax carbon emissions curve crosses the horizontal line. This finding draws attention to the substantial monitoring cost savings that are available, holding the level of observed emissions fixed.[44]

'Spillover' Case. Next we discuss the ex-post optimal lists when spillovers are incorporated into both the estimation procedure and the minimization problem. Here we draw attention to how the results differ in the 'spillover' case.

[44]The same type of reasoning applies when we change the number of municipalities allowed on the optimal list, $\overline{M}$, with the minimum number of municipalities generating the same amount of observed emissions as the Priority List being 21, or around 60 percent of the original list.

First, the bottom panel of Table 5 compares the Priority List with the ex-post optimal lists based on the two different constraints we consider, but now allowing for spillovers. Under the total area constraint, the optimal list replaces a greater number of large municipalities with small municipalities, as compared to the optimal list with no spillovers. This leads to a smaller overlap in the municipalities that appear in both the optimal list and original Priority List – now 72 percent. Intuitively, such an assignment takes advantage of the fact that a larger number of small municipalities treated can have wider impacts because of spillover effects.

Figure 6 presents the geographic distribution of the resulting optimal lists. When compared to the no-spillover case in Figure 4, the area-constrained optimal list is now more geographically dispersed, with fewer municipalities being contiguous. Such a pattern is also attributable to the operation of spillover effects: placing all targeted municipalities together does not exploit the potential reduction in deforestation in adjacent locations that arises when spillovers are operating. (A similar pattern is observed for the number-constrained optimal list.)

The bottom panel of Table 6 then compares the levels of deforestation and emissions associated with the alternative lists. Because the optimal lists now take advantage of potential spillovers, they can achieve lower levels of forest loss in the worst-case scenario. The Priority List now results in around 13 percent more deforestation and 12 percent higher carbon emissions in 2009-10 than the area-constrained optimal list. This places a lower bound on the additional social value of the optimal list of approximately \$1.29 billion (approximately 36% of the social value of the original blacklist), leading to a total value of at least \$4.92 billion. From a suggestive benefit-cost perspective, if we assume that the same combined budgets for INPE and IBAMA observed in 2009-10 were allocated to implement the counterfactual list, the corresponding benefit-cost ratio for the optimal list would be at least 42 – a substantial increase when compared to the ratio under the Priority List. (Results are similar when we consider the optimal list constrained by the number of municipalities that can be included.) In turn, randomly selected municipalities now result in 30 percent more emissions than the number-constrained optimal list.

More generally, by relaxing the constraints of the minimization problem, we find that the minimum area the optimal list would need to cover in order to generate the same amount of carbon emissions as the Priority List is approximately 250 thousand km^2, only around 33 percent of the area covered by the original list. This points to even greater monitoring cost savings that become available once spillovers are taken into account. (Similarly, the minimum number of municipalities generating the same amount of observed emissions as the Priority List is 19 – just over half the original blacklist.)

It is worth rehearsing several reasons why we find a mismatch between the Priority List and our optimal lists. First, the authorities used a simple threshold rule to select municipalities, while we minimized total deforestation explicitly. The explicit solution does not necessarily result in a simple threshold rule – that is, we made better use of the information about the geographic distribution of the protected areas than the government's threshold rule could. Second, the selection rule may have considered (even if implicitly) effective monitoring costs, while the lack of information about these constraints led us to proxy for these costs in the form of restricting the total area covered or the number of municipalities included on the list. Third, regulators did not seem to have incorporated potential benefits from spillovers in the design of the policy, while our optimal blacklist takes explicit account of spillover effects. Fourth, we made use of ex-post treatment effect estimates that were not available to the government at the time of their decision. Fifth, partial identification of treatment effects, combined with the minimax decision rule, results in a list that minimizes the (expected) *maximum* deforestation (i.e., that minimizes the worst-case scenario), which is different from a list that minimizes expected deforestation (possible under point-identification of treatment effects). Thus, even if the government tried to minimize expected deforestation, that list would not necessarily coincide with our minimax list.

8 Conclusion

In this paper, we have developed a new approach for assessing the efficacy of targeted, blacklist-type policies for slowing deforestation, a primary contributor to global carbon emissions and a source of considerable concern worldwide. Focusing on the Priority List introduced by the federal government in the Brazilian Amazon in 2008, we first showed that the policy had a substantial causal impact, utilizing the flexible changes-in-changes approach of Athey and Imbens (2006): deforestation was cut by 43 percent in municipalities placed on the list (relative to the case in which no policy was introduced), and also generated non-trivial spillover effects in the form of lower deforestation elsewhere.

We then used the treatment effect estimates in a counterfactual policy framework that allowed us to compute an ex-post optimal list, reflecting realistic resource constraints faced by the regulatory agency in its ability to monitor behavior and enforce environmental protections. The framework also accommodated the possibility that the treatment effect estimates that serve as inputs to the policy calculations may be only partially identified.

