University of California
Santa Barbara

Evaluating impacts of forests and forest policy: methods and applications using satellite data

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Environmental Science and Management

by

Alberto Gabriel Garcia

Committee in charge:

Professor Robert Heilmayr, Chair
Professor B. Kelsey Jack
Professor Andrew Plantinga

June 2023

The Dissertation of Alberto Gabriel Garcia is approved.

Professor B. Kelsey Jack

Professor Andrew Plantinga

Professor Robert Heilmayr, Committee Chair

May 2023

Evaluating impacts of forests and forest policy: methods and applications using satellite data

by

Alberto Gabriel Garcia

Acknowledgements

I would like to thank my committee members Kelsey Jack and Andrew Plantinga. Both encouraged me to think more critically about the economics in the environmental questions I asked, and this work would not be possible without their support. I would especially like to thank my advisor, Robert Heilmayr, for his patience and guidance throughout the dissertation. He has been such a kind and thoughtful mentor throughout the entire dissertation process, and there were many times when he provided much needed support in the trenches of the Ph.D.. I feel extremely lucky to have had Robert as an example of a mentor.

I received helpful advice and encouragement from a number of people on the UCSB campus. I am particularly grateful for feedback and friendship from members of the Conservation Economics Lab and the UCSB Environmental Economics group. In particular, my Bren economics cohort-mates, Jacob Gellman and Vincent Thivierge, made the first year of the Ph.D. survivable and the subsequent years entertaining.

Thank you to my parents Judy and Al, and my sister Courtney, for their continued support and for always encouraging me to 'follow my dreams'. I would also like to thank Michelle Lee. I am grateful to have had her by my side and for her love, patience, and support.

I was fortunate to receive funding from several sources including the Bren School of Environmental Science & Management, the Environmental Markets Lab, the UCSB Graduate Division, the Schmidt Family Foundation, and the Arnhold Foundation.

Curriculum Vitæ

Alberto Gabriel Garcia

Bren School of Environmental Science & Management
2400 Bren Hall, office no. 3031
Santa Barbara, CA 93101

Education

Ph.D. Environmental Science and Management — expected 2023
University of California, Santa Barbara (UCSB)
Fields: Environmental and natural resource economics; Applied econometrics

B.A. Mathematics and Economics — 2016
Willamette University

Grants and Awards

Graduate Division Dissertation Fellowship, UCSB	2023
Schmidt Family Foundation Research Grant	2020-2021
Regents Fellowship, UCSB	2018-2020
Bren Fellowship, UCSB	2017-2018

At Willamette University: Phil Hanni Scholar 2016; Willamette University Academic Leadership Award 2016; Jason Lee Scholarship Award 2016

Publications

Garcia, A. (2020) "The Environmental Impacts of Agricultural Intensification". Technical Note N.9. Rome. Standing Panel on Impact Assessment (SPIA)

Presentations

2023

Stanford University, Center on Food Security and the Environment; University of Rhode Island, Department of Environmental and Natural Resource Economics; University of Utah, Department of Economics; Environmental Defense Fund

2022

FLARE 8th annual meeting, Rome, Italy; Oregon State University, College of Forestry; UCSB Bren School Symposium

2021

BIOECON XXII conference, Jackson, Wyoming; UCSB ERE Seminar

Teaching

Undergraduate courses:

ENVS 30: Environmental Economics, UCSB (TA) Fall 2019, Fall 2020, Spring 2021, Fall 2021

ECON 140A: Introduction to Econometrics, UCSB (TA) Fall 2018, Winter 2019

Graduate courses:

EDS 223 (Masters): Geospatial Analysis and Remote Sensing, UCSB (TA) Fall 2022

Service

Bren School Seminar Speaker Committee 2020-2021

Bren PhD Events Committee 2019

Abstract

Evaluating impacts of forests and forest policy: methods and applications using satellite data

by

Alberto Gabriel Garcia

Forests provide ecosystem services at a variety of scales, from local to global. Recent attention has focused on forests' potential to mitigate climate change because of their ability to store carbon. They also provide substantial local benefits, as trees can mitigate pollution, reduce extreme temperatures, and enhance psychological well-being. Unfortunately, threats such as logging and commodity agriculture have led to substantial deforestation of primary forest. Other factors such as drought, fire, and insects also pose a threat, impeding the provision of ecosystem services and undermining climate mitigation potential. In order to design policies that can effectively deliver on the promise of forests, better evidence is needed to 1) quantify the social costs and benefits that forests provide; and 2) understand how policy can be best designed to support both forest ecosystems and local actors.

Advances in earth observation have made more data detailing the dynamics of land cover and land use change available than ever before. In response, a growing body of work has emerged that integrates econometric methods of causal inference with these big data in order to evaluate the effectiveness of various policy designs. This has become particularly true in the context of forest conservation, where these quasiexperimental impact evaluations are increasingly used to inform new policy. However, factors inherent to their structure such as measurement error and irreversibility may affect the performance of common econometric approaches. Further work is needed to help researchers grapple

with how these data can be best integrated with econometric methods of causal inference, thereby providing more informative insight into policy design.

This dissertation seeks to answer three distinct but intertwined questions. In the first chapter, my co-author, Robert Heilmayr, and I ask how binary and irreversible data, a structure found in most deforestation datasets, affects the performance of panel econometric estimators. While the application of quasiexperimental impact evaluation to remotely sensed measures of deforestation has yielded important evidence detailing the effectiveness of conservation policies, researchers have paid insufficient attention to structure of these deforestation datasets. We use analytical proofs and simulations to demonstrate that many commonly employed panel econometric models are biased when applied to binary and irreversible outcomes. The significance, magnitude and even direction of estimated effects from many studies are likely incorrect, threatening to undermine the evidence base that underpins conservation policy adoption and design. To address these concerns, we provide guidance and new strategies for the design of panel econometric models that yield more reliable estimates of the impacts of forest conservation policies.

The second chapter asks how policymakers can best design forest policy in order to capitalize on the potential of forest-based climate solutions, while supporting livelihoods. In order to address the intertwined challenges posed by climate change, biodiversity loss, and rural poverty, policymakers throughout the world have begun to adopt policies that pay private landowners to protect or restore forests. However, relatively little evidence exists documenting the impacts these payments have had, and how incentives can be best structured to achieve multiple objectives. I evaluate the land cover impacts of a forest restoration subsidy included in Chile's Native Forest Law, which prioritized the participation of rural smallholders and indigenous communities. I find that 68.12% of landowners who had applied for the subsidy did not comply with their stated forest management commitments. However, verification protocols included as part of the conditional cash transfer program prevented nearly $30 million USD in unconditional transfers to these non-compliant landowners. Compliant landowners who were paid for forest restoration did expand native forests on their properties relative to a robust counterfactual. I find that the program has expanded Chile's native forests and paid the average landowner an estimated $36.78 USD per tonne CO_2 stored in aboveground biomass. In contrast to many studies on avoided deforestation, complying smallholders in regions of high poverty generated the greatest tree cover gains per enrolled hectare. These findings illustrate that, in contrast to payments for avoided deforestation, targeting for social development may enhance the environmental effectiveness of payments for restoration.

In the third and final chapter, I ask whether ecosystem degradation, specifically invasive species induced tree cover loss, has impacts on education outcomes in metropolitan United States. I leverage variation from the introduction of an invasive insect that exclusively targets ash trees, the emerald ash borer, to the Chicago Metropolitan area. Exploiting the staggered and idiosyncratic spread of the borer, I show how tree canopy cover and education outcomes were affected using difference-in-differences methods robust to general treatment effect heterogeneity. My findings indicate that ash borer infestation reduced canopy cover in affected areas by 1.4% on average, stemming from both increases in tree cover loss and declines in tree cover gain. Further, the ash borer reduced standardized test performance at exposed schools. Infestation exposure led to an average 1.86% fewer students meeting or exceeding the state benchmark at the typical school, with impacts concentrated among low-income students. This paper shows that invasive species can substantially impact ecosystem service provision and ultimately, education outcomes, adding to the damages known to be caused by human-induced environmental change.

Contents

Chapter 1

Conservation impact evaluation using remotely sensed data

1.1 Introduction

Policymakers often need to understand the causal impacts of conservation interventions. Can payments for ecosystem services encourage lasting reforestation? Do marine protected areas stop unsustainable harvesting of fish? While randomized experiments are the gold standard for the identification of causal relationships (Edwards et al., 2020; Jayachandran et al., 2017), conservation often poses questions that are prohibitively expensive, unethical or impossible to pursue through experimentation. In such settings, a growing portfolio of statistical techniques enable researchers to draw causal conclusions using observational data (Larsen et al., 2019; Ferraro and Hanauer, 2014; Miteva et al., 2012). Increasingly, these econometric approaches to impact evaluation are being used to disentangle the causal relationships that underpin conservation decision-making (Butsic et al., 2017a; Baylis et al., 2016; Williams et al., 2020).

Conservation has seen a proliferation of panel impact evaluation studies that detail the impacts of conservation interventions on land use change. This has been enabled, in part, by the increasing prevalence of remotely sensed datasets detailing forest cover outcomes through time (Blackman, 2013; Jones and Lewis, 2015). As a result, a scientist hoping to quantify the impacts of a land use policy adopted decades ago can assemble data for treated and control units that span multiple years both pre- and post-implementation periods (Jain, 2020). Deforestation is often measured using data with a similar structure to the Hansen Global Forest Change product (Hansen et al., 2013). These data yield binary observations detailing the first year in which each 30 by 30m pixel was deforested. Most studies that evaluate the effectiveness of forest conservation policies use these remotely sensed pixels as the unit of observation (Börner et al., 2020). Importantly, the data are unable to detect repeated deforestation events in the same location, so when converted to a panel structure, researchers have no information about a pixel beyond the year it is first cleared. In response, it has often been advised to drop pixels in the years after they are first cleared (Jones and Lewis, 2015; Alix-Garcia and Gibbs, 2017). This yields an unbalanced panel dataset of binary, irreversible deforestation events.

In this paper, we investigate whether this irreversible, binary data structure affects the performance of panel econometric methods typically used in conservation impact evaluation. Specifically, we focus on methods typically used in difference-in-differences (DID) settings, where the researcher observes land use change in treatment and control areas, both prior to and after some intervention of interest. DID and two-way fixed effects (TWFE) regression are among the most popular econometric approaches in observational panel data settings (de Chaisemartin and D'Haultfœuille, 2020; Sant'Anna and Zhao, 2020; Goodman-Bacon, 2021), and this extends to the conservation context.

We use a combination of analytical proofs and Monte Carlo simulations to demonstrate that many econometric analyses are likely biased when applied to panel data with this structure- significance, magnitude and even direction of estimated effects might be incorrect. The resulting biases arise even when researchers follow common guidance to adopt "rigorous" research designs with valid counterfactuals (Blackman, 2013; Jones and Lewis, 2015). While we focus primarily upon impact evaluation of forest conservation policies, our findings are relevant to a wider audience. Specifically, these results apply to diverse settings in which the outcome of interest represents an irreversible, binary event. Such events include recidivism (e.g. Agan and Makowsky, 2018; Mastrobuoni and Pinotti, 2015), mortality (e.g. Friedman and Schady, 2013), and unilateral technology adoption (e.g. Bollinger et al., 2022).

Our core result shows that TWFE regressions with individual unit fixed effects do not identify the desired treatment effect parameter when applied to panel datasets with binary, irreversible outcomes. This is the case even when typical common trends assumptions hold and in the absence of general treatment effect heterogeneity. It has been widely stated that the standard DID estimator is numerically equivalent to the linear TWFE estimator in the case of two-periods (each of which can consist of multiple years) and two-groups (2x2), where treatment is administered to only some units in the second period (Imai and Kim, 2021; Goodman-Bacon, 2021). However, when applied to binary, irreversible panel data with the structure we describe, we demonstrate that the TWFE estimator is distinct from the standard DID estimator. We show that the coefficient of interest in TWFE specifications instead recovers an ex-post difference in deforestation rates between treatment and control areas. This is particularly worrisome in the context of forest conservation, where interventions often target areas with high deforestation rates (e.g., Brazil's blacklisted Priority Municipalities) or low opportunity costs (e.g., protected areas). Papers published in both conservation science and economics journals frequently use this problematic specification to recover treatment effect estimates. This concern also extends to many recently developed DID estimators that estimate treatment effects in staggered adoption settings (e.g. Callaway and Sant'Anna, 2020; Borusyak et al., 2021;

Gardner, 2021), meaning that researchers cannot simply adopt a different estimator to avoid the issue. To help guide future impact evaluations, we identify multiple ways in which this bias can be reduced or even eliminated. In the land use context, one easily implemented solution is to aggregate the binary pixels spatially. Both pixel-level TWFE specifications with spatially aggregated unit fixed effects and TWFE specifications with spatially aggregated units of analysis recover the expected treatment effect parameter.

Researchers often account for irreversible, binary outcomes by adopting non-linear survival models. However, many invoke the traditional linear common trends assumption in this context. We demonstrate that this is not appropriate, as the survival analog to the traditional DID estimator relies instead on a proportional trends assumption. In general, this proportional trends assumption cannot simultaneously hold with the common trends assumption typically invoked in linear models. We develop a new survival-based DID estimator that relies on the more traditional common trends assumption.

We then explore non-random selection that arises due to irreversibility in the deforestation setting and how this feature of the data may lead to bias. Finally, we reflect on the econometric benefits that emerge when researchers are able to match their model structure to the relevant scale of the deforestation process and explore how area weighting may change the interpretation of researchers' treatment effect estimates.

This work makes important contributions to several literatures. First, our research contributes to an emerging literature aimed at better understanding how best to integrate remotely sensed data and econometric methods of causal inference (Jain, 2020; Gibson et al., 2021). Several studies have begun to document measurement error in satellite-based measurements and explore its implications for econometric analysis (e.g. Proctor et al., 2023). For example, Alix-Garcia and Millimet (2022) address misclassification in the context of a remotely sensed binary forest cover outcome and propose a solution for unbiased causal inference. Our paper implicitly assumes away issues of non-classical measurement error but demonstrates that researchers may still recover biased treatment effect estimates when using an irreversible, binary forest cover outcome.

We contribute to work in environmental economics and conservation science seeking to rigorously estimate the impacts of conservation policy. Causal impact evaluation in nature conservation has emerged relatively recently (Börner et al., 2020), and a wide array of papers has called for researchers to improve the rigor of their approaches in this space (Ferraro et al., 2019a; Avelino et al., 2016; Baylis et al., 2016; Miteva et al., 2012). While a burgeoning literature of quasiexperimental conservation impact evaluation has sought to meet this call, this study identifies key obstacles when integrating these econometric methods with frequently used sources of land cover data.

Finally, we add to the emerging literature on empirical practices in difference-in-differences settings. We complement the growing literature on the causal interpretation of coefficients from TWFE models, showing that individual unit-level TWFE regressions may not identify the expected treatment effect parameter when applied to binary, irreversible outcomes. While much of the recent work on TWFE models identifies concerns in the staggered adoption setting (e.g. Callaway and Sant'Anna, 2020; de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021), our results apply even in the 2x2 setting. We further show that the TWFE and DID estimators are numerically distinct in the two-period, two-group context in this setting. Further, we show that recently developed estimators suffer from similar issues to TWFE regressions in the context of irreversible, binary outcomes.

1.2 Empirical context

1.2.1 Analysis setting

We focus on the case in which a researcher would like to quantify the impact an intervention has had on deforestation rates. We assume that the intervention has clearly defined boundaries (e.g., a protected area, certified concession, or indigenous territory), and that the researcher has access to spatially explicit observations of forest cover and forest loss spanning multiple years over the periods before and after the intervention was adopted. This general setting describes a broad array of studies that apply panel methods to remotely sensed data. Table 1.1 shows an unsystematic review of avoided deforestation impact evaluation studies using panel econometric methods and data accessed as pixelated binary deforestation. There also exist a vast array of studies using matching as the primary identification approach (e.g. Pfaff et al., 2014; Robalino et al., 2015; Pfaff et al., 2015), but Table 1.1 focuses on only those incorporating panel approaches.

In each of the studies detailed in Table 1.1, the researcher's goal is to measure the impact that a specific policy had on deforestation within treated units, also known as the *average treatment effect on the treated* (ATT). The ATT estimates the difference between the average deforestation rate of treated units with treatment, and the average deforestation rate of treated units without treatment. The fundamental challenge is that, for every treated unit, the researcher is unable to observe the value that the outcome would have taken in the absence of treatment (Holland, 1986). In our case, this means that the researcher cannot observe the deforestation that would have occurred in treated units had they not received treatment.

Table 1.1: Panel econometric methods used in avoided deforestation impact evaluations. All included studies use data accessed as pixelated binary deforestation.

Paper	**Panel method**	**Unit of analysis**	**Unit FE**
Alix-Garcia and Gibbs 2017	TWFE	binary point/pixel	pixel
Alix-Garcia et al 2018	TWFE	binary point/pixel	pixel
Anderson et al. 2018	matched DID	binary point/pixel	
Araujo et al. 2009	TWFE instrument	state	state
Arriagada et al. 2012	matched DID	farm	
Baehr et al. 2021	TWFE	binary pixel/grid	pixel
Baylis et al. 2012	DID	grid cell	
BenYishay et al. 2017	TWFE	grid cell	grid cell
Blackman 2015	unit FE model	binary point/pixel	county
Blackman et al. 2017	TWFE	community	community
Blackman et al. 2018	matched TWFE	mgmt. unit	mgmt. unit
Busch et al. 2015	matched TWFE	grid cell	grid cell
Butsic et al. 2017	TWFE	binary point/pixel	pixel
Carlson et al. 2018 (1)	matched TWFE	plantation	plantation
Carlson et al. 2018 (2)	Cox PH DID	pixel	
Heilmayr and Lambin 2016	matched DID	property	
Heilmayr et al. 2020	Triple DID/FE	binary point/pixel	municipality
Herrera et al. 2019	matched regression	binary point/pixel	
Holland et al. 2017	matched TWFE	property	property
Jones and Lewis 2015 (1)	matched TWFE	binary point/pixel	pixel
Jones and Lewis 2015 (2)	matched TWFE	property	property
Jones et al. 2017	matched TWFE	household	household
Koch et al. 2018	matched DID	municipality	
Nolte et al. 2017	DID	property	
Panlasigui et al. 2018	TWFE	binary point/pixel	pixel
Rico-Straffon et al. 2022	staggered DID	grid cell	
Ruggiero et al. 2022	TWFE	municipality	municipality
Sales et al. 2022	Cox PH DID	pixel	
Shah and Baylis 2015	DID	grid cell	
Sims and Alix-Garcia 2017	TWFE	locality	locality
Tabor et al. 2017	TWFE	fokontany	fokontany
Wendland et al. 2015	matched TWFE	binary point/pixel	pixel

We model deforestation (y_{ivt}) as a binary choice by a landowner to clear a small plot of land i within their larger property v in year t, where $t \in T$. The decision to deforest depends upon a latent variable (y^*_{ivt}) that represents the returns from the plot of land in its cleared state ($V^{cleared}_{ivt}$) relative to the returns from its forested state ($V^{uncleared}_{ivt}$), such that:

$$y^*_{ivt} = V^{cleared}_{ivt} - V^{uncleared}_{ivt} \tag{1.1}$$

$$y_{ivt} = \begin{cases} 1 & \text{if } y^*_{ivt} > 0 \\ 0 & \text{otherwise} \end{cases} \tag{1.2}$$

This generic clearing rule underpins a broad class of more specific static and dynamic models that have been used to explore the determinants of deforestation (e.g. Pfaff, 1999; Kerr et al., 2003; Pfaff and Sanchez-Azofeifa, 2004).

However, this basic model makes an assumption that the decision to deforest is reversible. In reality, a number of characteristics of both the process of deforestation, as well as the methods used to detect deforestation in individual plots, complicate this assumption. First, the goal of many conservation interventions is to prevent the loss of mature forests that may take decades, if not centuries, to regrow. In such cases, deforestation itself may be considered irreversible in human time scales, focusing the researchers' attention upon the first instance in which a plot is deforested. Even when deforestation of secondary forests is an object of interest, constraints imposed by remotely sensed datasets may force empirical researchers to treat deforestation as irreversible. Gradual processes of reforestation are inherently harder to identify than abrupt losses of forest cover (Hansen et al., 2013). In addition, determining the precise year in which the extended process of forest regrowth began is currently an active area of research for the remote sensing community, and often requires many years of post-regrowth observations. As a result, commonly used deforestation datasets such as the Global Forest Change product often only identify the first year in which a pixel was cleared. Whether desired, or due to technical limitations, the resulting inability to observe repeated deforestation means that deforestation is, in effect, an irreversible process in most conservation impact evaluations. To incorporate this irreversibility into our model, we denote C_i as the first year in which $y^*_{ivt} > 0$, where y_{ivt} is not observed when $t > C_i$.

In response to this irreversibility, studies including Jones and Lewis (2015) and Alix-Garcia and Gibbs (2017) have suggested that deforested pixels should be dropped in the periods after they are first cleared. We follow this guidance, further modifying our outcome, y^o_{ivt}, the observed binary deforestation variable:

$$y^o_{ivt} = \begin{cases} 1 & t = C_i \\ 0 & t < C_i \\ - & t > C_i \end{cases} \tag{1.3}$$

Here, $-$ indicates that the outcome for pixel i in time t has been dropped from the panel in years where $t > C_i$.

Our parameter of interest, the ATT, is the average effect of an intervention on treated pixels. Let $y_{ivt}(1)$ and $y_{ivt}(0)$ denote the potential outcomes of pixel i in property v in year t with and without the treatment, respectively. In addition, let t_0 denote the first year in which the intervention of interest is implemented and let D_i represent a dummy indicating whether pixel i is ever treated. The ATT can now be expressed as:

$$ATT = E[y_{ivt}(1) - y_{ivt}(0) | t \geq t_0, D_i = 1] \tag{1.4}$$

DID and TWFE methods have become popular in part, because the researcher does not need random assignment of treatment to generate convincing estimates of a program's impact on avoided deforestation. Instead, the researcher must make a common trends assumption, under which we evaluate each method.

Assumption 1: (Common trends)

$$E[y_{ivt}(0)|t \geq t_0, D_i = 1] - E[y_{ivt}(0)|t < t_0, D_i = 1]$$
$$=$$
$$E[y_{ivt}(0)|t \geq t_0, D_i = 0] - E[y_{ivt}(0)|t < t_0, D_i = 0]$$

Assumption 1 requires that pixels in treated and untreated areas would have experienced the same change in their probability of deforestation across the two periods had no intervention occurred. While fundamentally untestable, researchers can take steps to address the plausibility of this assumption (Butsic et al., 2017a; Roth, 2022).

We also make the following stable unit treatment value assumption (SUTVA)

Assumption 2: (SUTVA)

$$\forall d \in \{0, 1\} : \text{ if } D_i = d \text{ and } t \geq t_0, \text{ then } y_{ivt}(d) = y_{ivt}$$

Assumption 2 requires that the potential outcomes for pixel i, $y_{it}(1)$ and $y_{it}(0)$, do not depend on the treatment status of any other pixel. There also cannot exist unobserved versions of treatment that may affect the potential outcomes.

1.2.2 Candidate empirical models

We present several empirical model specifications that we will evaluate and refer to throughout the remainder of the paper. These models have all been used in the forest conservation literature to estimate the ATT of specific conservation interventions. While some approaches, such as survival-analysis, are only beginning to emerge in the deforestation case, they are popular in other literatures in which the researcher wants to estimate the impact of an intervention on the occurrence of binary, irreversible events.

Traditional difference-in-differences estimator

Under the above two assumptions, perhaps the most popular panel approach to estimate the ATT in conservation impact evaluation is the traditional DID regression:

Regression 1: (DID regression) Let β_{DID} denote the coefficient of the interaction between D_i and an indicator for whether the intervention has been implemented in time t, $\mathbb{1}\{t \geq t_0\}$, in the following (population) OLS regression:

$$y^o_{ivt} = \alpha_0 + \alpha_1 D_i + \alpha_2 \mathbb{1}\{t \geq t_0\} + \beta_{DID} \times D_i \mathbb{1}\{t \geq t_0\} + \epsilon_{ivt}$$

Conceptually, the DID estimator calculates the treatment effect as the difference between the differences of the treated and untreated observations before and after treatment (Butsic et al., 2017a).

$$\begin{aligned}\beta_{DID} = {} & E[y^o_{ivt}|t \geq t_0, D_i = 1] - E[y^o_{ivt}|t < t_0, D_i = 1] \\ & -(E[y^o_{ivt}|t \geq t_0, D_i = 0] - E[y^o_{ivt}|t < t_0, D_i = 0])\end{aligned}$$

When the y^o_{ivt}s are i.i.d. and Assumptions 1 and 2 hold, it is straightforward to show that

$$\beta_{DID} = ATT$$

Individual unit-level TWFE regression

Researchers often want to estimate the ATT in a setting that does not fit the two-group, two-period case covered by the standard DID model. In such cases, TWFE regressions are frequently used to apply DID methods to multiple groups or treatment periods (Imai and Kim, 2021; Goodman-Bacon, 2021). This amounts to estimating a regression that includes individual unit and time fixed effects to control for unobservable confounding variables that vary across units or through time.

Regression 2: (Individual unit TWFE regression) Let β_{TWFE} denote the coefficient of the interaction between D_i and $\mathbb{1}\{t \geq t_0\}$ in the following (population) OLS regression:

$$y^o_{ivt} = \beta_{TWFE} \text{ x } D_i \mathbb{1}\{t \geq t_0\} + \lambda_t + \gamma_i + \epsilon_{ivt}$$

Here λ_t and γ_i represent the year and individual unit fixed effects, respectively. In the context of forest conservation, the individual unit i represents the pixel.

In the case of two groups and two time periods, the TWFE regression typically yields an estimate of the ATT that is numerically equivalent to the estimate generated by the DID model (Wooldridge, 2010; Imai and Kim, 2021). With this in mind, many researchers have used the TWFE model as a "generalized DID" that can be estimated not only in the 2x2 case, but also in settings where different units are exposed to treatment in more than two distinct time periods (Table 1.1). For example, a researcher may use a TWFE regression model to examine the effectiveness of a network of protected areas where the protected areas were created at different times, or a payment for ecosystem services (PES) program that enrolled properties in annual cohorts.

Aggregated unit fixed effects

One can also use aggregated units in the TWFE specification. Regression 3 outlines the use of unit fixed effects at the level of an aggregated unit rather than the individual binary unit. In the forest conservation context, aggregation is generally done spatially, aggregating pixels into larger units such as a grid cell (e.g. Rico-Straffon et al., 2023), property (e.g. Heilmayr and Lambin, 2016), or larger administrative unit (e.g. Pfaff, 1999). Individual unit-level TWFE models with aggregated fixed effects are all in the form of Regression 4.

Regression 3: (Individual unit-level TWFE regression with aggregated unit fixed effects) Let $\beta_{FE,j}$ denote the coefficient of the interaction between D_i and $\mathbb{1}\{t \geq t_0\}$ in the following (population) OLS regression:

$$y^o_{ivt} = \beta_{FE,j} \text{ x } D_i \mathbb{1}\{t \geq t_0\} + \lambda_t + \gamma_j + \epsilon_{ivt}$$

, where λ_t denotes year fixed effects and γ_j denotes fixed effects at the level of an aggregated unit. If k differs from the level of treatment assignment, one must also include a treatment group indicator or fixed effects at the level of the unit at which treatment is assigned.

Aggregated units of analysis

Another potential solution to the bias associated with TWFE models is for researchers to aggregate multiple binary unit-level observations into larger units of analysis. Regression 4 details specifications of this type. The researcher must now calculate the event rate (i.e., deforestation rate) in each time period within the relevant unit.

Regression 4: (TWFE regression with aggregated unit of analysis) Let β_j denote the coefficient of the interaction between D_j and $\mathbb{1}\{t \geq t_0\}$ in the following (population) OLS regression:

$$z_{jt} = \beta_j \text{ x } D_j \mathbb{1}\{t \geq t_0\} + \lambda_t + \gamma_j + \epsilon_{jt}$$

, where λ_t denotes year fixed effects and γ_j denotes fixed effects at the level of the aggregated unit. Regression 5 differs from Regression 4 in both the treatment variable, D_j and the outcome variable, z_{jt}. The treatment variable $D_j = \frac{1}{N_j} \sum_{i=1}^{N_j} D_i$, is the average treatment value amongst all pixels in unit j. If treatment is assigned at the level of j, $D_j = 1$. The outcome variable, z_{jt}, denotes the event rate within unit j in period t.

Arguably the most commonly used formula in the deforestation literature to calculate z_{jt}, the deforestation rate in unit j in time t, uses the share of unit j with forest cover and its lag (e.g. Carlson et al., 2018; Busch et al., 2015):

$$z_{jt} = \frac{F_{j,t-1} - F_{j,t}}{F_{j,t-1}} \tag{1.5}$$

, where $F_{j,t}$ and $F_{j,t-1}$ are the share of spatial unit j with forest cover at times t and $t-1$, respectively.

Survival analysis

Survival analysis has emerged as a common approach to modeling irreversible events (Emmert-Streib and Dehmer, 2019). It used frequently to model events such as mortality (e.g. Puterman et al., 2020) and recidivism (e.g. Luallen et al., 2018), but can also be applied to deforestation contexts. Survival models, such as the Cox Proportional Hazards model, quantify how covariates relate to changes in the length of time that a unit remains in a sample. In the case of deforestation, survival analyses can be used to explore how policy adoption changes the duration that treated, forested pixels survive until they are first cleared.

Despite the theoretical appeal of using survival models to study deforestation, they are still relatively uncommon in conservation impact evaluation. One emerging approach introduces the intuition of a difference-in-differences research design into a Cox Proportional Hazards model (e.g. Heilmayr et al., 2020b; Sales et al., 2022). Specifically, researchers estimate a Cox proportional hazards model of the following general form (e.g. Mastrobuoni and Pinotti, 2015):

Regression 5: (Cox DID regression) Let β_{coxDID} denote the coefficient of the interaction between D_i and $\mathbb{1}\{t \geq t_0\}$ in the following (population) OLS regression:

$$h(t) = \delta_0(t)exp(\alpha_0 + \alpha_1 D_i + \alpha_2 \mathbb{1}\{t \geq t_0\} + \beta_{coxDID} \text{ x } D_i \mathbb{1}\{t \geq t_0\} + \epsilon_{it})$$

, where $h(t)$ is the hazard rate of deforestation, t years into the study period; and $\delta_0(t)$ is the baseline hazard function.

1.3 Methods

1.3.1 Analytical proofs

The rapid growth of the conservation impact evaluation literature has resulted in a diversity of model structures that all attempt to estimate the effectiveness of conservation interventions (Table 1.1). However, researchers have not adequately considered how the irreversible, binary structure of these data may affect the properties of some estimators. We use analytical proofs to demonstrate several concerns that arise with the use of specific model specifications in the context of an irreversible, binary outcome. All of our analytical results apply to any setting with an irreversible, binary outcome, and we have indeed identified several studies outside of the forest conservation literature that fall prey to the issues we present.

For our analytical results, we restrict ourselves to the case where $t \in \{1, 2\}$ and $t_0 = 2$. We can then denote y^0_{iv1} and y^0_{iv2} as the outcomes for individual unit i in the first and second periods, respectively. In this setting, $y^0_{iv1} \in \{0, 1\}$ and $y^0_{iv2} \in \{0, 1, -\}$, meaning that although only outcomes in the second period are dropped, there exists variation in deforestation across both periods. We present the core analytical results in the main text, and the detailed proofs can be found in the Appendix.

1.3.2 Simulation models

To verify our analytical results and explore the relative performance of different models frequently used in conservation impact evaluation, we employ a series of Monte Carlo simulations. Specifically, we randomly generate synthetic landscapes with known policy effectiveness and analyze the performance of different econometric models in estimating the policy's known impact.

Landscape configuration

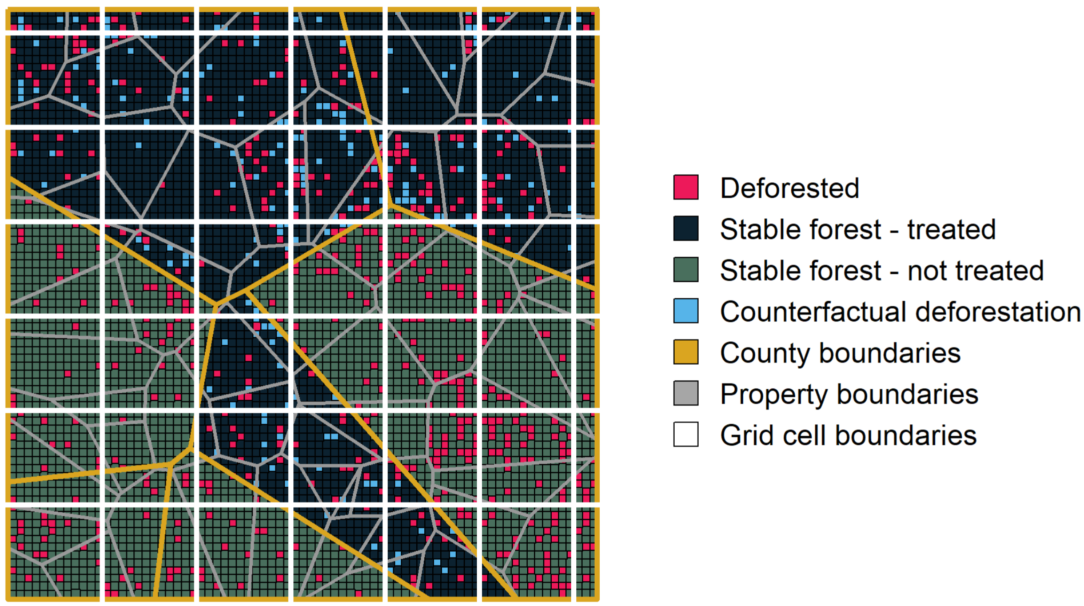

Figure 1.1: A map of a simulated landscape depicting patterns of deforestation under an effective conservation intervention, as well as counterfactual deforestation illustrating what would have happened in the absence of the intervention.

Figure 1.1 displays a simulated conservation intervention that reduced deforestation rates in treated areas — the landscape is depicted as observed by the researcher at the end of the observation period, including the unobservable counterfactual of what would have happened if the conservation intervention had not been adopted. Note that in untreated areas, there is no counterfactual deforestation, since no intervention ever took place. We begin each Monte Carlo simulation by creating a synthetic landscape consisting of 150 rows and 150 columns of square pixels (22,500 total pixels), equivalent to a raster that is four times larger than what is illustrated in Figure 1.1. We assume that each pixel has a resolution of 30 meters, comparable to the resolution of many Landsat-based, remote sensing analyses. The landscape thus represents an area of approximately

20.25 km2. We then divide this landscape into a variety of spatial units, composed of either uniform aggregations of pixels (i.e. large or small “grid cells"), or randomly spaced Thiessen polygons (i.e. “counties" or “properties"). Grid cells are intended to represent arbitrary units of spatial aggregation imposed by the researcher. In contrast, counties and properties are intended to represent simulated administrative units over which policy or land use decisions are made. Table 1.2 summarizes the relative scale of each of these spatial units under our baseline specifications.

Table 1.2: Spatial unit structure and size

Spatial unit	**Spatial structure**	**Avg. number of pixels**	**Area (hectares)**
Property	Thiessen polygons	100	9
County	Thiessen polygons	900	81
Large grid	Uniform square	900	81
Small grid	Uniform square	100	9

Data generating process

Each of our simulated landscapes consists of administrative units that are either untreated ($D_i = 0$) or are assigned to a conservation treatment ($D_i = 1$). We observe deforestation in two even-length periods, each of which consists of multiple years.

We follow Equation 1.1 and model these binary deforestation events as a function of each pixel's unobservable value along the continuous, latent variable (y^*_{ivt}) indicating the return to clearing pixel i, in property v, in year t.

$$\begin{aligned} y^*_{ivt} =& V^{cleared}_{ivt} - V^{uncleared}_{ivt} \\ =& \beta_0 + \beta_1 D_i + \beta_{2,0}(1 - D_i)\mathbb{1}\{t \geq t_0\} + (\beta_{2,1} + \beta_3) D_i \mathbb{1}\{t \geq t_0\} + \alpha_i + u_{it} + \rho_v \end{aligned}$$

That is, the returns to deforestation evolve over the two time periods ($\mathbb{1}\{t \geq t_0\}$), and differ across the control ($D_i = 0$) and treated pixels ($D_i = 1$). In addition, we assume that the value of deforestation is influenced by time-invariant random disturbances at the scale of individual pixels ($\alpha_i \sim N(0, \sigma_a^2)$) or properties ($\rho_v \sim N(0, \sigma_p^2)$), as well as time-varying, pixel-scale disturbances ($u_{it} \sim N(0, \sigma_u^2)$). These disturbances can represent a variety of spatial and temporal processes including, for example, the biophysical characteristics of a location, or the preferences of a property owner.