Comparing ex-post optimal lists with the actual Priority List, we showed that optimal targeting can generate significant additional gains. Carbon emissions would be at least 10 percent lower than under the Priority List, resulting in savings of at least \$1.29 billion (more than a third of the total value of the Priority List). The estimated benefit-cost ratio under the optimal list is over a third higher than under the Priority List, itself already very high.

From a regulation perspective, our approach provides a means to quantify the gains to the environment from optimally targeted policies aimed at countering tropical deforestation, based on credible econometric estimates. More generally, our counterfactual approach using ex-post treatment effects is applicable in a variety of other settings where targeted regulations have been introduced. In such contexts, the approach can help policy makers assess which policy configurations (accounting for resource and information constraints) are likely to have most environmental impact.

References

Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2020). Sampling-based versus design-based uncertainty in regression analysis. *Econometrica 88*(1), 265–296.

Alix-Garcia, J. M., K. R. E. Sims, and P. Yañez Pagans (2015). Only one tree from each seed? Environmental effectiveness and poverty alleviation in Mexico's payments for ecosystem services program. *American Economic Journal: Economic Policy 7*(4), 1–40.

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

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

Andrade, L. and A. L. S. Chagas (2016). Spillover effects of blacklisting policy in the Brazilian Amazon. Working Paper 2016-32, Department of Economics, University of São Paulo.

Arima, E. Y., P. Barreto, E. Araújo, and B. Soares-Filho (2014). Public policies can reduce tropical deforestation: Lessons and challenges from Brazil. *Land Use Policy 41*, 465–473.

Assunção, J., C. Gandour, and R. Rocha. (2017). DETERing deforestation in the Amazon: Environmental monitoring and law enforcement. CPI Working Paper, Climate Policy Initiative.

Assunção, J., C. Gandour, and E. Souza-Rodrigues (2019). The forest awakens: Amazon regeneration and policy spillovers. CPI Working Paper, Climate Policy Initiative.

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

Athey, S. and G. W. Imbens (2006). Identification and inference in nonlinear difference-in-differences models. *Econometrica 74(2)*, 431–97.

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

Bhattacharya, D. and P. Dupas (2012). Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics 167*(1), 168–196.

Bonan, G. B. (2008). Forests and climate change: Forcings, feedbacks, and the climate benefits of forests. *Science 320*(5882), 1444–1449.

Borner, J., S. Wunder, S. Wertz-Kanounnikoff, G. Hyman, and N. Nascimento (2014). Forest law enforcement in the Brazilian Amazon: Costs and income effects. *Global Environmental Change 29*, 294–305.

Burgess, R., F. J. M. Costa, and B. A. Olken (2019). The Brazilian Amazon's double reversal of fortune. Working Paper, London School of Economics.

Burgess, R., M. Hansen, B. A. Olken, P. Potapov, and S. Sieber (2012). The political economy of deforestation in the tropics. *The Quarterly Journal of Economics 127(4)*, 1707–1754.

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

Cassidy, R. and C. F. Manski (2019). Tuberculosis diagnosis and treatment under uncertainty. *Proceedings of the National Academy of Sciences 116*(46), 22990–22997.

Cisneros, E., J. Hargrave, and K. Kis-Katos (2013). Unintended consequences of anti-corruption strategies: Public fiscal audits and deforestation in the Brazilian Amazon. Working Paper, BIOECON Conference.

Cisneros, E., S. L. Zhou, and J. Börner (2015). Naming and shaming for conservation: Evidence from the Brazilian Amazon. *PloS One 10*(9), e0136402.

Dehejia, R. H. (2005). Program evaluation as a decision problem. *Journal of Econometrics 125*(1), 141–173.

Foley, J. A., R. DeFries, G. P. Asner, C. Barford, G. Bonan, S. R. Carpenter, F. S. Chapin, M. T. Coe, G. C. Daily, H. K. Gibbs, J. H. Helkowski, T. Holloway, E. A. Howard, C. J. Kucharik, C. Monfreda, J. A. Patz, I. C. Prentice, N. Ramankutty, and P. K. Snyder (2005). Global consequences of land use. *Science 309*(5734), 570–574.

Friesen, L. (2003). Targeting enforcement to improve compliance with environmental regulations. *Journal of Environmental Economics and Management 46*(1), 72 – 85.

Ginther, D. K. (2000). Alternative estimates of the effect of schooling on earnings. *The Review of Economics and Statistics 82(1)*, 103–116.

Goodman-Bacon, A. (2018). Difference-in-differences with variation in treatment timing. NBER Working Paper 25018, National Bureau of Economic Research.

Gray, W. B. and J. P. Shimshack (2011). The effectiveness of environmental monitoring and enforcement: A review of the empirical evidence. *Review of Environmental Economics and Policy 5*, 3–24.