The potential outcomes for the latent variable, y^*_{ivt}, are as follows:

$$y^*_{ivt}(0) = \beta_0 + \beta_1 D_i + \beta_{2,0}(1 - D_i)\mathbb{1}\{t \geq t_0\} + \beta_{2,1} D_i \mathbb{1}\{t \geq t_0\} + \alpha_i + u_{it} + \rho_v \tag{1.6}$$

and

$$y^*_{ivt}(1) = \beta_0 + \beta_1 + \beta_{2,1} + \beta_3 + \alpha_i + u_{it} + \rho_v \tag{1.7}$$

The ATT in our simulated setting is, therefore, defined:

$$\begin{aligned} ATT =& P(y^*_{ivt}(1) > 1 | D_i = 1, t \geq t_0) - P(y^*_{ivt}(0) > 1 | D_i = 1, t \geq t_0) \\ =& P(\beta_0 + \beta_1 + \beta_{2,1} + \beta_3 + \alpha_i + u_{it} + \rho_v > 1) \\ & - P(\beta_0 + \beta_1 + \beta_{2,1} + \alpha_i + u_{it} + \rho_v > 1) \end{aligned}$$

Assumed parameter values and evaluation criteria

For the remainder of the paper, we explore a guiding example that has been parameterized to represent an impactful intervention in a high deforestation setting. Conservation interventions often have annual treatment effects smaller than a 1 percentage point reduction in the annual deforestation rate (e.g. Robalino and Pfaff, 2013; Jones et al., 2017). These modest reductions in the annual deforestation rate, however, can amount to large landscape-scale effects. For example, Alix-Garcia et al. (2018) find that environmental land registration in Brazil's Amazonian states of Mato Grosso and Para reduced the annual deforestation rate by an average of 0.5 percentage points, which has resulted in an overall reduction in deforestation of 10%.

Our initial simulated landscape has the following characteristics: 6 years in each of the pre-treatment and post-treatment periods ($T = 12$, $t_0 = 7$); a pre-treatment deforestation rate of 2% in the control area; a pre-treatment deforestation rate of 5% in the intervention area; a decrease in the deforestation rate of 0.5 percentage points between the first and second period in the absence of treatment; and an average reduction of 1 percentage point in the deforestation rate in treated units due to the intervention ($ATT = -0.01$). We assume that $\sigma_u = 0.5$. Finally, we begin by assuming away time invariant pixel ($\sigma_a = 0$) and property-level disturbances ($\sigma_p = 0$) but relax this assumption in Section 1.5.6. Note that Assumptions 1 and 2 are satisfied by construction. The derivations detailing the mapping from the landscape characteristics to the corresponding parameters in y^*_{ivt} can be found in Appendix A.1.7.

We compare econometric models using a combination of estimate bias, root mean squared error (RMSE), and coverage probability based on 500 replications of each set of parameters. Using our Monte Carlo simulations, we calculate estimate bias as the difference between each model's mean estimate of the ATT and the known ATT parameter. RMSE describes the distribution of estimates around the ATT. Coverage probability is defined as the proportion of simulations in which the true ATT lies within the simulation's 95% confidence interval (CI). As such, we would expect the ATT to lie within this CI 95% of the time, however, factors such as the bias of the estimates, their distribution, and treatment of standard errors may impact coverage.

1.4 Analytical results

1.4.1 Irreversible Point FE bias: TWFE does not identify *ATT* with binary, irreversible outcomes

Despite widespread use of pixel-level analyses of deforestation, the application of TWFE models to a binary, irreversible process yields a biased estimate of the ATT. Specifically, we show that, the coefficient of interest from the point TWFE model (β_{TWFE}) estimates the post-treatment difference in outcomes (single difference), rather than the desired ATT (full proof contained in Appendix A.1.2).

Result 1: (Irreversible unit FE bias)

$$\beta_{TWFE} = ATT + \underbrace{E[y_{iv1}(0)|D_i = 1] - E[y_{iv1}(0)|D_i = 0]}_{\text{pre-treatment difference in deforestation rates}} \tag{1.8}$$

Regression 2 thus forgoes the benefits that panel methods provide, and if the treated area has a different baseline deforestation rate than the control, will generate a biased estimate of the intervention's impact. Many conservation interventions are specifically targeted towards locations with either low opportunity costs for conservation or high threats of conversion. As a result, it is likely that many conservation impact evaluations will have treatment and control units that experienced different pre-treatment deforestation rates. It is important to note that this bias could even lead to changes in the estimated treatment effect's sign, in addition to errors in the effect's magnitude and significance.

Intuition for this result stems from the use of point unit fixed effects in the regression specification. By including point fixed effects in TWFE regressions, researchers hope to control for local confounders, including pre-treatment differences in the outcome. However, when following common guidance to drop observations in the periods after the irreversible event is first realized, these fixed effects do not behave as the researcher expects. Observations that realize the event (i.e. are deforested) in years prior to treatment are by definition not observed for the entire panel. Implicitly, the pre-treatment outcomes get assigned as 0 for both the treatment and control groups.

The next section verifies this result and shows that the traditional DID does not suffer from similar concerns in the context of our simulated landscapes and forest conservation intervention. In section A.1.6, we also show that $\hat{\beta}_{TWFE}$ is equivalent to the coefficient from a TWFE regression on a dataset where any pixel deforested in years prior to treatment is simply dropped from the dataset completely.

1.4.2 Alternative construction of y^o_{ivt}

Although dropping previously deforested pixels from the panel introduces bias into the TWFE estimate of the ATT, keeping observations in the panel after initial deforestation introduces its own questions, particularly if researchers are interested in deforestation rates. The ATT as defined in Section 1.2 is an estimate of the impact of an intervention on the frequency of deforestation events (i.e. the decision to clear). Keeping the deforested pixel in the panel beyond the first period in which it was observed as deforested would incorrectly imply that it has actively been deforested in each subsequent time period, when in fact, no new deforestation event or clearing decision has occurred. This is intuitively problematic, because the deforestation rate in each period would be monotonically increasing by construction, which is not necessarily the case. In the context of deforestation, this approach estimates an intervention's impact on deforested area, rather than on deforestation rates. However, economists are often interested in an interventions effect on clearing behavior, which is better represented by the deforestation rate or clearing event. Further, when reforestation timing is unaccounted for, as is the case with most data products using a binary measure of deforestation, it is not clear that this approach of keeping deforested pixels recovers the true measure of interest. Rather than measuring deforested area at any given time, this outcome variable measures the stock of ever-deforested area through the current time period. Primary forest may be the exception, where this loss truly is irreversible. In any case, this should not be interpreted as the deforestation rate (Appendix A.1.4).

There may be cases in which researchers do want to keep observations in the panel dataset after an irreversible event is realized. For example, consider educational attainment, where a researcher is interested in how an intervention impacts the education levels of a population. In this case, one likely wants to construct an outcome reflecting the proportion of the population educated rather than an “education rate". An example reflecting the opposite scenario may be mortality, where the researcher is interested in mortality rates, not overall mortality.

1.4.3 Survival analysis

Hazard rate ratios from a single survival model do not estimate the ATT under common trends

Multiple studies across a wide variety of settings have interpreted the resulting exponentiatated coefficient $exp(\beta_{coxDID})$ as a hazard ratio that is indicative of the impact that treatment has had on the relative likelihood of survival. Specifically, based on the way this hazard ratio is interpreted in multiple papers, it appears that many researchers expect this hazard ratio to represent the ratio of the hazard rates in the treatment group post-treatment, relative to the counterfactual in that group had treatment not occurred. This desired hazard ratio measuring the relative impact of treatment on the treated, which we denote as the $HRTT$, can be considered a reframing of the traditional ATT as a ratio rather than a difference:

$$HRTT = \frac{E[y_{iv2}(1)|D_i = 1]}{E[y_{iv2}(0)|D_i = 1]}$$

Both in conservation and alternative settings, researchers using Regression 3 have made Assumption 1, and evaluate whether it is plausible in their setting. However, Appendix A.1.3 shows that $exp(\beta_{coxDID})$ only identifies the $HRTT$ under an alternative assumption:

Assumption 3: (Proportional trends)

$$\frac{E[y_{iv2}(0)|D_i = 1]}{E[y_{iv1}(0)|D_i = 1]} = \frac{E[y_{iv2}(0)|D_i = 0]}{E[y_{iv1}(0)|D_i = 0]}$$

Assumption 3 requires that pixels in treated and untreated areas would have experienced the same ratio of change in their probability of deforestation across the two periods had no intervention occurred. Note that Assumption 1 and Assumption 3 cannot simultaneously hold (unless there is no trend at all). This means that researchers estimating Regression 3 under the traditional common trends assumption (Assumption 1) will not recover the $HRTT$, the relevant treatment effect parameter.

Proposing a new survival analysis-based estimator of the ATT

To the best of our knowledge, no prior studies have successfully combined the Cox Proportional Hazards model and the difference in differences research design to recover an unbiased estimate of the ATT under the traditional common trends assumption (Assumption 1). Here we outline a new estimation approach that first recovers an unbiased estimate of the $HRTT$ and then translates this into an estimate of the ATT that holds under Assumption 1. First, we note that the desired $HRTT$ can be re-written as a combination of three different hazard ratios:

$$HR_1 = \frac{E[y_{iv2}(1)|D_i = 1]}{E[y_{iv1}(0)|D_i = 1]} \tag{1.9}$$

$$HR_2 = \frac{E[y_{iv2}(1)|D_i = 1]}{E[y_{iv2}(0)|D_i = 0]} \tag{1.10}$$

$$HR_3 = \frac{E[y_{iv2}(0)|D_i = 0]}{E[y_{iv1}(0)|D_i = 0]} \tag{1.11}$$

$$HRTT = \frac{E[y_{iv2}(1)|D_i = 1]}{E[y_{iv2}(0)|D_i = 1]} = \frac{1}{1/HR_1 + 1/HR_2 - 1/(HR_2 * HR_3)} \tag{1.12}$$

Each of the three hazard ratios, HR_1, HR_2, and HR_3, can be estimated through separate Cox Proportional Hazards models estimated on subsets of the larger dataset. Specifically:

* $HR_1 = exp(\alpha)$, where α is estimated by subsetting to observations from the treated group ($D_i = 1$), and estimating the hazard rate of deforestation at time t as $h(t) = \lambda_0(t)exp(\alpha 1\{t \geq t_0\})$; and

* $HR_2 = exp(\beta)$, where β is estimated by subsetting to observations from the post-treatment period ($t \geq t_0$), and estimating the hazard rate of deforestation at time t as $h(t) = \gamma_0(t)exp(\beta 1\{D_i = 1\})$;

* $HR_3 = exp(\delta)$, where δ is estimated by subsetting to observations from the untreated group ($D_i = 0$), and estimating the hazard rate of deforestation at time t as $h(t) = \psi_0(t)exp(\delta 1\{t \geq t_0\})$.

Because the numerator of $HRTT$, $E[y_{iv2}(1)|D_i = 1]$, can be estimated as the mean of post-treatment deforestation rates in the treated group (denoted $\widehat{defor}_{D_i:1,t\geq t_0}$), we can estimate the ATT using this estimated deforestation rate and our estimate of $HRTT$:

$$\widehat{ATT}_{Cox} = \widehat{defor}_{D_i:1,t\geq t_0} - \frac{\widehat{defor}_{D_i:1,t\geq t_0}}{\widehat{HRTT}} \tag{1.13}$$

We have shown that the simple extension of the traditional DID to the survival setting only recovers an easily interpretable measure of a policy's impact under an assumption that cannot simultaneously hold with the traditional common trends assumption, the "Proportional Trends" assumption. In contrast, our proposed estimator, which relies on separate estimation of relevant hazard ratios, does recover the relevant analog of the ATT under the traditional common trends assumption (Assumption 1). We explore the performance of $\widehat{ATT}_{Cox}$ relative to the proposed OLS regressions under various circumstances likely to arise in the deforestation setting in the next sections. If a researcher opts to use survival analysis to recover an intervention's impact, their choice of estimator should depend on which trends assumption is plausible in their specific setting.

1.4.4 Non-random sample selection can generate bias in irreversible settings

Irreversibility in observed deforestation creates the potential for non-random sample selection. Specifically, deforested pixels are no longer at risk of clearing in the periods after they are first deforested. This means the "at risk" set of pixels changes through time as more pixels become deforested. As such, the distribution that describes the returns to clearing the at-risk pixels may change through time as well, leading to non-random selection of the sample through time. For example, pixels with extremely high returns to clearing are more likely to be cleared early on, regardless of treatment status. In subsequent periods, therefore, these high return pixels are less likely to be present in the sample at all. In the context of two-groups and two-periods, only the second period suffers from this non-random sample selection. We express the bias introduced from non-random sample selection below.

Result: Under Assumptions 1 and 2, in the two-group, two-period case, β_{DID} suffers from non-random sample selection bias when the y^o_{ivt}s are not i.i.d.

$$\beta_{DID} = ATT + \underbrace{E[y^o_{iv2}|D_i = 1] - E[y_{iv2}|D_i = 1] - (E[y^o_{iv2}|D_i = 0] - E[y_{iv2}|D_i = 0])}_{\text{bias emerging from non-random sample selection}} \tag{1.14}$$

Proof: Appendix A.1.5 Q.E.D.

In essence, the first and third expectations in the bias term are conditional on the pixel remaining forested after the first period.

1.5 Simulation results

1.5.1 TWFE bias

TWFE models have risen to prominence due to their flexibility in applying DID methods to settings with multiple groups and variation in treatment timing. However, we have shown that TWFE models with pixel fixed effects are not a viable approach to estimate the ATT. Figure 1.2 shows the bias associated with Regression 2, a pixel-level TWFE regression with pixel unit fixed effects. In our guiding example, the ex-post single difference is 0.02 (the ATT plus the post-treatment group difference in deforestation rates), when the true ATT is equal to -0.01. This means that a positive bias of 0.03 results from the use of this regression model.

Figure 1.2 explicitly contradicts the common wisdom that 2x2 TWFE and DID estimators are numerically equivalent, a finding unique to this setting. We further examine this finding in section A.1.6 and show that the coefficient of interest from Regression 2 is numerically equivalent to that from the same regression on a dataset where all pixels deforested in the first period are dropped completely.

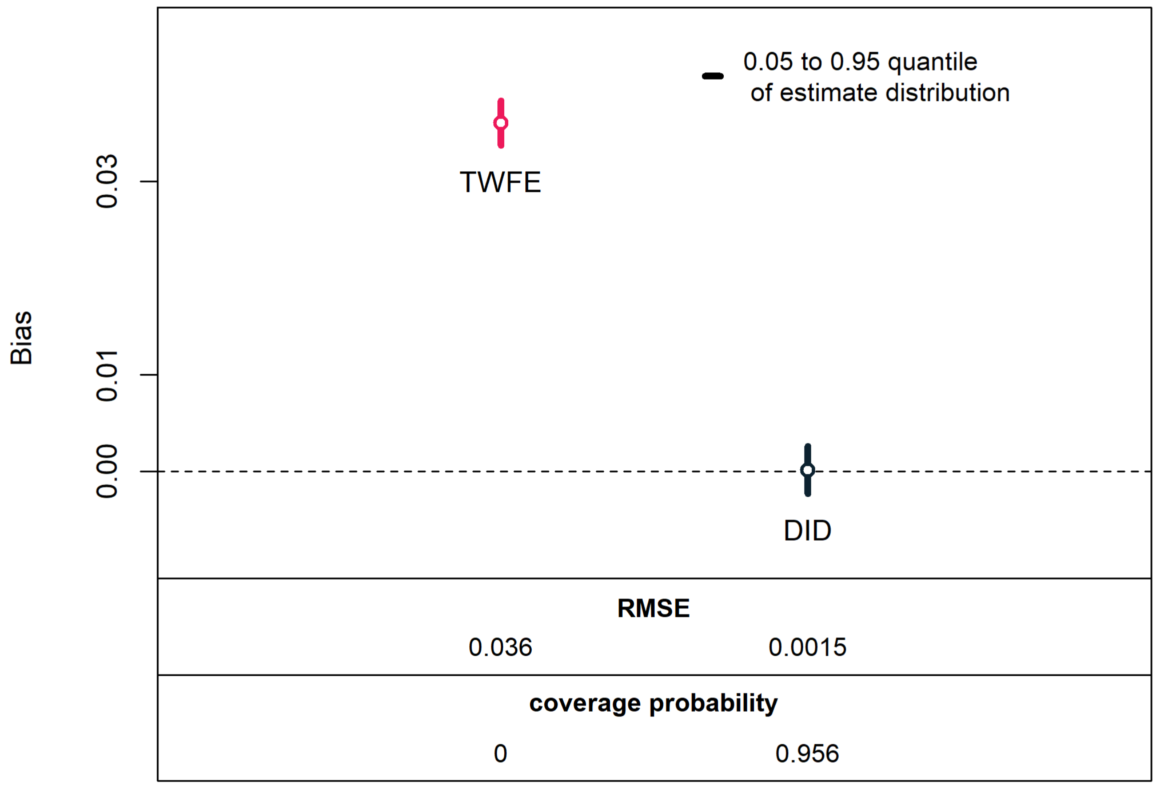

Figure 1.2: This figure shows the mean bias along with the 0.05 to 0.95 quantile range for estimates from the pixel-level DID and TWFE with pixel unit fixed effects. It is clear that the two are numerically distinct and that the coefficient of interest from the TWFE regression suffers from significant bias in this context. Standard errors are clustered at the pixel.

1.5.2 Traditional difference-in-differences model

In the two-group, two-period case the traditional DID (Regression 1) is an unbiased estimator of the ATT, as shown in Figure 1.2 and column 1 of Figure 1.3 (The traditional DID is equivalent to including treatment fixed effects). However, researchers often want to use TWFE models because of their flexibility in situations that do not fall under the simplest DID setting. Therefore, we explore the trade-offs below when using aggregated units of analysis and fixed effects when using TWFE models for impact evaluation with binary irreversible outcomes.

1.5.3 Models using spatial aggregation

We proposed using aggregation as a way to avoid the issues associated with integrating TWFE models with binary outcomes when the outcome is irreversible. Columns 2-4 of Figure 1.3 show that in pixel level TWFE regressions in the form of Regression 3, $\beta_{FE,j}$ is an unbiased estimator of the ATT. These regressions use grid, county, or property fixed effects rather than pixel fixed effects. We also see that, in the absence of property level perturbations (i.e. $\sigma_p = 0$ in the DGP), all three models provide similar estimates and estimate distributions.

Similarly, columns 5-7 show that that β_j from TWFE regressions in the form of Regression 4 is an unbiased estimator of the ATT. These TWFE regressions use an aggregated unit of analysis, where the deforestation rate for unit j is explicitly calculated in each time t. The following results are based on Equation 1.5. Neither the bias nor RMSE of the estimates vary dramatically across different levels of aggregation.

1.5.4 Survival analysis

Survival analysis provides an appealing alternative to traditional linear estimators when studying irreversible changes such as deforestation. The simple, single-regression DID framing of the Cox Proportional hazards model (Regression 5), however, is not a viable solution (Section 1.4.3) under the typical common trends assumption (Assumption 1). Column 8 of Figure 1.3 shows the bias associated with this model in our parameterization where Assumption 1 holds but Assumption 3 does not. In light of the fact that the simple Cox DID is not a viable analog to the traditional DID, we explore whether our proposed estimator, $\widehat{ATT_{Cox}}$, recovers relevant treatment effect parameters in the deforestation setting. We see in column 9 that $\widehat{ATT_{Cox}}$ indeed recovers our parameter of interest, the ATT.

Although our proposed, survival analysis-based approach to estimation yields a good estimate of the true ATT in this simple setting, multiple considerations raise questions about the utility of this non-linear model in more complex settings. One of the primary reasons for the use of survival analysis is censoring. This occurs when the researcher has partial information about the subjects' survival times but does not have access to precise event times. While the researcher may not observe all pixels until they are deforested, many other common forms of censoring are rarely a concern in the context of deforestation since remote sensing typically enables the creation of balanced panels. Further, the proposed strategy to drop pixels in the periods after they are first deforested successfully addresses irreversibility in deforestation events. Finally, as we show in the following sections, survival analysis is likely to suffer from bias when the data generating process underpinning deforestation is influenced by unobservable characteristics of more aggregated spatial units such as the preferences of a property-owner. While researchers can account for this in OLS specifications, there is no clear solution in survival analysis.

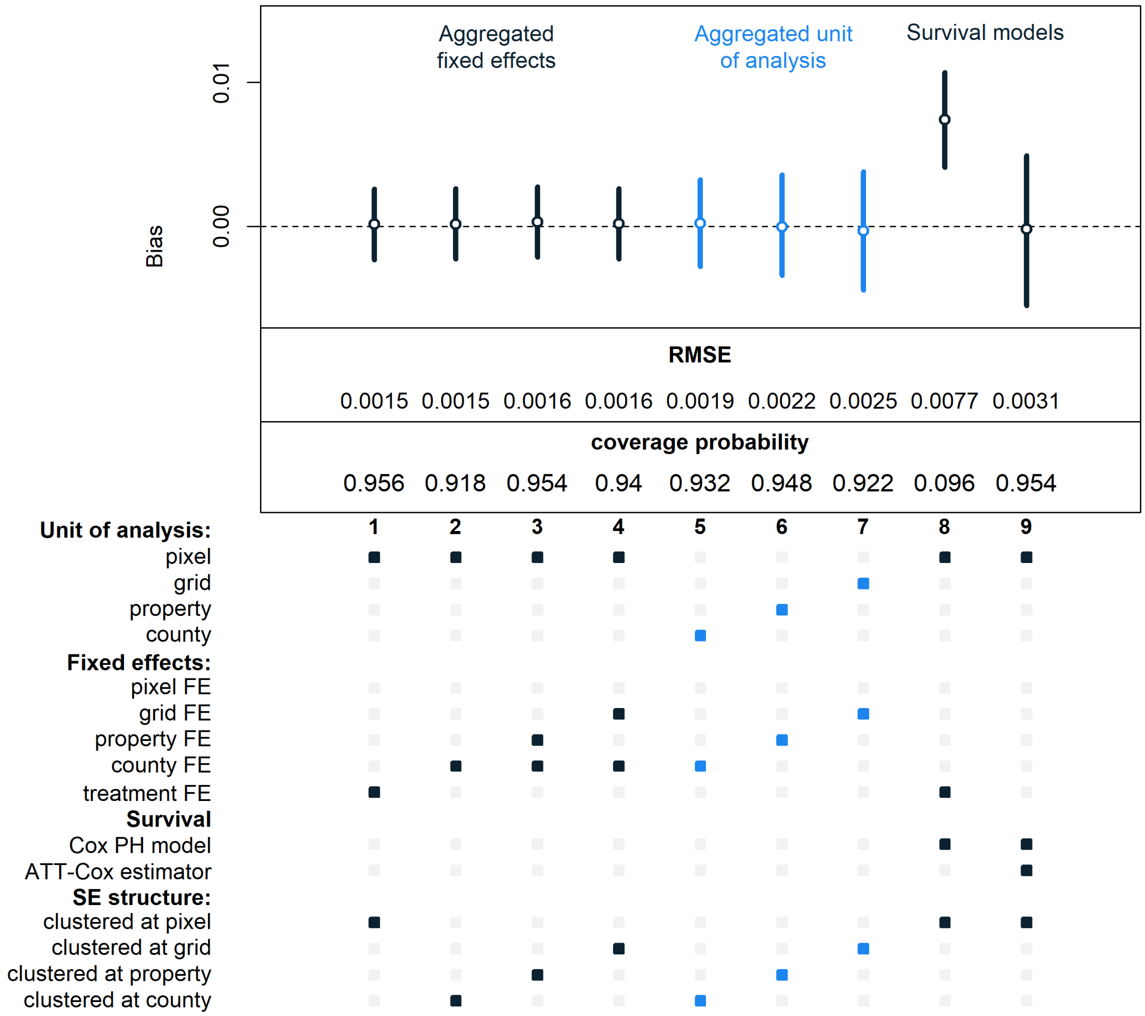

Figure 1.3: This figure shows the mean bias along with the 0.05 to 0.95 quantile range for estimates from each of our candidate models. RMSE and coverage of candidate models with clustered standard errors are shown below. Candidate models are separated by color according to whether they incoporate aggregated fixed effects in pixel-level specifications, aggregated units of analysis, or survival analysis.

1.5.5 Non-random sample selection

We now explore how the non-random sample selection described in section 1.4.4 may influence estimates in our simulated landscapes. Non-random sample selection did not bias our initial simulations as presented in Figure 1.3 because we assumed away time-invariant pixel and property-level disturbances. The sample of at risk pixels in each time period did not depend on the deforestation that occurred the previous period, since the y^o_{ivt}s were i.i.d.. However, once time-invariant disturbances enter the DGP, the distribution of the $y^o_{ivt}s$ is potentially different in each year of the panel. This is likely to be the case in reality, since each plot of land will have time-invariant characteristics such as slope, elevation, market access, agricultural suitability, etc. that impact its expected returns to clearing.

In order to see how this non-random selection influences estimates in our simulated setting, we set σ_a, the standard error of the time-invariant pixel-level disturbances equal to 0.1. Figure 1.4 shows that non-random selection introduces a slight downward bias across every specification.

In practice, researchers cannot recover the second and fourth terms of the bias term in equation 20, meaning that the magnitude of this bias is unknown to the researcher. However, this bias is likely to be of a smaller magnitude than in our simulated setting if deforestation rates are lower or more similar across treated and control groups.

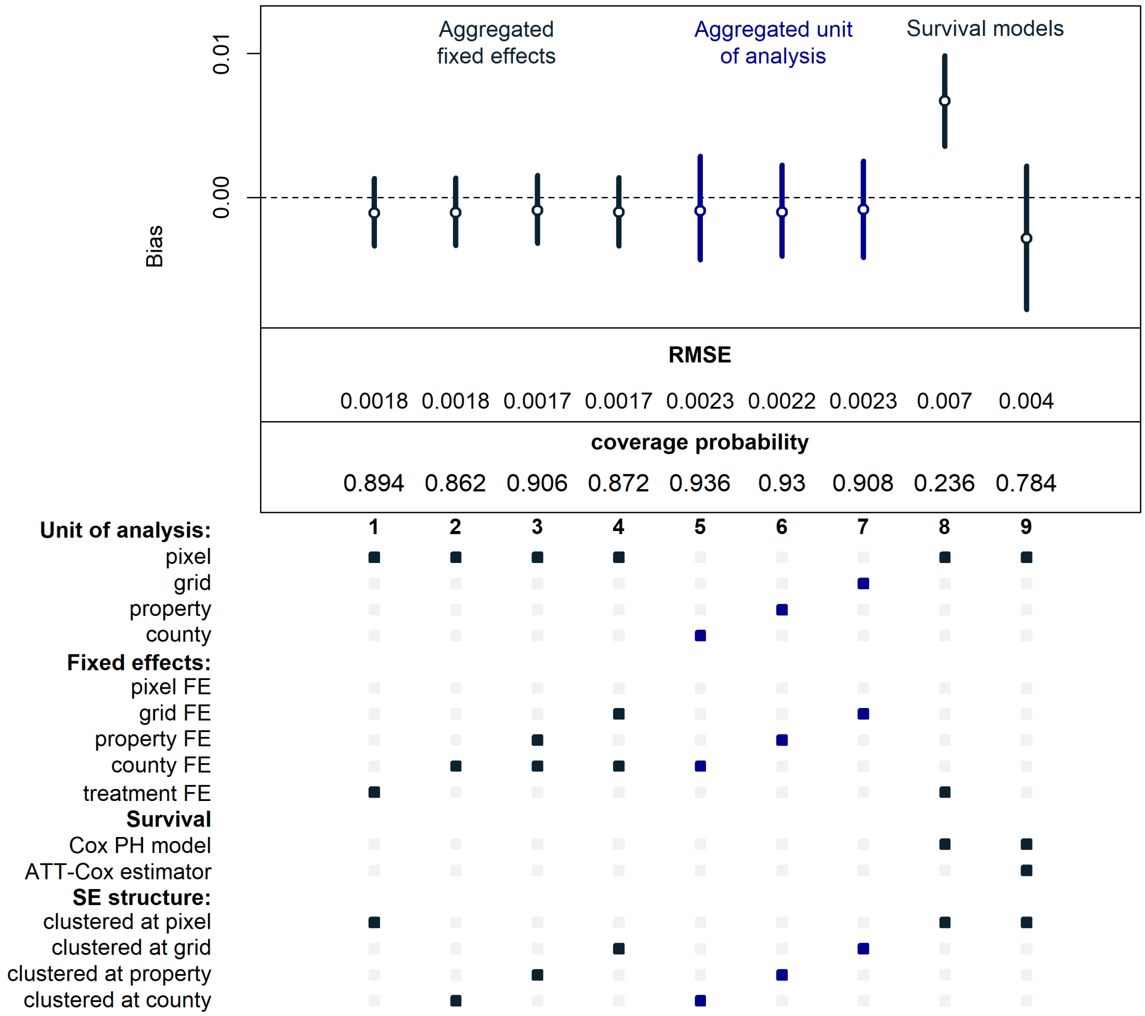

Figure 1.4: This figure shows the mean bias along with the 0.05 to 0.95 quantile range for estimates from each of our candidate models after allowing for non-i.i.d. pixel-level outcomes through time. RMSE and coverage of candidate models with clustered standard errors are shown below. Candidate models are separated by color according to whether they incoporate aggregated fixed effects in pixel-level specifications, aggregated units of analysis, or survival analysis.

1.5.6 Selecting the appropriate spatial structure

Model structures that match the spatial process of deforestation can reduce bias

Connecting the econometric model to the process by which land use change occurs on the ground has clear benefits for estimation and inference in deforestation impact evaluation. Table 1.1 shows that researchers often use an arbitrary spatial unit such as a point, pixel, or grid cell as the unit of analysis. While this may be a useful way of structuring data, it can lead to biased results if land use change is determined through a process that is mediated by other spatial structures.

In reality, property level unobservables such as the preferences and resources of a landowner may drive significant variation in land use across a landscape. These differences will impact both treatment effect estimates and coverage probabilities. To illustrate this effect, we introduce property-level perturbations to the returns from forest clearing by varying σ_p, the standard deviation of time-invariant property level disturbances in the DGP.

The introduction of σ_p changes the relative performance of each specification. The traditional DID does not account for the spatial nature of the deforestation process, and in Figure 1.5, we see that the pixel-level DID begins to suffer in terms of bias, RMSE, and coverage as these property-level unobservables play a larger role in the data generating process.

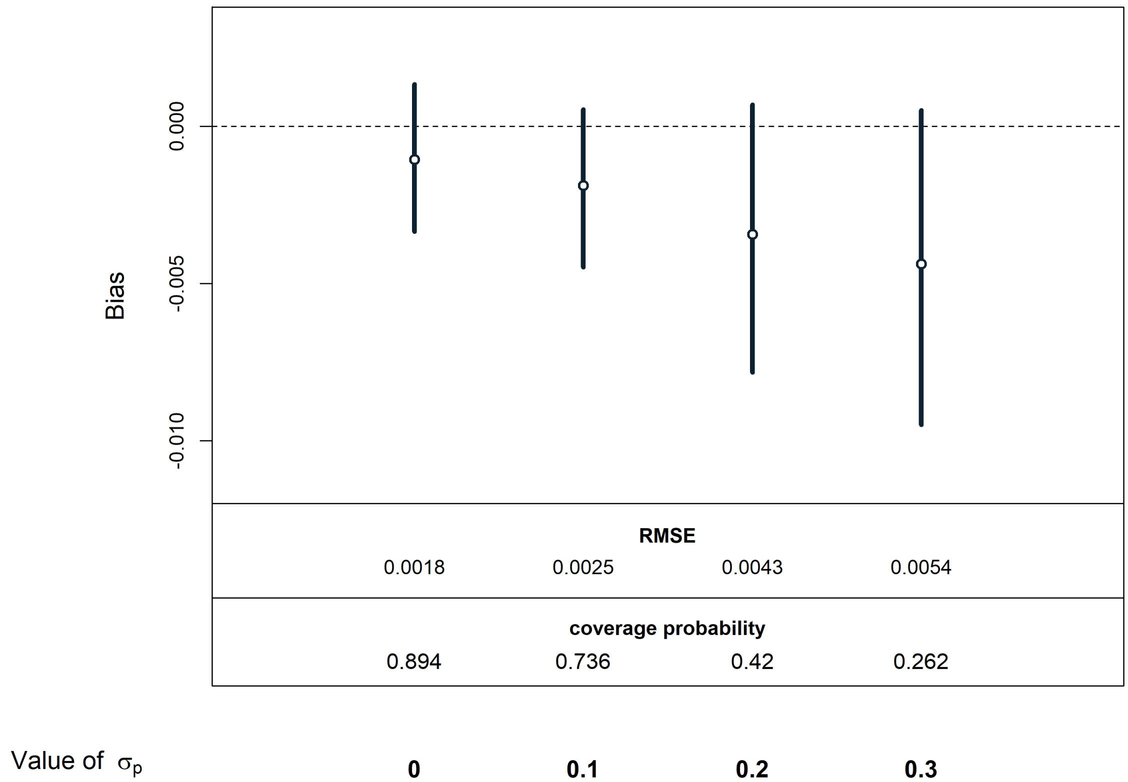

Figure 1.5: This figure shows how the mean bias and 0.05 to 0.95 quantile range from estimates based on the pixel-level DID model change as the relative scale of property level disturbances increases.

In Figure 1.6 we see that, by incorporating spatially aggregated units into the model structure, the researcher can reduce bias relative to the simple pixel-level DID in settings where property-level perturbations are relatively large ($\sigma_p = 0.3$). This improvement is apparent across specifications that either control for spatially aggregated fixed effects (Regression 3; left panel) or use a spatially aggregated unit of analysis (Regression 4; right panel). That said, specifications with a spatially aggregated unit of analysis consistently outperform their counterpart with spatially aggregated fixed effects. Further, we see that the scale of aggregation plays a role. Analyses incorporating the property or grid cells near the size of the average property outperform models using larger or smaller scales.

Ultimately, all models that incorporate spatial aggregation suffer from relatively little bias and incur less bias than the simple pixel-level DID. In Appendix A.1.9, we see that specifications incorporating the property as the unit of analysis continue to outperform other models in alternate parameterizations.

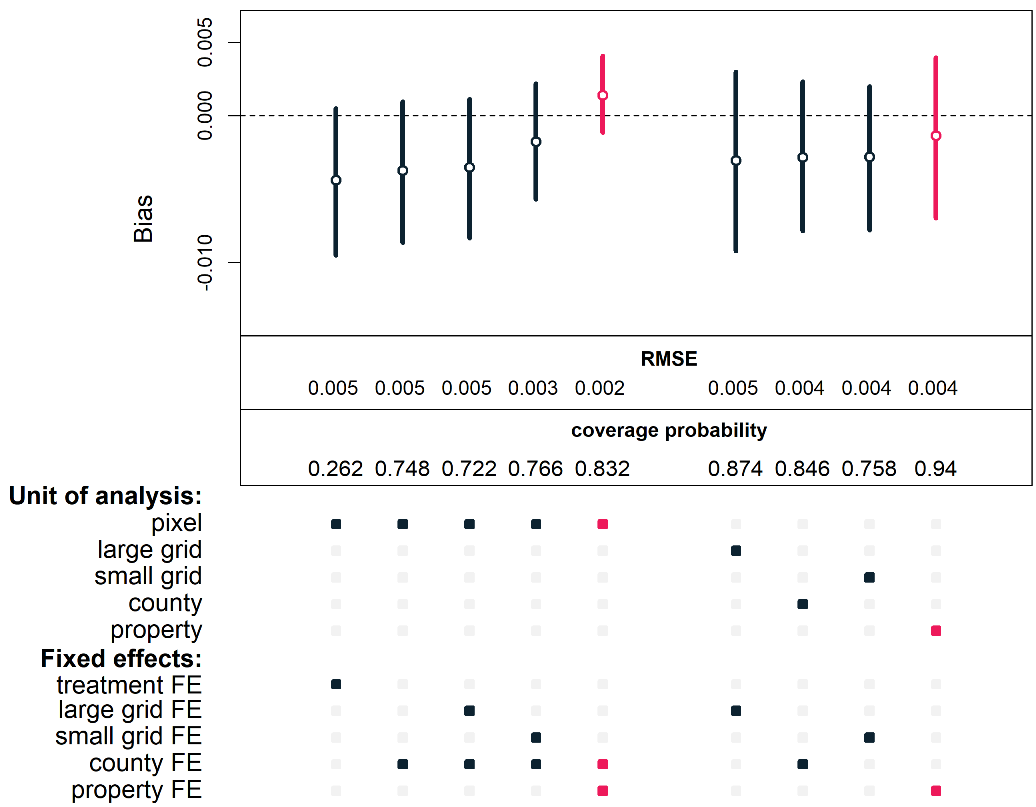

Figure 1.6: This figure shows the mean bias and 0.05 to 0.95 quantile range for specifications with aggregated unit fixed effects (left panel), and specifications with aggregated units of analysis (right panel) when $\sigma_p = 0.3$.