Greenstone, M. and R. Hanna (2014). Environmental regulations, air and water pollution, and infant mortality in India. *American Economic Review 104(10)*, 3038–3072.

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

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

Hanna, R. and B. A. Olken (2018). Universal basic incomes versus targeted transfers: Anti-poverty programs in developing countries. *Journal of Economic Perspectives 32*(4), 201–26.

Harding, T., J. Herzberg, K. Kuralbayeva, T. Harding, J. Herzberg, and K. Kuralbayeva (2021). Commodity prices and robust environmental regulation: Evidence from deforestation in Brazil. *Journal of Environmental Economics and Management 108*, 102452.

Harrington, W. (1988). Enforcement leverage when penalties are restricted. *Journal of Public Economics 37*(1), 29 – 53.

Hirano, K. and J. R. Porter (2020). Asymptotic analysis of statistical decision rules in econometrics. Technical report, Pennsylvania State University.

Imbens, G. W. and C. F. Manski (2004). Confidence intervals for partially identified parameters. *Econometrica 72*, 1845–1857.

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

Jack, B. K. and S. Jayachandran (2018). Self-selection into payments for ecosystem services programs. *Proceedings of the National Academy of Sciences*.

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

Kitagawa, T. and A. Tetenov (2018). Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica 86*(2), 591–616.

Koch, N., E. K. H. J. Ermgassen, J. Wehkamp, F. Oliveira, and G. Schwerhoff (2019). Agricultural productivity and forest conservation: Evidence from the Brazilian Amazon. *American Journal of Agricultural Economics 101*(3), 919–940.

Kottelenberg, M. J. and S. F. Lehrer (2017). Targeted or universal coverage? Assessing heterogeneity in the effects of universal child care. *Journal of Labor Economics 35*(3), 609–653.

Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies 76(3)*, 1071–1102.

Maia, H., J. Hargrave, J. J. Gómez, and M. Röper (2011). Avaliação do plano de ação para prevenção e controle do desmatamento na Amazonia legal: PPCDAm 2007–2010.

Manski, C. (2000). Identification problems and decisions under ambiguity: Empirical analysis of treatment response and normative analysis of treatment choice. *Journal of Econometrics 95*(2), 415–442.

Manski, C. (2003). *Partial Identification of Probability Distributions*. New York: Springer-Verlag.

Manski, C. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica 72(4)*, 1221–1246.

Manski, C. (2005). *Social Choice with Partial Knowledge of Treatment Response*. Princeton University Press.

Manski, C. F. (2006). Search profiling with partial knowledge of deterrence. *The Economic Journal 116*(515), F385–F401.

Manski, C. F. (2010). Vaccination with partial knowledge of external effectiveness. *Proceedings of the National Academy of Sciences 107*(9), 3953–3960.

Manski, C. F., A. H. Sanstad, and S. J. DeCanio (2021). Addressing partial identification in climate modeling and policy analysis. *Proceedings of the National Academy of Sciences 118*(15).

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

Matsuura, K. and C. Willmott (2012). Terrestrial precipitation: 1900–2010 gridded monthly time series (1900 - 2010). *University of Delaware. http://climate.geog.udel.edu/climate/*.

Melly, B. and G. Santangelo (2015). The changes-in-changes model with covariates. Working Paper, University of Bern.

MMA, M. d. M. A. (2009). Plano de ação para prevenção e controle do desmatamento na Amazonia Legal (PPCDAm). 2a fase (2009-2011), rumo ao desmatamento ilegal zero.

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

Pailler, S. (2018). Re-election incentives and deforestation cycles in the Brazilian Amazon. *Journal of Environmental Economics and Management 88*, 345–365.

Pfaff, A. (1999). What drives deforestation in the Brazilian Amazon? Evidence from satellite and socioeconomic data. *Journal of Environmental Economics and Management 37*(1), 26–43.

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

Rajão, R., B. Soares-Filho, F. Nunes, J. Börner, L. Machado, D. Assis, A. Oliveira, L. Pinto, V. Ribeiro, L. Rausch, H. Gibbs, and D. Figueira (2020). The rotten apples of Brazil's agribusiness. *Science 369*(6501), 246–248.

Simonet, G., J. Subervie, D. Ezzine-de Blas, M. Cromberg, and A. E. Duchelle (2019). Effectiveness of a REDD+ project in reducing deforestation in the Brazilian Amazon. *American Journal of Agricultural Economics 101*(101), 211–229.

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

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

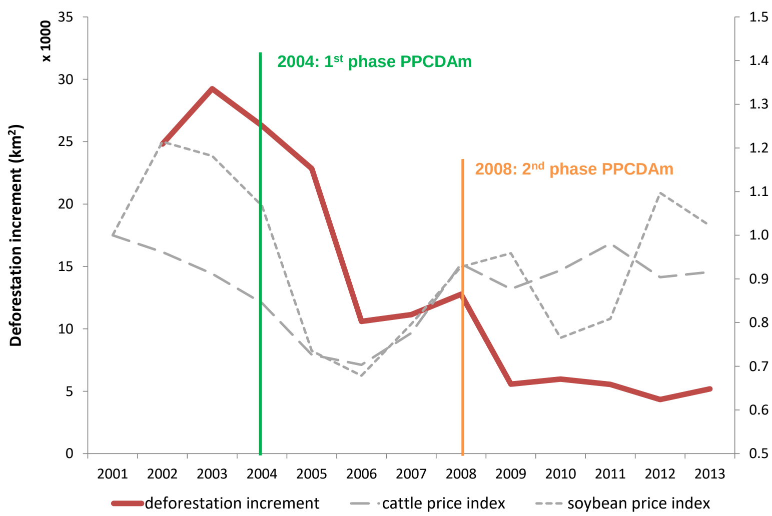

Figure 1: Incremental Deforestation, Key Policy Changes, and Agricultural Commodity Prices, by Year

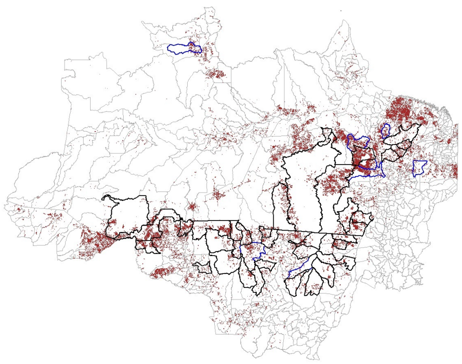

(a) Incremental Deforestation, 2006

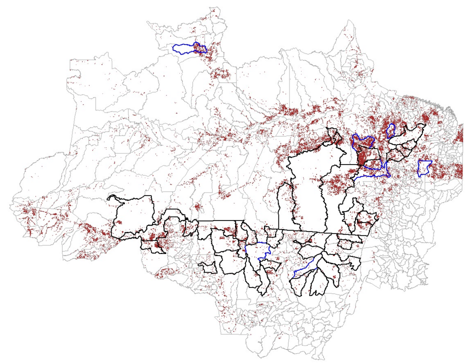

(b) Incremental Deforestation, 2010

Figure 2: Map of Deforestation in 2006 and 2010 (with Priority Municipalities overlaid)

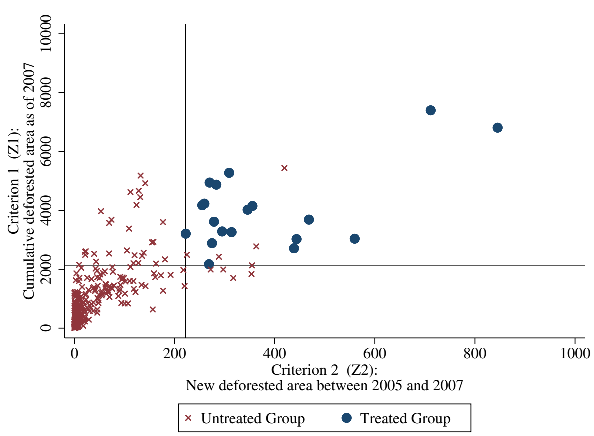

(a) Combinations of Z^1 and Z^2, given $Z^3 = 0$

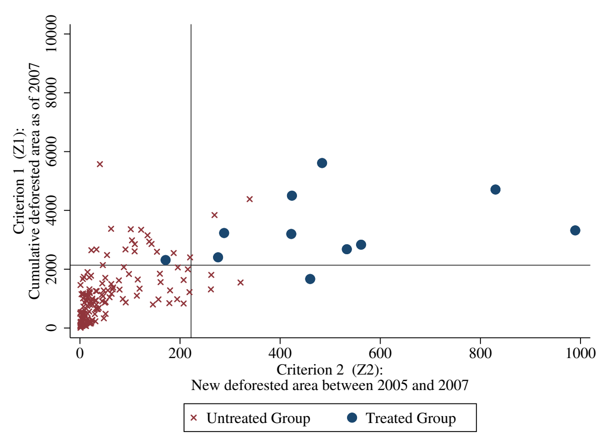

(b) Combinations of Z^1 and Z^2, given $Z^3 = 1$

Figure 3: Selection onto Priority List in 2008 – Combinations of Criteria Variables (Z_{mt-1})