Although spatial aggregation can improve the performance of OLS-based model specifications, there is no clear analog for survival models. Figure A.1 in the appendix shows that the performance of $\widehat{ATT}_{Cox}$ suffers as σ_p increases. When unobserved, spatial processes contribute to the underlying DGP, linear models that effectively control for these processes are likely to outperform survival analysis-based estimates of the impact of conservation interventions. Survival analysis-based estimators may perform better when the preferred unit of analysis is actually the individual point (i.e., mortality, recidivism).

1.5.7 Weighting by area recovers landscape scale estimates

As researchers transition towards spatially aggregated units of analysis, interpretation of the estimated ATT can become more complicated. Authors frequently choose to use a set of evenly-sized pixels or grid cells as their preferred units of analysis in order to simplify the interpretation of their estimated ATT (Alix-Garcia and Gibbs, 2017). For example, when researchers estimate a model with pixel-level units of analysis, the coefficient of interest can be interpreted as a population average for all treated, forested pixels. In contrast, if a property is used as the unit of analysis, the coefficient should be interpreted as the effect of the intervention on the characteristic property in the sample. In order to obtain a landscape-scale interpretation, one must weight the regression by the area of each unit of analysis (i.e. property).

Weighting does not have a large impact on bias, RMSE, or coverage probability when the treatment effect is constant across properties (even with property-level unobservables). The use of area weights is likely to be most useful when the treatment effect in the characteristic property differs from the landscape's ATT. To illustrate this effect, we consider a landscape in which treatment effects are correlated with property size. The full DGP for this case can be found in Appendix A.1.7.

The treatment effect now varies across properties, and properties with greater areas experience treatment effects of a lower magnitude than smaller properties. For clarity of definitions, we assign treatment at the property level in this subsection. We consider two sample ATTs: the landscape ATT and the property-level ATT. They can be defined as follows:

* $ATT_{ls} = \frac{1}{n_{i:D_i=1}} \sum_{i:D_i=1}(y_{iv2}(1) - y_{iv2}(0))$, where $n_{i:D_i=1}$ is the number of treated pixels in the simulated landscape; and

* $ATT_{property} = \frac{1}{n_{v:D_v=1}} \sum_{v:D_v=1}(\frac{1}{n_{iv}} \sum_{i=1}^{n_{iv}}(y_{iv2}(1) - y_{iv2}(0)))$, where $n_{v:D_v=1}$ is the number of treated properties in the simulated landscape; and n_{iv} is the number of pixels in property v.

Note that neither ATT_{ls} nor $ATT_{property}$ can be calculated directly, because $y(0)_{iv2}$ is not observable for treated units.

Because the treatment is more effective in properties of a smaller size, the treatment effect for the average property is greater than the average treatment effect experienced across the landscape. Figure 1.7 shows the sample ATTs for both the property and landscape. In our simulation, $ATT_{property} = -0.0136$, and $ATT_{ls} = -0.0092$. The property-level TWFE regression recovers the ATT relative to the characteristic property when area weights are not used and the landscape scale ATT when they are used. Researchers should use these area weights when they are interested in the impact of the intervention across the landscape. In cases where the researcher is interested in how an intervention affects incentives at the property level, using these weights may not be necessary.

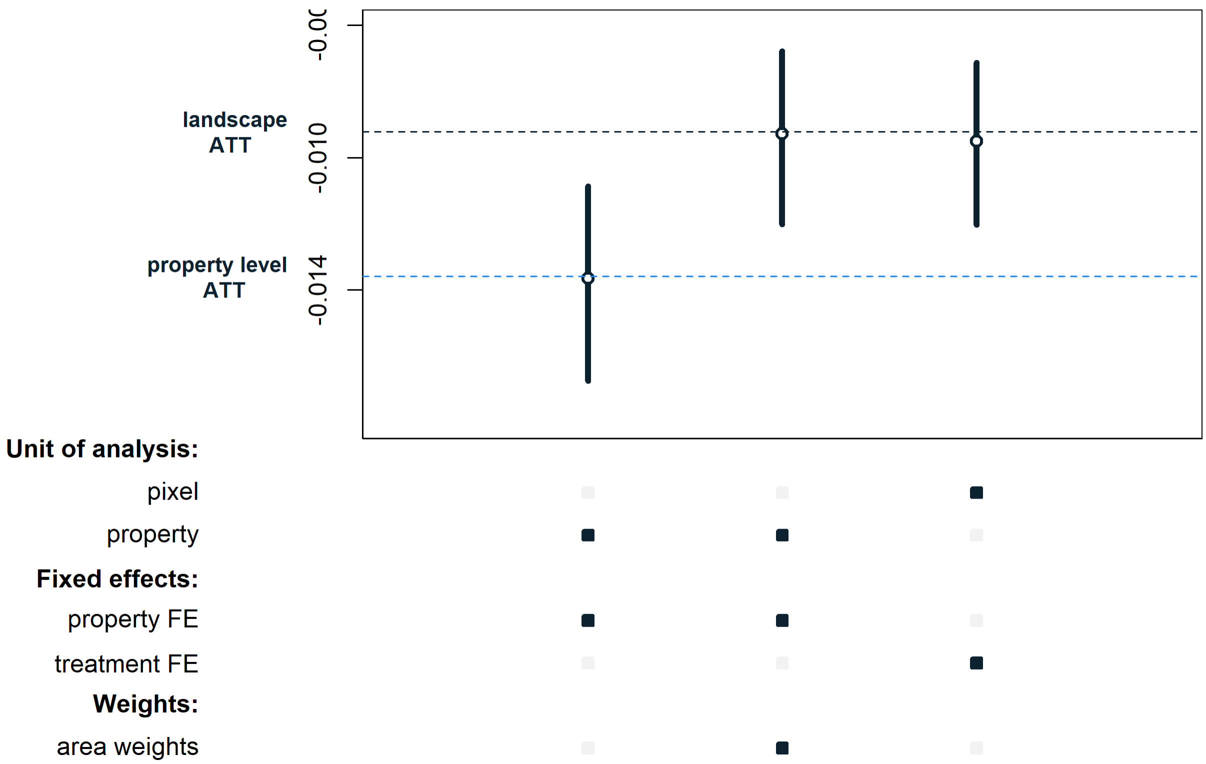

Figure 1.7: Weighting recovers landscape scale interpretation.

1.6 Estimating the ATT under staggered treatment

1.6.1 Staggered setup

The traditional DID regression applies to settings with two groups and two time periods. However, researchers often use TWFE regressions to exploit variation across groups of units that receive treatment at different times. Recent work has shown that, in these staggered treatment settings, TWFE regressions identify a weighted average of all possible two-group/two-period DID estimators in the data (Goodman-Bacon, 2021).

Further, when estimating the ATT, some weights on each group-time treatment effect parameter may actually be negative (de Chaisemartin and D'Haultfœuille, 2020). Newly developed DID estimators seek to produce unbiased estimates of the ATT in settings with multiple groups and time periods. These estimators do so through a variety of strategies including imputation (Borusyak et al., 2021), two-stage least squares (Gardner, 2021), and the re-weighting of group-time ATTs (Callaway and Sant'Anna, 2020). Some researchers might hope that these new estimators would solve the bias detailed in Section 1.4.1.

1.6.2 New DID estimators are biased when applied to binary, irreversible outcomes

Although the new class of DID estimators effectively address concerns about staggered treatment timing and heterogeneous treatment effects, they continue to yield biased treatment effect estimates when applied to binary, irreversible, outcomes. To illustrate this, we introduce a setting in which groups of units receive treatment at different times (full DGP can be found in Appendix A.1.10). We consider three groups: an early group, a late group, and a never-treated group, where the early and late groups undergo treatment in years three and four, respectively. Each group experiences differing pre-treatment deforestation rates (7%, 4%, and 2% for the early, late, and never-treated groups, respectively) and no time trend. The ATT for both treated groups is -0.02. Common trends is satisfied by construction, and we do not introduce any dynamic effects. Figure 1.8 shows the observed deforestation rates ($E[y^o_{ivt}]$) from one iteration of our simulation in this setting.

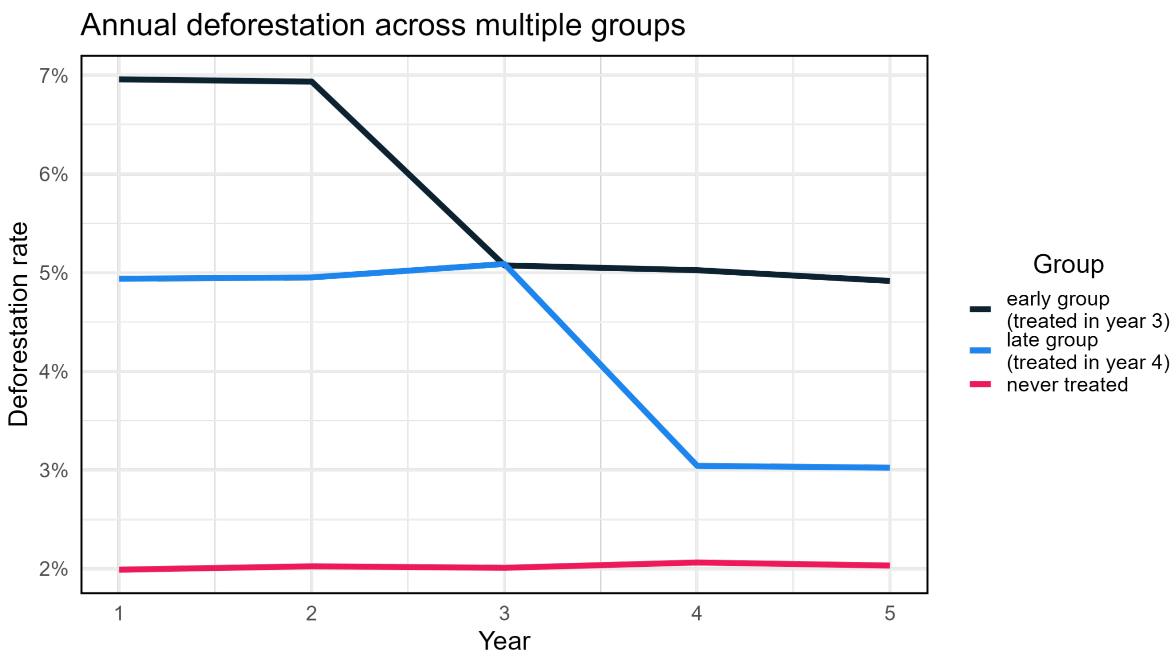

Figure 1.8: Observed deforestation in simulated example with multiple groups and periods.

The left panel of Figure 1.9 shows that the estimators developed in Callaway and Sant'Anna (2020), Gardner (2021), and Borusyak et al. (2021) suffer from similar bias to TWFE regressions with pixel unit fixed effects if the pixel is used as the unit of analysis. All methods yield a treatment effect greater than or equal to 0 in all post-treatment periods, reflecting the fact that pre-treatment period deforestation rates are unaccounted for by the estimators. This is particularly clear in the Callaway and Sant'Anna (2020) estimator in which pre-treatment periods are all precisely zero, indicating that the estimator could only compute treatment effects using pixels that survived until the end of the observation period. The right panel of Figure 1.9 shows that this bias is eliminated when one uses an aggregated unit of analysis with binary treatment (e.g., county). We do not include pixel-level TWFE regressions with spatially aggregated fixed effects, because most recently developed estimators do not allow for a comparable implementation. Therefore if the researcher wants to obtain a landscape-scale interpretation, they must aggregate to a uniform spatial area (e.g. grid cell) or weight the regression by the unit area.

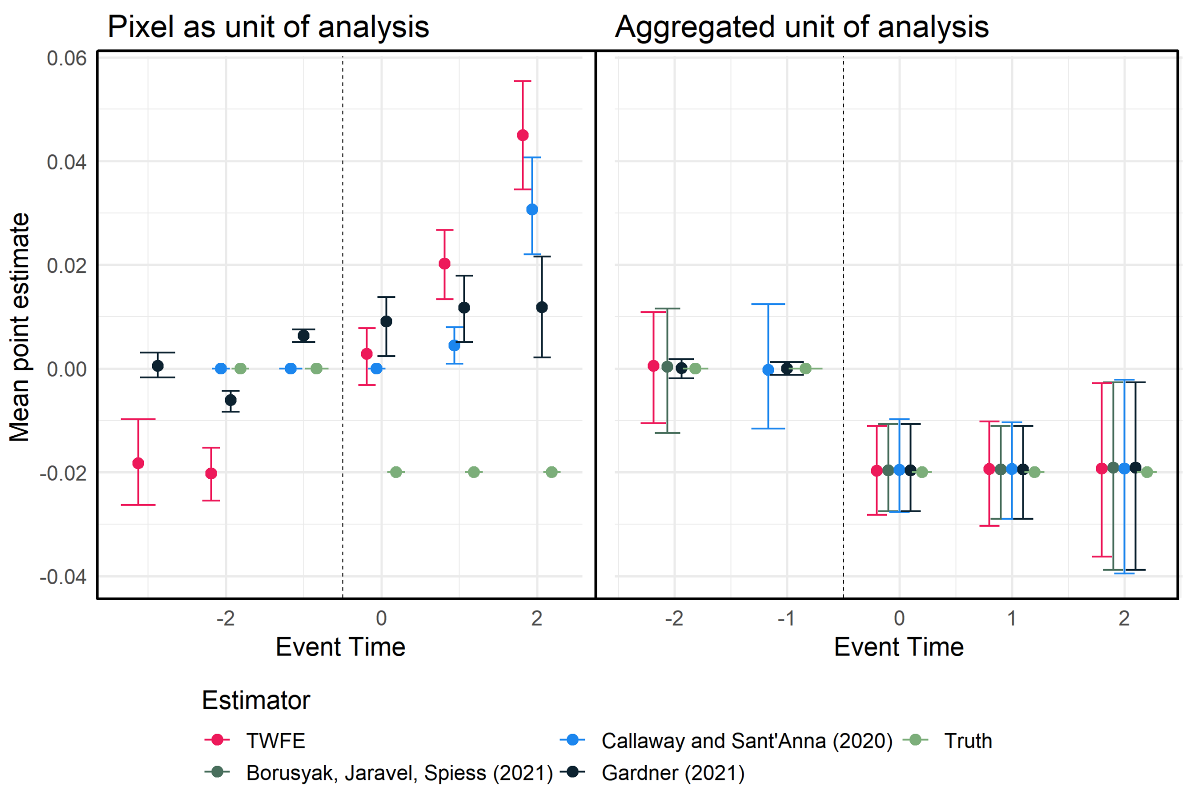

Figure 1.9: New estimators, similar to TWFE regressions with pixel unit fixed effects, cannot identify ATT with pixel as unit of analysis.

1.6.3 New DID estimators can yield unbiased estimates of heterogeneous treatment effects

Finally, we examine the performance of the new DID estimators relative to a traditional TWFE regression when treatment effects vary across time and across groups. We again work with an early, late and untreated group. The full parameterization and DGP can be found in Appendix A.1.10. Figure 1.10 shows deforestation rates in each of the three groups through time.

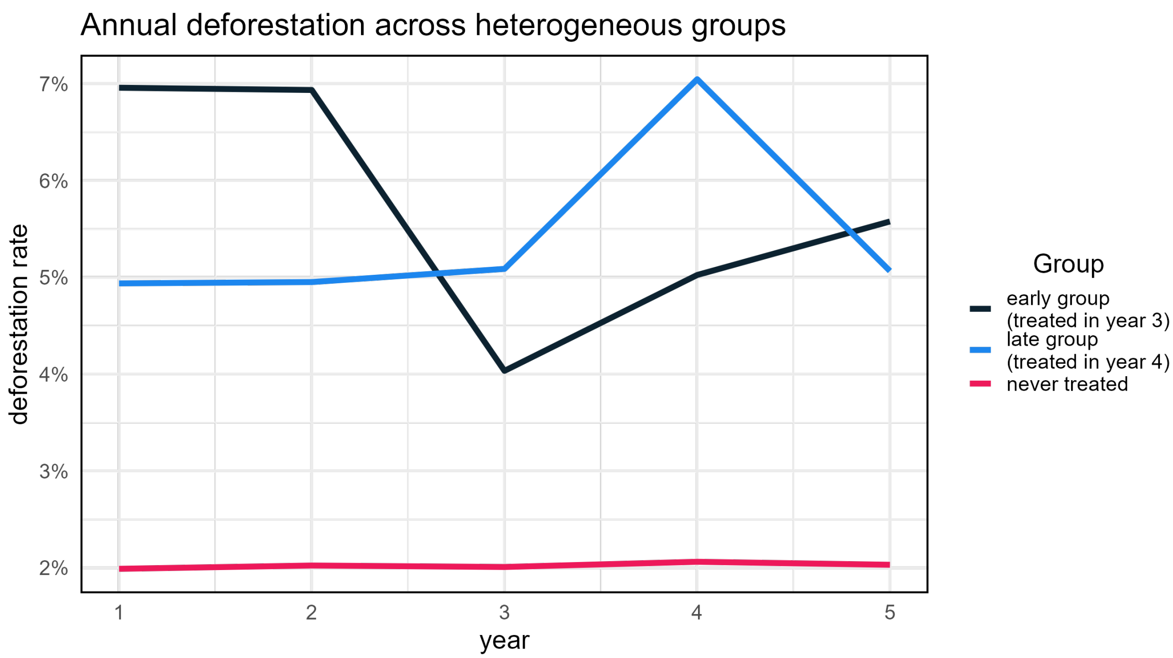

Figure 1.10: Observed deforestation in simulated example when treatment effects vary across groups and through time.