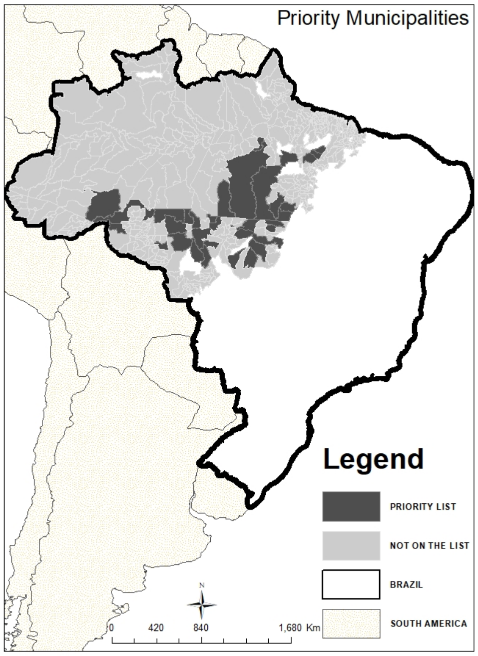

(a) Priority List

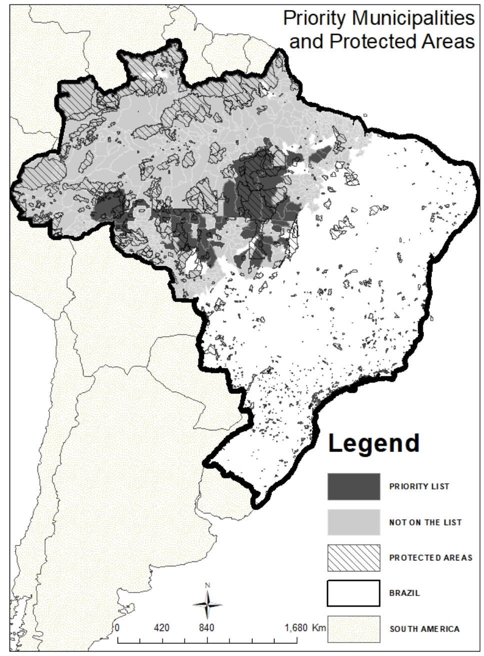

(b) Priority List and Protected Areas

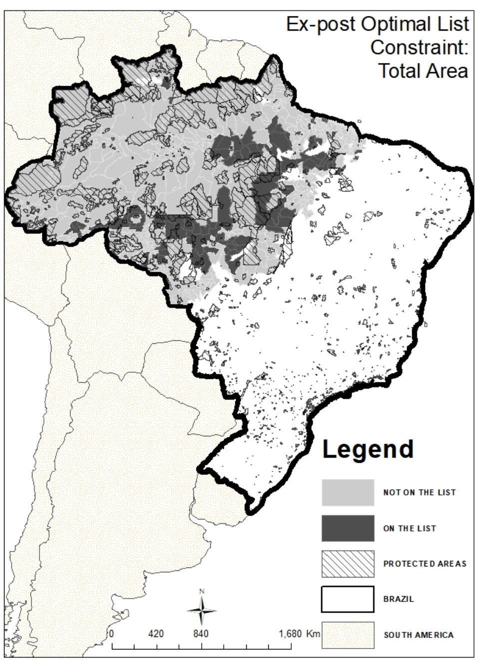

(c) Optimal List based on Total Area

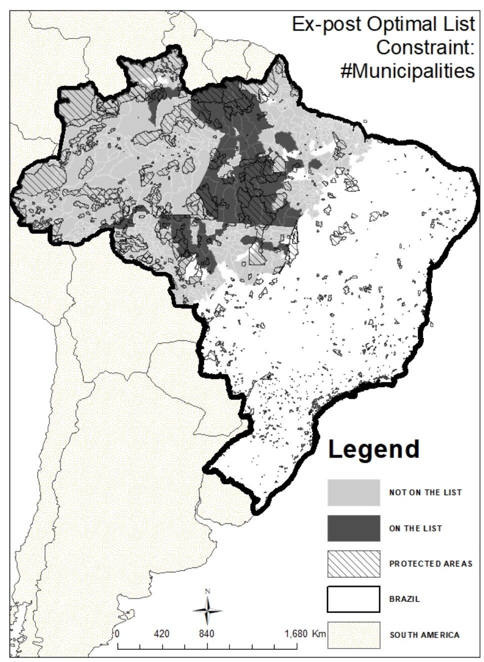

(d) Optimal List based on Number of Municipalities

Figure 4: Location of Priority List, Protected Areas, and Ex-post Optimal Lists without Spillovers

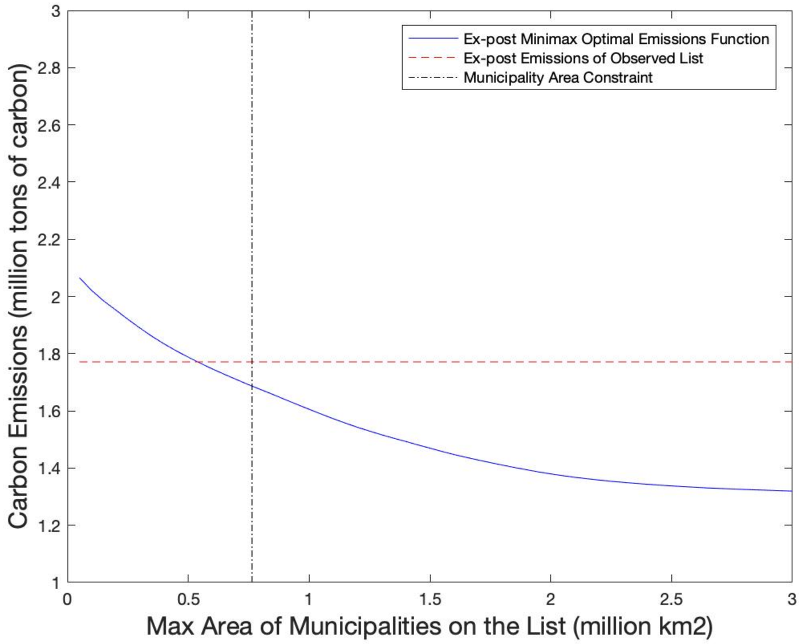

Figure 5: Ex-Post Minimax Carbon Emissions, Varying the 'Total Area' Constraint

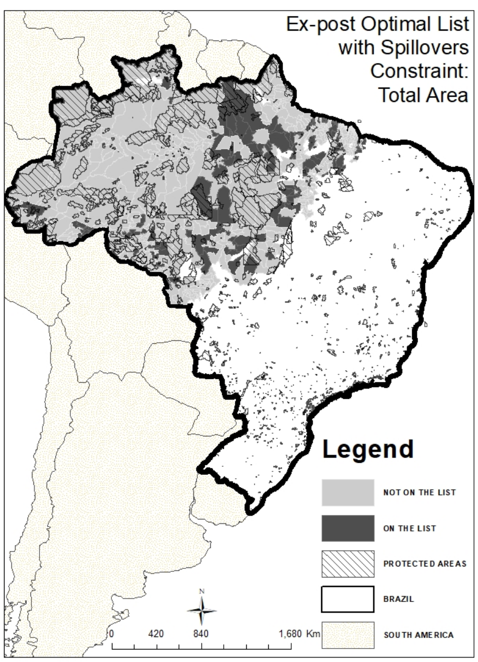

(a) Spillover Optimal List based on Total Area

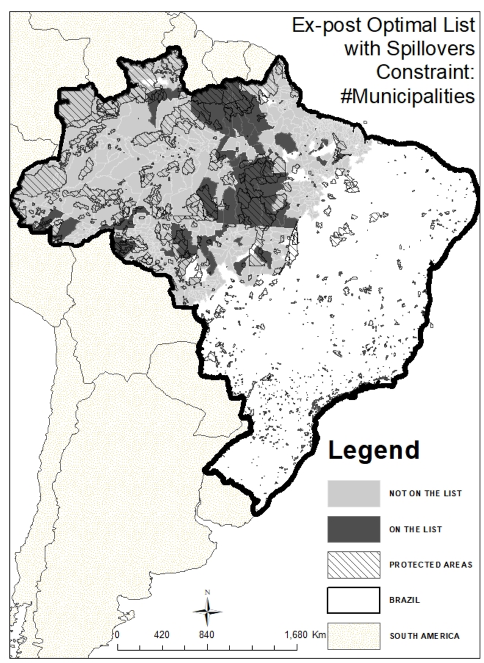

(b) Spillover Optimal List based on Number of Municipalities

Figure 6: Location of Ex-post Optimal Lists with Spillover Effects

Table 1: Summary Statistics for 2007 Cross-Section by Group

	Total Sample ($N = 490$)		Treated Group ($N = 35$)		Untreated Group ($N = 455$)		Spillover Group ($N = 34$)		Control Group ($N = 421$)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Land Use (km^2)										
Deforested Area	21	60	148	169	12	20	48	34	8.60	15
Cumulative Deforested Area	1,270	1,436	4,413	2,437	1,028	978	2,513	933	907	879
Forested Area	6,499	15,507	15,990	27,549	5,769	13,953	5,909	10,037	5,757	14,231
Municipal Area	8,726	16,716	21,815	29,409	7,719	14,899	10,053	10,910	7,530	15,170
Deforested Share (%)	1.21	1.91	1.72	1.41	1.17	1.94	1.96	2.09	1.11	1.92
Policy Measures										
Number of Alerts	57	213	509	616	22	60	129	133	13	39
Fines Issued	9	19	40	44	7	14	26	27	5	10
Share of Protected Area (%)	28	33	22	22	28	34	18	21	29	34
Agriculture and Ranching										
GDP (million Reais)	179	1,013	180	399	179	1,046	186	328	178	1,083
Agricultural GDP (million Reais)	19	24	39	22	18	24	38	39	16	21
Cattle (thousands)	105	148	363	290	85	108	233	118	73	98
Crop Area	109	455	289	563	95	443	256	386	82	445
Total Rural Credit (million Reais)	7	13	19	16	6	12	15	15	5	12
Other Variables										
Rainfall (mm)	2,206	613	1,948	195	2,227	633	2,031	287	2,242	650
Temperature ($°C$)	26	1	26	1	26	1	26	1.55	26	1
FAO-GAEZ Soy	33.7	4.32	35.7	2.07	33.5	4.41	35.7	2.40	33.4	4.49
FAO-GAEZ Corn	54.7	8.65	52.6	1.22	54.9	8.95	53.8	3.66	55	9.25
Distance to Port (km)	909	723	1,174	462	888	735	1,255	563	859	741
Carbon Stock in Forested Areas (tC/ha)	189	68	212	35	187	69	204	45	186	71
Carbon Stock in Deforested Areas (tC/ha)	117	48	101	23	118	49	99	26	119	50

Notes: This table reports municipality-level means and standard deviations (SD) for the variables used in the empirical analysis for the year 2007. An observation is a municipality in the Brazilian Amazon. The Untreated Group combines the Spillover and Control Groups. The Spillover Group consists of the untreated municipalities that (i) share a border with a treated municipality, and (ii) have the 'selection criteria' variables Z^1_{mt-1} and Z^2_{mt-1} above the thresholds values: $Z^1_{mt-1} \geq 0.7 \times 2,137$ km^2 and $Z^2_{mt-1} \geq 0.7 \times 222$ km^2. Land use data are from satellite images from PRODES (areas are measured in square kilometres). Deforested Area measures incremental deforestation during the year; Cumulative Deforested Area adds past deforestation