Figure 1.11 shows the event study estimates produced by each of the three estimators as well as the “truth" for both pixel and county-level analyses. Again, none of the estimators yield the ATT with pixel-level analyses. In the county-level estimates, we see that the newer estimators slightly outperform the TWFE estimator. This is evidence of the weighting that has become a concern with TWFE estimators in these type of settings. While TWFE estimates represent a weighted average of all possible 2x2 DID estimates, the weights may not always be intuitive (Goodman-Bacon, 2021). In contrast, newer estimators do not suffer from this concern.

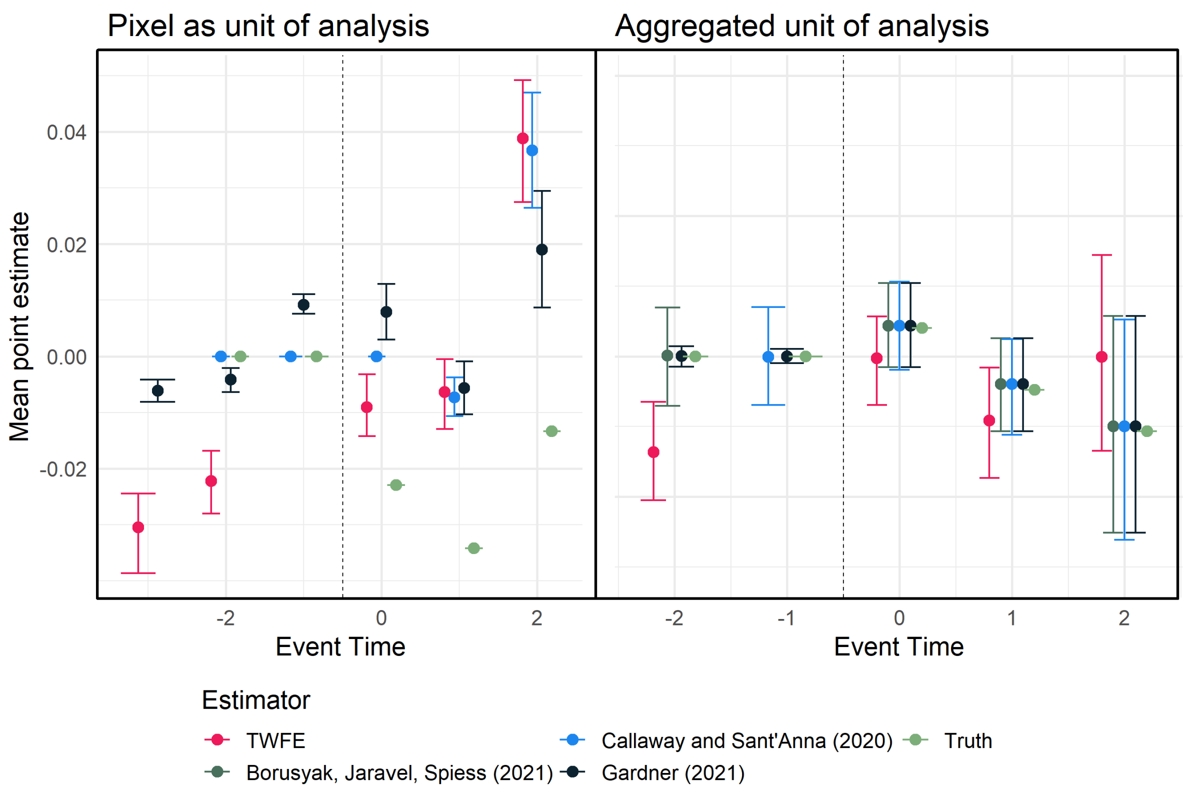

Figure 1.11: TWFE regressions suffer from weighting concerns when treatment effects vary across groups and through time.

1.7 Conclusions

By applying econometric methods of causal inference to remotely-sensed measurements of land use change, researchers have advanced society's understanding of the impacts of conservation interventions. However, this interdisciplinary research community has insufficiently considered how the data generating processes underpinning land use change and its measurement might affect the performance of standard econometric models. The analytical proofs and simulations presented in this paper highlight that the conclusions made in many prior studies may be biased.

Researchers can take several practical steps in the design of their econometric models to more accurately measure the impacts of conservation policies. First, despite past guidance to the contrary, researchers should recognize that pixel-level, TWFE models are unable to yield unbiased estimates of a policy's impact when applied to irreversible, binary outcomes. Researchers can easily avoid this bias by aggregating either the units of observation, or the scale at which fixed effects are estimated. Second, while survival models provide an appealing empirical framework with which to study deforestation, past studies have typically overlooked implicit assumptions made when applying survival models to the difference-in-differences research design. To resolve this challenge, we propose a new, survival-based estimation procedure that enables researchers to recover an unbiased estimate of the ATT under the traditional parallel trends assumption. Because this estimator underperformed when heterogeneity entered the landscape at a scale different than the unit of analysis, it is likely to be most useful when the researcher can justify the individual unit as the desired unit of analysis. Finally, we provide evidence suggesting that researchers should seek to align the structure of their econometric models to match the real-world units at which land use decisions are being made. For example, if unobservable, property-level characteristics are thought to be an important driver of deforestation, the inclusion of property-level fixed effects can improve the accuracy of model estimates and inference. Ultimately, context plays a role in what is feasible, and researchers should make clear the limits to their impact evaluation strategy.

This paper complements an emerging literature calling for a deeper understanding of the interdependencies between the creation of remotely sensed data, and the interpretation of that data through econometric models (Jain, 2020; Alix-Garcia and Millimet, 2022; Proctor et al., 2023). However, we have largely abstracted away from prior concerns that characteristics of the remote sensing data collection process, including sensor properties, atmospheric conditions, and image processing methods, may influence the structure of output data products. Of particular concern is the potential for these processes to give rise to non-classical measurement error, which can lead to biased estimates of the ATT (Wooldridge, 2010). Importantly, our study implicitly assumes that pixel-level outcomes are measured without any non-classical measurement error.

This paper focuses upon efforts to identify the impacts of conservation policies on deforestation. However, the lessons we highlight are relevant to a wider audience, and many of our key findings apply to diverse settings in which the outcome of interest represents an irreversible, binary event. For example, studies exploring the drivers of unidirectional technology adoption or recidivism may fall prey to the same issues we have identified in the context of deforestation. Moving forward, researchers should carefully consider the underlying structure of their data, and ensure that their chosen models minimize bias and allow inference at expected levels of confidence. Misleading causal inference may lead policymakers to avoid effective policies, or to adopt interventions that worsen environmental damages.

Chapter 2

Targeting in payments for forest restoration: evidence from Chile's Native Forest Law

2.1 Introduction

In order to achieve the warming targets set by the IPCC, both emissions reductions and removals of carbon from the atmosphere will be necessary (Masson-Delmotte et al., 2018). Reforestation and forest restoration have been lauded as potentially near term, large scale and low-cost options to achieve these carbon removals (Busch et al., 2019). Global initiatives such as the Bonn Challenge, Trillion Trees Campaign, and UN Decade on Ecosystem Restoration highlight the enthusiasm for forest restoration as a climate solution (Chazdon and Brancalion, 2019). However, there is limited evidence documenting the effectiveness of policies seeking to encourage restoration at large scales. Payments for ecosystem services (PES), which pay private landowners in exchange for providing environmental benefits, have been used extensively in the context of avoided deforestation (Börner et al., 2017), and are likely to play a major role as countries begin to scale up forest restoration in line with global commitments (Gichuki et al., 2019).

Political processes often lead policymakers to design programs aimed at multiple objectives (Tinbergen, 1952). PES programs are no different, and are frequently designed to achieve social development goals in addition to environmental gains (Zilberman et al., 2008; Lipper et al., 2009; Pagiola et al., 2008). This is done by targeting payments toward landowners or communities with the greatest need for poverty alleviation (Wunder, 2008). However, subsidies that reflect political processes have the potential to undermine environmental benefits (Jack et al., 2008), and this strategy of targeting for social development has proven to create tradeoffs in payments for avoided deforestation programs (Alix-Garcia et al., 2015).

The effectiveness of payments for avoided deforestation is driven by the presence of deforestation risk (Alix-Garcia and Wolff, 2014), but the drivers of effectiveness in payments for restoration is less clear. In the deforestation context, targeting payments toward households unlikely to deforest in the absence of payments generates relatively little additional forest cover relative to the no-payment scenario (Pfaff, 1999; Cisneros et al., 2022). It is well documented that poverty is associated with lower levels of deforestation (Busch and Ferretti-Gallon, 2017), and as a result of this broad finding, it has been shown that enrollment of landowners in high-poverty areas generates relatively fewer environmental benefits (Pfaff et al., 2007). Further, credit constrained and low-income landowners may be prone to leakage, meaning that they deforest unenrolled land upon enrolling part of the property (Jack and Cardona Santos, 2017; Alix-Garcia et al., 2012).

While evidence detailing the pitfalls of targeting for "win-wins" in payments for avoided deforestation is widespread, there is a need for rigorous evidence in the restoration setting. This lack of evidence is important because theory indicates that targeting within the context of payments for reforestation might have fundamentally different impacts than in the context of payment for avoided deforestation. Private information about the costs to providing forest often make targeting difficult (Jack, 2013), and the underlying correlation between targeted characteristics and costs of producing forest may be different in the deforestation and restoration cases. For example, relatively poor households may be unlikely to deforest in the absence of payments, but it may also be infeasible for them to produce new forest without payments due to credit constraints or lack of technical capacity. In this scenario, payments may not generate any behavior change in the deforestation case, but lead landowners to begin providing forest in the restoration case. Second, non-compliance often plays a role in the restoration context. Landowners may enroll in a PES program but fail to complete contracted activities if, for example, landowners are uncertain about restoration costs and benefits (Oliva et al., 2020). If targeted characteristics are positively correlated with non-compliance, tension may arise between participation of priority groups and program effectiveness.

Chile's Native Forest Law provides us a unique setting in which to explore the potential of payments for restoration to achieve low-cost carbon removals and whether targeting for social development undermined environmental effectiveness. The Native Forest Law provided subsidies for native forest restoration and explicitly sought co-benefits such as economic development of smallholder, indigenous, and rural communities. Program administrators did so through an annual competition for subsidies to support new projects aimed at either restoring existing native forest or planting new forests. Applications to the competition were split into two separate contests based on property size and landowner assets. They were then explicitly scored to determine priority, where the score is based not only on project-specific characteristics, but also social characteristics that we show to have a clear positive association with poverty. Few PES programs explicitly score applications by social characteristics, giving us a unique look at the explicit targeting of certain priority characteristics.

In this paper, we use novel data to quantify the land cover impacts of Chile's Native Forest Law for ten cohorts (2009–2018) using annual data from 2000 to 2018. Temporally consistent annual landcover data detailing tree cover and other land classes such as grasslands and croplands provide a unique opportunity to evaluate restoration outcomes using panel difference-in-differences approaches. Using the estimates of program impact, we calculate the carbon impacts of the policy and the price paid to the average landowner per tonne of CO2 stored in aboveground biomass through native forest restoration. We then evaluate the Native Forest Law's strategy to prioritize participation of marginalized landowners by examining treatment effect heterogeneity across contest, social score, and poverty.

A key concern with evaluations of programs with voluntary participation is that apparent effectiveness may actually be due to different participation costs (Jack and Jayachandran, 2019). The typical enrollee to the Native Forest Law competition was already increasing forest cover prior to enrollment in the program. If we were to simply compare enrolled and unenrolled properties after enrollment, we would likely overestimate the effect of the subsidies, since enrollees were producing new forest cover even without payments. In order to alleviate concerns surrounding selection, we use pre-processing techniques to construct a set of control properties from a pool of all rural properties in the major forested regions of Chile. These matched control properties are much more comparable to enrollees than the typical unenrolled property based on a detailed set of land use and property characteristics. These characteristics include pre-program land-uses, which are likely to drive enrollment decisions. Difference-in-differences methods further allow us to control for fixed differences in land cover between enrollees and the control group as well as time trends affecting both groups. Using this strategy to estimate treatment effects for properties that dropped out after the first six months and never received payment yields no statistically or economically meaningful treatment effect, lending creedence to our estimates for program beneficiaries.

We find that the program increased tree cover on the properties of landowners who received payment, coming largely from land that was previously cropland or grassland. Subsidization of the typical property through the Native Forest Law contest led to a 0.21% increase in the share of the property covered in forest. This estimate suggests that the average landowner was paid an estimated $36.78 USD per tonne CO2 stored in aboveground biomass through the subsidy competition. Event study estimates based on weaker identification assumptions show that dynamic effects play a significant role and that the program led to a more substantial increase in tree cover for cohorts observed for several years post-enrollment. Based on only the earliest cohorts, enrollment ultimately increased the share of the typical property covered in forest by 1.00%. Based on this fact and several other moderating factors, we feel comfortable interpreting the estimated $36.78 USD per tonne CO2 as an upper bound estimate. On one hand, this estimated price assumes that newly established forest matures. On the other hand, this price ignores the fact that recent cohorts may continue to provide new forest in a similar fashion to the earliest cohorts. Moreover, subsidized restoration in already standing native forest is not likely to be reflected in our satellite-derived measures of tree cover extent, and environmental benefits may be understated if landowners substitute away from plantation forest into native forest, which is not observable in our study.

Non-compliance in the Native Forest Law was high, as over two-thirds of applicants never received payment. Landowners with higher application social scores complied at relatively lower rates, however, program administrators avoided unconditional payments by requiring verification of activity completion. This allowed the program to target priority characteristics without increasing the probability of unconditional cash transfers. That said, recent work by Jack et al. (2022) shows that conditional payments in PES can reduce compliance, meaning that the conditional nature of the Native Forest Law payments may have had tradeoffs.

We find that smallholders drive the main treatment effect, while larger, wealthier landowners in the other interested party contest provided no positive tree cover benefits at all. Increased community-level poverty was associated with greater tree cover impacts, for which the program's scoring mechanism was a good predictor. This implies that the strategy to prioritize the participation of smallholders and those with priority social characteristics improved the program's environmental effectiveness. These findings are in contrast to the expectation in payments for avoided deforestation, where wealthier landowners with high deforestation risk often provide the greatest additionality.

Contributions. Our results make several contributions to the literature. First, we add to the limited existing evidence on the environmental effectiveness of large-scale native forest restoration programs. Between 2009 and 2019, the National Forest Corporation (CONAF) allocated approximately US $58 million to enroll more than 235 thousand hectares of land through the Native Forest Law, making it one of the largest native forest restoration programs in the world. There are relatively few studies evaluating payments for forest restoration and even fewer quantifying the impacts of large-scale or national policies. Notable exceptions include Uchida et al. (2005), who examine land enrolled in China's Grain for Green program. While unable to examine land cover outcomes explicitly, they find that payments were relatively well targeted toward plots with low costs of conversion and high potential for erosion reduction. Heilmayr et al. (2020a) use econometric simulation to evaluate the land cover impacts of Chile's Decree Law 701, which subsidized monoculture plantations of exotic eucalyptus and pine. They find that additionality was quite low, and further, that native forest extent was reduced as landowners substituted toward plantation forest. España et al. (2022) is one of the few studies of which we are aware that quantifies land cover impacts of a national-scale afforestation or reforestation policy using quasiexperimental methods. They quantify the effects of Chile's Decree Law 701 using difference-in-difference methods and find that the policy led to a 13% increase in forest plantations amongst subsidy participants. To our knowledge, our paper is the first to provide an ex-post, causal assessment of the land cover impacts of a national payment for native forest restoration program.

Our second contribution is to highlight key differences between targeting in payments for reforestation in contrast to payments for avoided deforestation. Several studies find that targeting for social development has come at the cost of environmental gains in payments for avoided deforestation. Alix-Garcia et al. (2015) find that the landowners for which PES most effectively reduced poverty provided the least avoided deforestation. In another study, Alix-Garcia et al. (2012) identify leakage amongst capital-constrained landowners in Mexico's National PES program, reducing program benefits in high-poverty areas. No study to our knowledge has addressed whether targeting for social development is costly in the reforestation or restoration setting. In this study, we identify a potential opportunity for win-wins in the restoration context. Smallholders in high-poverty areas actually yielded greater tree cover gains than larger, wealthier landowners. Descriptive evidence suggests that this may be due to the fact that wealthier landowners were already expanding tree cover at relatively higher rates prior to program enrollment. This highlights that poor landowners may experience relatively greater obstacles to producing new forest in the restoration setting. Although prioritized characteristics were associated with increased non-compliance, non-complying landowners were never paid, avoiding unconditional cash transfers.

2.2 Background: Chile's Native Forest Law

Chile's long history of public policies supporting tree cover expansion provides an incredibly useful setting in which to measure the impacts of payments for restoration. Chile's decree law no. 701 (DL 701) is one of the world's longest operating afforestation subsidies, but mainly promoted even-aged mococulture plantations of eucalyptus and pine that had negative effects on biodiversity and native forest extent (Heilmayr et al., 2020a).

In an attempt to encourage the recovery and protection of native forests, Chile sought to pass the Ley de Recuperación del Bosque Nativo y Fomento Forestal (Native Forest Law) as a successor to DL701 (Clapp, 1998). Initially expected in 1994, it became frozen in legislature before finally passing in 2008. In addition to protections for native forests, the law established an annual competition for grants to support private landowners in their efforts to manage, restore, or reforest their land using native species. Between 2009 and 2019, more than $58 million were allocated through these competitions for projects covering 235 thousand hectares. Much of this allocated funding has not been paid to landowners, however, as program follow-through is low.