up to and including 2007; Forested Area measures the total area covered by forests at the beginning of the year; Deforested Share divides incremental deforestation in 2007 by the forested area at the beginning of the year. The Number of Alerts comes from the DETER system. Fines Issued are obtained from IBAMA. Share of Protected Area is the proportion of the municipal area that is under legal protection (either indigenous land or conservation units), based on data from the National Register of Conservation Units. GDP consists of the municipalities' total GDPs, from IBGE's account system. Agricultural GDP includes crop and livestock production, from IBGE's Municipal Crop Survey and Municipal Livestock Survey. The variable Cattle measures the total number of cattle from IBGE's Municipal Livestock Survey. Crop Area measures the total area (in square kilometres) used to produce crops from IBGE's Municipal Crop Survey. Total Rural Credit measures the values of all contracted rural loans, aggregated up to the municipality-year level, based on data from the Brazilian Central Bank. All monetary amounts are expressed in December 2011 Reais. Annual rainfall is measured in millimetres (mm), while annual temperature is measured in degrees Celsius ($°C$) — both data are from Matsuura and Willmott (2012). FAO-GAEZ Soy and FAO-GAEZ Corn consist of maximum attainable crop yields at the field level, aggregated up to the municipality level. Distance to the nearest port is measured in kilometres using data from the Brazilian Ministry of Transportation. Carbon stocks are measured in tons of carbon per hectare (tC/ha), based on data developed by Baccini et al. (2012).

Table 2: Average and Cumulative Treatment Effects, without Spillovers

	Average Treatment Effects: Deforestation (km^2)		
	ATT	ATU	ATE
2009	-24.49	[-4.04, -3.67]	[-5.50, -5.15]
	(-27.64, -21.34)	(-4.16, -3.58)	(-5.65, -5.04)
2010	-57.97	[-6.84, -5.95]	[-10.49, -9.66]
	(-62.95, -52.99)	(-6.97, -5.85)	(-10.66, -9.52)

	Cumulative Treatment Effects, 2009-2010		
	Deforestation *(km^2)*	Carbon Emissions *(millions tC)*	Value *(billions U\$)*
Total Effects	-2886	-34.60	2.54
	(-3145, -2627)	(-37.96, -31.24)	(2.29, 2.78)

Notes: 95% confidence intervals are in parentheses. For ATT and the cumulative effects, the confidence intervals are computed based on the standard i.i.d. nonparametric bootstrap, where the i.i.d. resampling occurs in the cross-sectional dimension. For ATU and ATE, they are based on Imbens and Manski (2004). We implemented 500 bootstrap replications. Deforestation is measured in square kilometres. Emissions are measured in millions of tons of carbon. Values are measured in billion US$, assuming a social cost of carbon of US$ 20/tCO2. The calculations use the fact that 1 tC = (44/12) tCO2.