The subsidy component of the law encourages three types of activities: 1) the regeneration, restoration or protection of native forests for conservation; 2) silvicultural activities aimed at restoring native forests for timber production purposes; and 3) silvicultural activities aimed at restoring native forests for non-timber production purposes. Of the 12,889 projects enrolled between 2009 and 2019, 10,912 (84.66%) restored native forest for the purposes of timber production. Examples of subsidized activities include thinning, enrichment planting (introduction of trees to degraded forest), and establishment of new forest via tree planting. Few estimates of the impacts of the Native Forest Law on land cover currently exist. CONAF estimated the carbon impacts of the Native

Forest Law through 2019 as it relates to Chile's Nationally Determined Contribution (NDC) as part of the Paris Agreement (CONAF, 2020). These estimates, however, assume that the carbon stored by every eligible subsidized hectare is the direct result of the law. These types of estimates ignore the concept of additionality, since some of this forest would likely exist even in the absence of the law.

The Native Forest Law prioritizes not only forest cover in line with Chile's NDC goal of managing 200,000 hectares of native forest, but also the participation of underrepresented groups. By prioritizing native forest rather than monocultures of pine or eucalyptus, the law seeks to incentivize the preservation of biological diversity. Prioritizing carbon-plantings without consideration of other co-benefits may result in negligible biodiversity co-benefits (Bryan et al., 2016). In fact, DL701 resulted in the decline of native forest and biodiversity, landowners replaced native forest with plantation forest, which provide significantly less biodiversity than native forests in Chile (Heilmayr et al., 2020a).

The Native Forest Law explicitly states that it seeks to meet the dual objectives of rural economic development in addition to protection of the environment. In the case of the Native Forest Law subsidy competition, this is done by prioritizing applications of landowners deemed to be of social importance. This is done in two main ways: 1) holding separate contests for smallholders and wealthier, larger properties; and 2) a scoring mechanism that gives higher scores to indigenous peoples, smaller properties, and those in rural regions. We further explore the ramifications of this targeting strategy in Section 2.5.

Upon having an application selected through the Native Forest Law's competition, landowners must clear two key administrative hurdles in order to receive payment. First, landowners must provide an updated native forest management plan, detailing the specific activities to be performed on the property, a timeline for activity completion, and a georeferrenced map of the property. Landowners generally have six months from the date of selection to provide the management plan or be dropped from the program. The second hurdle is to have the project's completion verified by a third party. The timeframe for activities generally lasts multiple years, and landowners are not eligible to receive payment until a minimum of two years after enrollment. Conditional on submitting the management plan, the rate of payment is relatively high at 75.44%, however, only 37.05% of enrollees submit the management plan within the required six month window. This means that ultimately, only 31.88% of projects enrolled between 2009 and 2019 have actually been paid out on. Program administrators have cited both credit constraints and a lack of labor supply as potential causes for such high rates of non-compliance. Submission of cartography with the management plan may also deter some landowners who struggle with the administrative hurdle itself.

2.3 Data and descriptive statistics

We have obtained property boundaries for all rural properties in the major forested regions of Chile as of the year 2009. Data on the awarded properties are available through CONAF and reflect aspects of the property and projects such as project objective, project surface area, bonus amount, and whether a management plan was submitted within six months. Also included is each property's parcel identifier, which is unique to each property within a comuna, Chile's level 3 administrative unit. We match the enrolled properties to their corresponding boundaries via this unique parcel identifier. In addition, we observe payment recipients, which are matched to the corresponding program application, indicating whether a project was successfully completed.

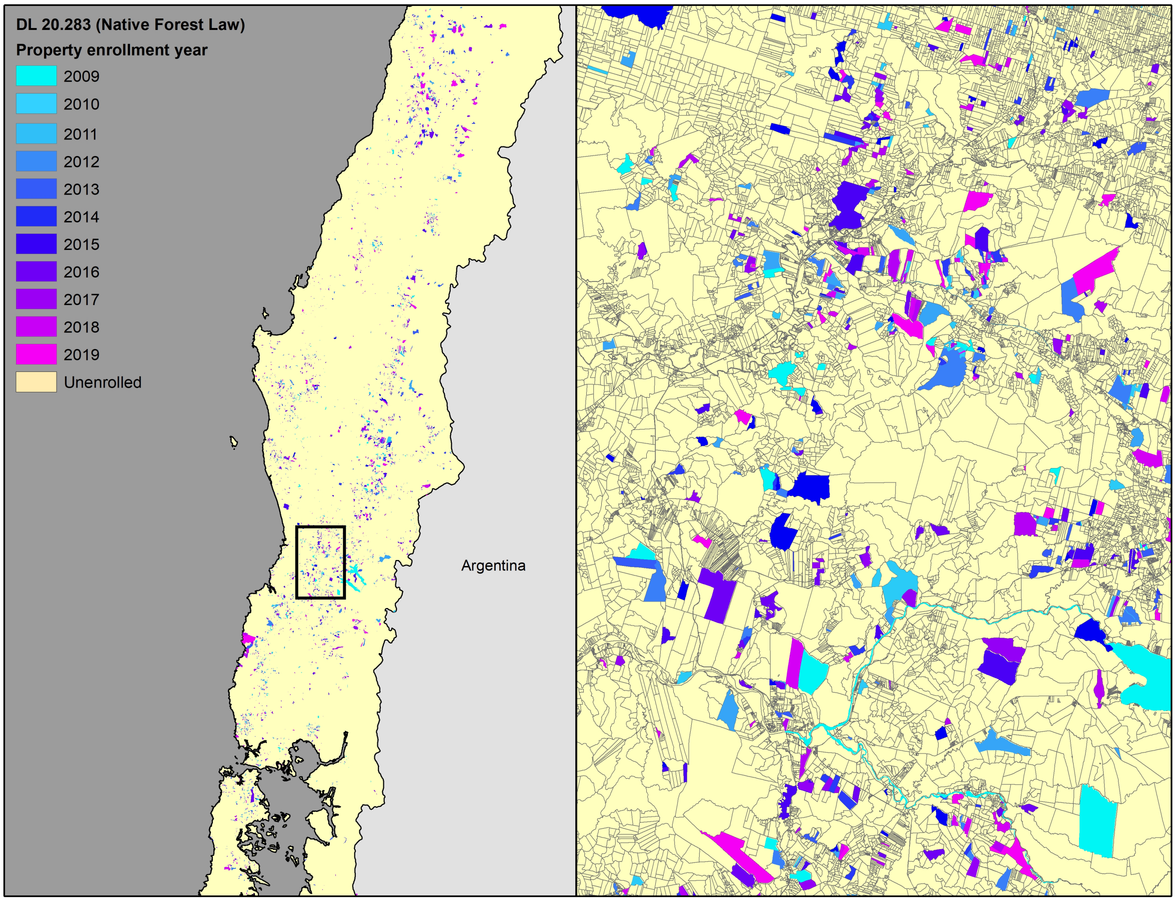

Figure 2.1: Properties enrolled through the Native Forest Law subsidy competition.

In order to quantify environmental impacts of the Native Forest Law, we use annual measures of land cover. Our primary outcomes of interest come from annual Landsat (30m) resolution maps of land cover developed in Graesser et al. (2022). These maps classify pixels into one of the following classes for each year between 2000 and 2018: forest, crop, grassland, shrub, and bareground. These land cover maps provide a unique opportunity to examine restoration outcomes. While many studies use satellite derived measures of deforestation to generate annual panels of forest loss, very few studies leverage annual variation in specific land cover types outside of those focused on North America and Europe. Further, the Graesser et al. (2022) product was developed specifically to produce more consistent estimates of land cover change over long time periods and gradual change events such as restoration, making it the ideal data product for our analyses.

Second, Landsat resolution land-use classification maps of these regions in Chile developed in Heilmayr et al. (2020a) allow us to distinguish the proportion of each property engaged in specific land uses prior to the existence of the Native Forest Law. Of particular interest is the distinction between plantation forest and native forest, which cannot be distinguished in the Graesser et al. (2022) product. In contrast to native forest, high levels of plantation forest may indicate greater ability to manage forest and undertake contracted activities. All of our satellite derived measures cover the extent of the major forested areas of Chile, representing the regions that contain the vast majority of Native Forest Law enrollees.

Summary statistics for enrolled properties are shown in Table 2.1. Most enrollees enroll less than 15% of their property through the competition, and typical enrollee properties already have quite a large area of native forest cover and some plantation forest prior to the existence of the Native Forest Law. Section A.2.1 in the Appendix shows how descriptive statistics vary across subgroupings of the data such as contest type and compliance. Notable statistics include the fact that enrollees in the other interested party contest have a greater share of the property with both native forest and plantation forest prior to enrollment, possibly indicating that it is less costly for larger wealthier properties to maintain tree cover.

Table 2.1: Summary statistics for enrollees in Native Forest Law subsidy contest

Statistic	Mean	Median	Std. Deviation
Property size (ha)	201.116	45.120	514.733
Subsidized surface (ha)	15.460	5.840	39.632
Proportion of property subsidized	0.233	0.143	0.240
Bonus amount (UTM)	84.589	34	202.193
Received payment	0.319	0	0.466
Submitted management plan	0.371	0	0.483
Timber production objective	0.842	1	0.365
Received extensionist support	0.499	0	0.500
Pct. tree cover change (00-08)	0.074	0.001	1.365
Native forest	0.430	0.432	0.305
Plantation	0.151	0.026	1.223
Tree cover	0.683	0.774	0.302
Crop	0.024	0	0.102
Grassland	0.272	0.191	0.267
Dist. to native timber processing (km)	21.235	15.141	17.359
Dist. to any timber mill (km)	10.730	9.092	7.774

Figure 2.2 shows the average change in tree cover in the 5 years leading up to enrollees' application to the subsidy competition. The left panel shows the raw trends in tree cover for the typical property across both the smallholder and other interested party contests. The right panel shows the rate of tree cover change in the years leading up to enrollment across the two contests. In both contests, we see that tree cover was already increasing on the typical enrollee's property prior to enrollment. This is evident by the raw trends (left panel) and the positive rate of tree cover change (right panel) for all years immediately leading up to enrollment. The right panel of Figure 2.2 shows that the rates of tree cover gain across both contests was similar until the three years leading up to enrollment. At this point, the typical property in the other interested party contest increases the rate of tree cover gain. This means that by the time of application, landowners in the other interested party contest are already increasing forest cover at a rate nearly four times higher than enrollees in the smallholder contest. Smallholders on the other hand do

not change their rate of tree cover gain at any point leading up to the enrollment year. This observation raises the question of whether landowners in the other interested party contest applied to the competition with the hopes of subsidizing tree cover expansion that they were already undertaking without any subsidization.

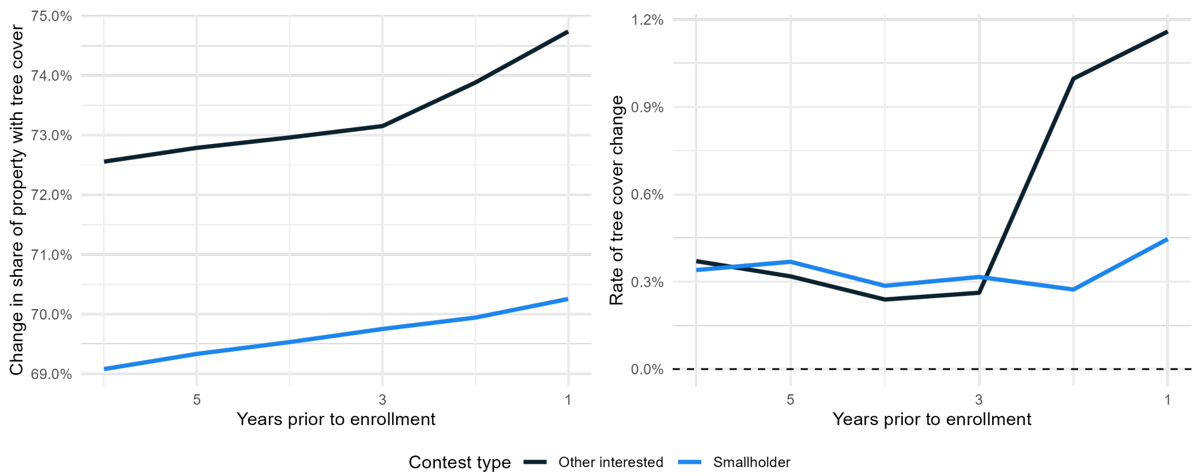

Figure 2.2: This figure shows rates of change by land cover type for enrollees in the years leading up to their application and enrollment in the Native Forest Law subsidy competition.

2.4 Program evaluation

2.4.1 Constructing a counterfactual

To quantify the environmental impacts of the Native Forest Law subsidy contest, we focus on landowners who complied with program requirements and received payment. As is the problem with many PES impact evaluations, enrollment is non-random. Landowners choose to enroll in the program and, in theory, have an opportunity cost equal to or lower than program payment. This means that the average enrollee likely has lower participation costs than the average unenrolled landowner. It is then ill-advised to simply use unenrolled properties as the counterfactual, since unobservable factors affecting enrollment could drive changes in forest cover outcomes, not enrollment.

In order to move toward a more convincing counterfactual, we first use matching as a pre-processing technique to generate a control group from amongst all unenrolled rural properties in the major forested regions of Chile. This should yield control properties with more similar opportunity costs to enrollees than amongst the general population. This approach is similar to many recent studies in the literature (e.g. Cisneros et al., 2022).

The covariates used for matching include environmental and economic characteristics likely to determine enrollment decisions and project performance. We include pre-enrollment property land-use including levels of native forest, plantation forest, and pasture. Landowners with similar levels of plantation forest, native forest, and other land uses on the property should face a similar decision about whether to enroll in a program involving native forest management. Other included covariates give a sense of a property's productive potential, remoteness, and timber market access. We also match on land cover pre-trends to help us build confidence in the common trends identification assumptions we make in the next section. In doing so, we use land cover trends leading up to the first Native Forest Law subsidy competition. This helps to avoid concerns of overfitting and should still allow us to see any major anicipatory land cover changes in our pre-trend analyses. Thus, seeing that pre-enrollment trends hold should lend further creedence to the matching process. Matches are made with replacement based on nearest neighbor propensity scores from a logit model. We include the two unenrolled nearest-neighbors for each program enrollee in the control group.

Prior to pre-processing, enrolled and unenrolled properties differ significantly. The typical enrollee had significantly less land engaged in pasture or agriculture and significantly more native forest already relative to the typical unenrolled property. Table 2.2 displays balance checks for all covariates used. The normalized mean difference and vari-

Table 2.2: Covariate balance before and after matching

	Norm. mean diff.		Variance diff.	
Covariate	**Unmatched**	**Matched**	**Unmatched**	**Matched**
Tree cover trend (05-08)	0.01	0.00	5.66	4.17
Crop trend (05-08)	-0.06	-0.01	6.78	3.51
Grassland trend (05-08)	0.06	0.02	13.90	3.12
Native forest (2001)	-1.32	0.12	0.37	1.35
Plantation forest (2001)	0.09	0.04	2.31	1.48
Tree cover	-1.03	0.20	1.42	1.04
Grassland	-0.37	-0.13	0.06	1.25
Crop	-0.15	-0.09	0.04	0.35
Shrubs	-0.12	-0.06	0.77	0.77
Development	0.02	-0.03	292.44	0.10
Water	0.01	-0.02	17819.17	1.78
Slope	-1.04	0.04	0.34	1.57
Elevation	-1.12	-0.08	0.24	1.41
Lattitude	-0.15	-0.09	0.80	1.00
Area	-0.37	-0.10	0.15	1.93
Dist. to industry	-0.60	-0.01	0.53	1.98
Dist. to native specific industry	-0.07	-0.02	0.98	1.10

ance were reduced for nearly every covariate after the matching process. Figure A.6 in the Appendix shows how comparability between selected covariate distributions drastically improved between treatment and control properties after matching. After processing, balance improved on every included covariate, and the normalized mean difference fell below the often-used threshold of 0.25 for every covariate.

One concern may be our decision to exclude rejected applicants from the control group. Given that these properties have revealed their intention to enroll in the program, it seems that they may have opportunity costs similar to program enrollees. However, the composition of the rejected applicants differs between smallholders and other interested parties depending on year, and the composition of the rejected applicant group is relatively unstable through time. Further, it would be difficult to make claims about the differences between the smallholders and other interested parties, it is much easier to be rejected in the other interested party contest for most contest years. We discuss the rejected applicants in more detail in Appendix A.2.7.

2.4.2 Main specification

We take advantage of our panel data setting and estimate the following equation to reveal the land cover impacts of the Native Forest Law subsidy contest:

$$outcome_{it} = \beta_0 + \beta_1 \times intensity_{it} + \gamma_i + \lambda_t + X_{it} + \epsilon_{it} \tag{2.1}$$

where $outcome_{it}$ represents the share of property i engaged in a specific land cover outcome in year t; $intensity_{it}$ represents the proportion of property i enrolled through the Native Forest Law subsidy contest in year t; and γ_i and λ_t represent property and year fixed effects, respectively. Property fixed effects (γ_i) control for unobserved time invariant characteristics such as landowner preferences. Year fixed effects (λ_t) control for time-varying shocks that are common across all properties such as changes in other environmental policies. Conditional on covariates and fixed effects, β_1 recovers the impact of enrollment in the Native Forest Law contest, conditional on compliance.

Because Equation 2.1 relies on property and year fixed effects, it falls under the umbrella of two-way fixed effects (TWFE) estimators. This literature has received ample attention in recent years, particularly in the case of binary treatment (i.e., whether a property enrolled) (Roth et al., 2022; de Chaisemartin and D'Haultfœuille, 2022). In this context, binary treatment would ignore the proportion of the property enrolled through the contest. Equation 1 is valuable in our context, because the median landowner enrolls less than 15% of their land in the program, with significant variation across properties (Table 2.1).

Importantly, $intensity_{it}$ represents a continuous treatment in the context of TWFE estimation. Callaway et al. (2021) decompose TWFE estimators when treatment is continuous and show that β_1 represents the weighted average change in outcomes from incremental changes in land enrollment across and within periods. Thus, our identification relies on the following assumption: properties that enrolled an additional increment of land in the Native Forest Law contest, must experience the same evolution in outcomes as properties that never enrolled the increment. We evaluate the plausibility of this common trends assumption based on both raw trends and an event study approach in Appendix A.2.4.