Table 3: Average and Cumulative Treatment Effects, with Spillovers

	Average Treatment Effects: Deforestation (km^2)			
	ATT	ATU	ATS	ATE
2009	-28.39	[-3.99, -3.52]	[-11.57, -11.53]	[-6.26, -5.86]
	(-31.73, -25.05)	(-4.08, -3.45)	(-13.57, -9.64)	(-6.40, -5.73)
2010	-58.76	[-6.29, -5.75]	[-20.02, -16.77]	[-10.99, -10.30]
	(-63.68, -53.84)	(-6.40, -5.66)	(-21.90, -14.68)	(-11.16, -10.15)

	Cumulative Treatment Effects, 2009-2010		
	Deforestation *(km^2)*	Carbon Emissions *(millions tC)*	Value *(billions U$)*
Direct Effects	-3050	-36.49	2.68
	(-3313, -2787)	(-39.94, -33.04)	(2.42, 2.93)
Indirect Effects	-1102	-12.90	0.95
	(-1232, -973)	(-14.50, -11.30)	(0.83, 1.06)

Notes: 95% confidence intervals are in parentheses. For ATT and the cumulative direct effects, the intervals are computed based on the standard i.i.d. nonparametric bootstrap, where the i.i.d. resampling occurs in the cross-sectional dimension. For ATU, ATS, ATE, and the cumulative indirect effects, they are based on Imbens and Manski (2004). We implemented 500 bootstrap replications. Deforestation is measured in square kilometres. Emissions are measured in millions of tons of carbon. Values are measured in billion US$, assuming a social cost of carbon of US$ 20/tCO2. The calculation uses the fact that 1 tC = (44/12) tCO2.

Table 4: Tests based on the Changes-in-Changes Model

	$\text{ATT}_{2009} = -24.49$, $\text{ATT}_{2010} = -57.97$					
	Placebo Test		'No Effect' Test		Stochastic Dominance Test	
	KS	CM	KS	CM	KS	CM
Unconditional	0.734	0.498	0.052	0.010	1.000	1.000
Residuals	0.706	0.724	0.000	0.000	1.000	1.000

Notes: The estimated treatment effects, ATT_{2009} and ATT_{2010}, are based on the 'no-spillover' case – consistent with all the tests presented in this table. The Placebo Test compares the factual and counterfactual distributions when we wrongly impose that the policy intervention was set in 2007; the null hypothesis states that the two distributions are equal to each other. The 'No Effect' Test is similar to the Placebo Test, but uses the correct timing of the intervention. The Stochastic Dominance Test assesses whether the counterfactual distribution is everywhere below the factual distribution. The test statistics are the Kolmogorov-Smirnov (KS) and the Cramer-von Mises (CM) statistics. We apply each test both on the log odds ratio of deforestation shares not conditioning on covariates (corresponding to the Unconditional row), and on the residuals, after partialling the covariates out (corresponding to the Residuals row). The cells present the p-values based on 500 bootstrap replications. The tests are proposed and developed by Melly and Santangelo (2015).

Table 5: Percent Correctly Predicted, Ex-post Optimal List

No Spillovers Case						
Constraint:	*Total Area*			*Number of Municipalities*		
	Optimal			*Optimal*		
	0	1		0	1	
Observed			*Percent Correct*			*Percent Correct*
0	382	73	83.96	439	16	96.48
1	10	25	71.43	16	19	54.29
		Overall	83.06		*Overall*	93.47

Spillovers Case						
Constraint:	*Total Area*			*Number of Municipalities*		
	Optimal			*Optimal*		
	0	1		0	1	
Observed			*Percent Correct*			*Percent Correct*
0	339	116	74.51	435	20	95.60
1	20	15	42.86	20	15	42.86
		Overall	72.24		*Overall*	91.84

Note: Authors' calculations

Table 6: Comparing Ex-Post Optimal, Priority, and Randomly Selected Lists

No Spillovers Case						
Constraint:	*Total Area*		*Number of Municipalities*			
	Observed vs Optimal		Observed vs Optimal		Random vs Optimal	
	Ratio	Value	Ratio	Value	Ratio	Value
Total Deforestation	1.06	-	1.06	-	1.22	-
Total Carbon Emissions	1.05	622	1.08	975	1.26	3,094

Spillovers Case						
Constraint:	*Total Area*		*Number of Municipalities*			
	Observed vs Optimal		Observed vs Optimal		Random vs Optimal	
	Ratio	Value	Ratio	Value	Ratio	Value
Total Deforestation	1.13	-	1.10	-	1.26	-
Total Carbon Emissions	1.12	1,292	1.12	1,278	1.30	3,271

Note: 'Ratio' divides total deforestation (total emissions) evaluated at the observed list by the ex-post optimal total deforestation (total emissions). 'Value' takes their difference. Values are measured in million US$, assuming a social cost of carbon = US$ 20/tCO2.