2.4.3 Event study

The dynamics of payments for ecosystem services are important, and perhaps moreso in the restoration context. Tree cover is not established instantaneously, including in satellite-derived measures of tree cover, where a pixel is generally only classified as tree cover if it exceeds a threshold of canopy cover. Survivorship of trees is also key, as many planting initiatives have led to minimal long-term success (Coleman et al., 2021). These factors are echoed in the Native Forest Law payment scheme, where landowners are not even eligible to receive their payment in the first year of enrollment.

We use the estimator developed in Callaway and Sant'Anna (2020) to generate event study treatment effects. It is important to note that this estimator relies on binary treatment. While this cannot account for the fact that properties enroll only a selected proportion of a property, it provides robustness properties not true of Equation 1. First, our event study estimates rely on a relatively weaker conditional common trends assumption.

Common trends need only hold after conditioning on our detailed set of time-invariant pre-treatment characteristics. The Callaway and Sant'Anna (2020) estimator is also robust to general treatment effect heterogeneity, which, in severe cases, could flip the sign of our main specification's treatment effect estimates (de Chaisemartin and D'Haultfœuille, 2020).

2.4.4 Program evaluation results

Table 2.3 shows that using the matched control group, tree cover expands on the characteristic enrolled property after enrollment relative to the counterfactual. It also sees a decline in grassland and cropland. Since $intensity_{it}$ is the proportion of the property enrolled, β_1 can be interpreted as the impact of enrolling the average landowner's full property through the Native Forest Law competition. For the average property, enrolling the full property leads to a 0.63 percentage point (0.91%) increase in the share of the property with tree cover. However, since the average landowner enrolls 23% of their land (Table 2.1), the impact for the typical enrollee is closer to a 0.14 percentage point (.21%) increase. It also led to a 0.19 (6.55%) and 0.45 (1.79%) percentage point decline in the share of the typical enrollees' property with crop and grassland, respectively. Standard errors are clustered at the property, the level of the decision-making unit and at which treatment is assigned.

In order to gauge the plausibility of the matched control group as a counterfactual, we consider a similar estimation strategy using program enrollees who never submitted a management plan. These non-compliant properties had applications selected through the competition, however, they failed to provide a management plan within the required six-month window following selection. Therefore, they were dropped from the program. Because these landowners dropped from the program within such a short time-frame and

Table 2.3: Estimates of subsidy impact using matched control group

	Tree cover	Crop	Grassland
Intensity	0.00628**	-0.00188**	-0.0045*
	(0.00239)	(0.00076)	(0.00225)
Num.Obs.	183521	183521	183521
R2	0.970	0.899	0.958
Control group	matched 2-to-1	matched 2-to-1	matched 2-to-1
2008 mean	0.699	0.029	0.251

Standard errors are clustered at the property level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

never submitted a management plan indicating exact project details, it is unlikely they engaged in sustained restoration activity. We use the same pre-processing techniques to generate a matched control group for these non-compliers and again estimate Equation 2.1. Table 2.4 shows that the estimated tree cover treatment effect is not only statistically insignificant, but also of no meaningful magnitude. Although there is a statistically significant effect for crop, the magnitude is very small. It may be the case that landowners in this group shifted away from crop cover in response to having their application selected, but never could establish tree cover in its place. Finding no meaningful effect amongst properties that engaged with the program but dropped out soon after enrollment lends credence to our main estimation strategy.

Figure 2.3 shows event study treatment effects based on the estimator developed in Callaway and Sant'Anna (2020). Similar to results based on Equation 2.1, these estimates indicate that the Native Forest Law subsidies increased tree cover, while reducing crop and grassland. These graphs indicate that most of the increased tree cover seems to have come from grassland conversion. Pre-treatment estimates in Figure 2.3 represent pseudo-ATTs, which are all indistinguishable from zero across the three land cover types. The conditional common trends assumption on which these estimates rely, thus, seems

Table 2.4: Estimates of subsidy impact on non-compliers using matched control group

	Tree cover	Crop	Grassland
Intensity	6e-05	-1e-04**	5e-05
	(5e-05)	(4e-05)	(5e-05)
Num.Obs.	321309	321309	321309
R2	0.970	0.893	0.957
Control group	matched 2-to-1	matched 2-to-1	matched 2-to-1
2008 mean	0.699	0.023	0.258

Standard errors are clustered at the property level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

plausible. Our event study estimates are robust against concerns surrounding general treatment heterogeneity and do not rely on assumptions as stringent as those from our main specification, so seeing that these results are similar to those from Equation 2.1 lends confidence to our main results. Similar to the main TWFE specification, estimates are based relative to both the control group and not-yet-treated properties. In the Appendix (Figure A.9), we include similar results based on inclusion of only the never-treated control properties in the control group.

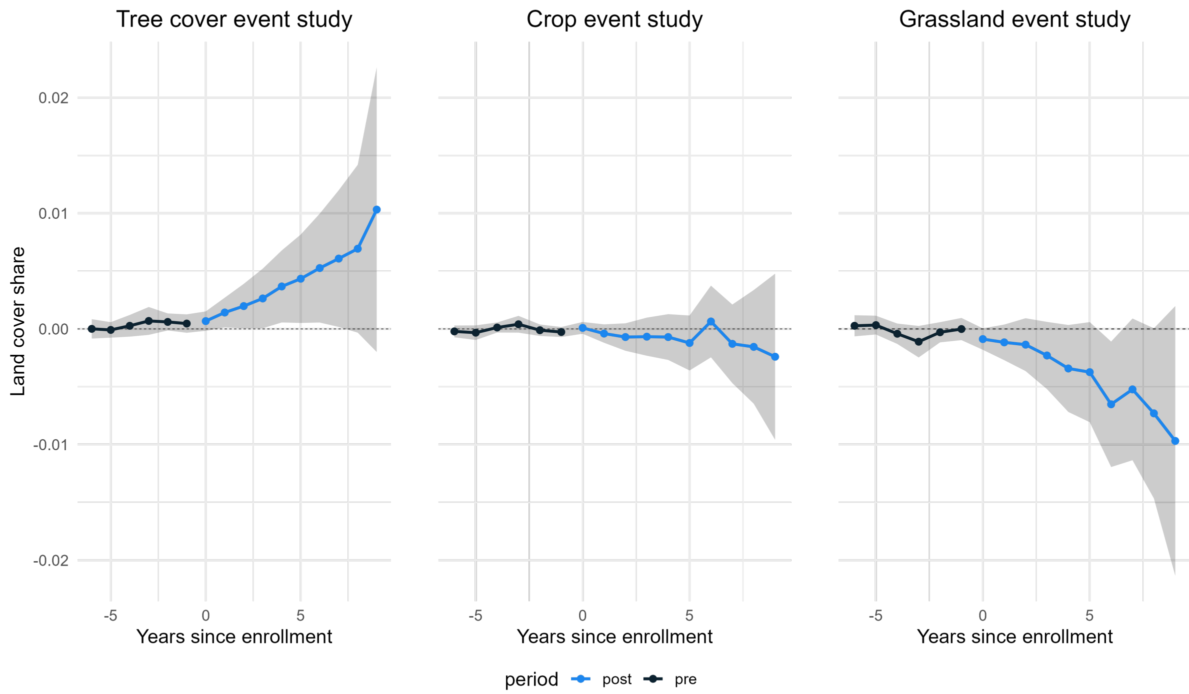

Figure 2.3: This figure shows event time treatment effects of the Native Forest Law subsidy contest for beneficiary properties. The three panels show event time treatment effects based on binary treatment for three land cover types: Tree cover, Crop, and Grassland.

Looking at Figure 2.3, the treatment effect has not leveled-off. This could be due to a couple different factors. First, landowners may still be establishing new tree cover 9 years after initial enrollment (10th year of treatment). In this case, the estimates from both Equation 2.1 and the event study will underestimate the carbon and land cover impacts of the subsidies. However, it may also be the case that this sustained effect is due to the changing composition of cohorts in each event-time window. Appendix A.2.6 shows that there is no clear systematic relationship between cohorts and treatment effect. We also present results of a "balanced" event study, proposed in Callaway and Sant'Anna (2020), in Figure 2.4. This balanced event study estimates event-time treatment effects for cohorts that are observed for at least 8 years post-enrollment (i.e., the 2009 and 2010 cohorts). The benefit of the balanced event study is that there is no change in treatment unit composition across event-time windows. The major cost of balancing is that fewer groups are used to compute these event-study-type estimands, which can lead to less informative inference (Callaway and Sant'Anna, 2020).

As expected, both types of event study results do imply that the treatment effect is increasing over time, suggesting that the estimates based on Equation 2.1 may underestimate the ultimate impact of the program. We believe that the balanced event-study estimates represent a more informative path of tree cover impacts through time than those presented in Figure 2.3. As expected, there is no clear effect for the first few years following enrollment. The subsidized activities often take several years to complete, and newly planted trees will not be picked up instantaneously in the outcome variable. While we expect biomass accumulation to continue through time, tree cover extent is not likely to continue to be established infinitely far in the future as a direct result of subsidized projects, as suggested by Figure 2.3. We see in Figure 2.4 that balanced treatment effect estimates do increase through time but seem to level off 7 years post-enrollment. Tree cover is largely established 4 to 7 years after enrollment. The final event-time estimate from the balanced event study suggests that enrollment led to a 0.70 percentage point (1.00%) increase in the share of the property with forest cover on average. If we assume that all of the forest cover gain ocurred in the subsidized proportion of the property, this effect is closer to 3.04 percentage points (4.34%), much greater than the effect suggested by results in Table 2.3.

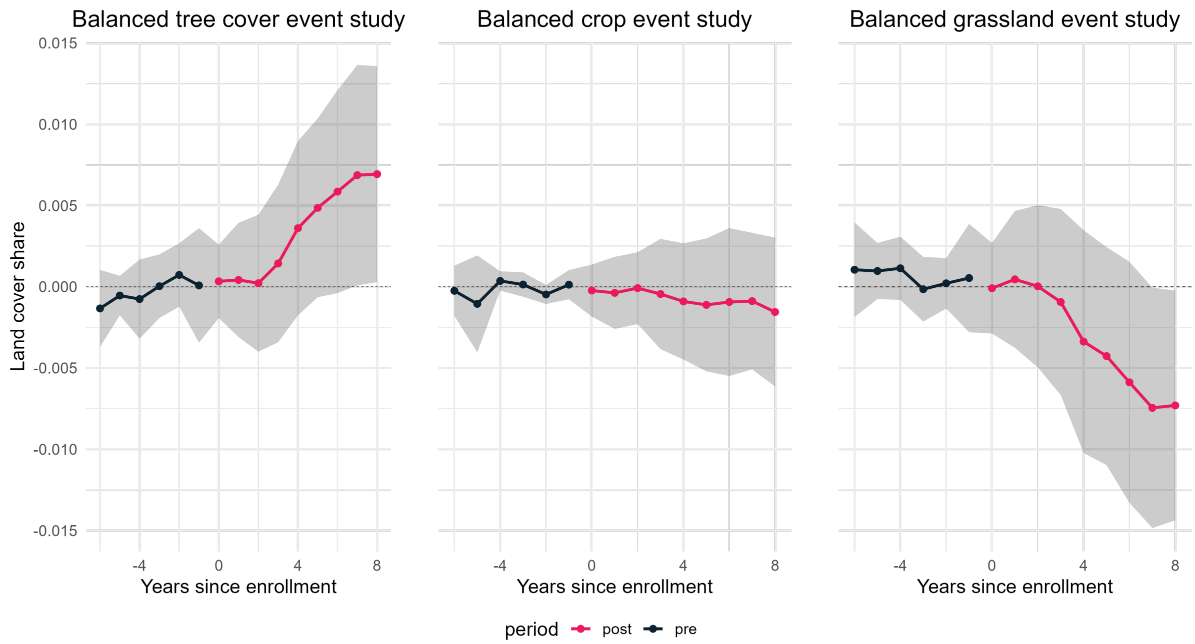

Figure 2.4: This figure shows balanced event time treatment effects of the Native Forest Law subsidy contest for beneficiary properties. Only estimates cohorts experiencing treatment for at least 8 years are included, so there is no change in the composition of treatment cohorts by event-time window. The three panels show event time treatment effects based on binary treatment for three land cover types: Tree cover, Crop, and Grassland.

2.4.5 Carbon impacts

Previous work identifies the potential for restoration to remove significant amounts of carbon from the atmosphere at a relatively low carbon price (e.g. Busch et al., 2019). Many of these estimates, however, do not account for social viability (Austin et al., 2020) and cannot be interpreted as causal. Moreover, there exists little evidence on the ex-post costs of real-world policy to achieve these removals. We present back of the envelope estimates of the carbon price achieved on the average property through the subsidy competition. We provide these estimates of the carbon price based on estimates of carbon per hectare for native forest cover, crop, and grasslands by region in Chile.

We estimate that the subsidy competition paid the average enrollee \$36.78 USD per tonne CO2 stored in aboveground biomass. While this estimate relies on assumptions about the carbon content of additional tree cover, we believe it to be relatively conservative. Estimates of carbon content per hectare of native forest used in this calculation are based on mature native forests by region in Chile. This means that this estimate assumes all additional tree cover extent is mature native forest. While this number may overestimate the current carbon content in newly established native forest subsidized through the Native Forest Law, we believe it is reasonable to assume subsidized forest will eventually mature and attain these levels of carbon content. Moderating factors include the fact that our estimates are unable to capture carbon benefits achieved through restoration of already standing forest, which is relatively common amongst subsidized projects. As seen in our event study estimates (Figures 2.3 and 2.4), tree cover impacts are increasing through time. Estimates from Equation 1 do not account for these dynamics and may further underestimate the ultimate carbon impacts of the program if recent cohorts see dynamic effects similar to the earliest cohorts. A similar exercise using only treatment effects from the 2009 and 2010 cohorts (Figure 2.4) yields a much lower price of $12.35 USD per tonne CO2. These factors lend confidence that our primary estimate is not a significant underestimate of the true final cost.

2.5 Targeting for social development

Payments for ecosystem services programs have often targeted payments toward priority groups with the hopes of achieving both environmental gains and poverty alleviation. The Native Forest Law follows a similar strategy, prioritizing the applications of smallholders and landowners with other priority social characteristics. In this section, we examine whether this strategy undermined the environmental effectiveness of payments, as it often has in the avoided deforestation setting.

Chile's Native Forest Law targets in two primary ways: 1) separating applicants into a smallholder and other interested party contest; and 2) scoring applications based on both social and project characteristics. In order to qualify for the smallholder contest, landowners must have assets and a property size below a set threshold for each region. The benefits of qualifying for the smallholder contest include 15% higher payments as well as increased odds of having an application selected through the contest. Landowners in the smallhoder contest are also significantly more likely to receive extensionist support. Extensionists often helped landowners understand the potential for different activities on the property and assisted landowners to complete and submit an application. In 2019, CONAF also began offering technical assistance to selected landowners through the smallholder contest. While the Native Forest Law used these separate contests to alleviate concerns that large corporations with significant assets would reap the rewards from the program, smallholder classification is still quite broad. In the typical region, properties up to 200 hectares can qualify for the smallholder contest. Figure A.4 shows that while many applicants in the smallholder contest are truly landowners with small properties, many larger properties are able to enter the smallholder contest (Full eligibility rules in Appendix A.2.2).

The program used a scoring system in order to assign project funding priority within each contest. Projects were granted funding in descending order of score until the allocated funding had been assigned. This meant that projects sometimes went unfunded because of a low score, although no ex-ante cutoff existed. This was particularly common in the other interested party contests, which were granted funding after the smallholders. In some years, a second smallholder contest was held, causing smallholders to go unfunded because of low scores. The scoring criteria include factors related both to landowner, property, and project characteristics. This score, although not always critical for smallholder applicants, provides insight to program administrators' preferences for project prioritization. It amounts to a sum of social and project specific characteristics:

$$score_i = social_i + project_i$$

where $social_i$ represents components of the score deemed to be of social importance by program designers; and $project_i$ represents components of the score representing project specific characteristics unrelated to the landowner themselves. The social score, $social_i$ and the project score, $project_i$ can be further broken down as follows:

$$social_i = \gamma_t VI + \beta_t VPS$$

$$project_i = \lambda_t VP + \psi_t VT$$

, where

- VI = social characteristics of interest, including property size (higher scores to smaller properties) and total subsidy amount (penalizing particularly large projects)

- VPS = other priority social characteristics, including indigenous status (higher scores to indigenous landowners or communities)

- VP = project characteristics

- VT = land characteristics

, and γ_t, β_t, λ_t, and ψ_t represent the weights given to each category in year t.

This score demonstrates the specific project characteristics prioritized by program administrators. While the score does not explicitly target high-poverty landowners, if the objective is to promote rural economic development and poverty alleviation, we are interested in whether the score actually helps administrators target low-income landowners. While we cannot observe individual landowner poverty, we examine whether comuna-level poverty is associated with social and project scores:

$$ln(ComunaPov_i) = \rho_0 + \rho_1 \times ln(social_i) + \rho_2 \times ln(project_i) + e_i \tag{2.2}$$

The results are shown in Table 2.5 and indicate that $social_i$ is associated with increased comuna-level poverty, holding $project_i$ constant. We also see that $project_i$ is associated with greater levels of poverty, perhaps indicative of less profitable alternate uses of the land in high-poverty comunas. Including region fixed effects in the model eliminates the clear association of the project score with poverty, however. While we cannot observe whether $social_i$ is predictive of landowner-specific poverty, it did help Native Forest Law administrators prioritize participation of landowners in high-poverty comunas.

Table 2.5: Social score is associated with relatively higher comuna level poverty

Outcome var.	ln(ComunaPov)			
	(1)	(2)	(3)	(4)
ln(Social score)	0.09601***	0.14969***	0.07363***	0.19827***
	(0.01847)	(0.02249)	(0.01293)	(0.01679)
ln(Project score)	0.13323***	0.13470***	-0.01432	-0.00567
	(0.01367)	(0.01721)	(0.00907)	(0.01097)
Num.Obs.	12828	8750	12828	8750
R2	0.009	0.009	0.617	0.645
Subsample	Full sample	Smallholder	Full sample	Smallholder
Region FE	No	No	Yes	Yes

Standard errors are clustered at the property level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

2.5.1 Targeting's impact on program effectiveness

To answer the question of whether the Native Forest Law's strategy to target for social development undermined environmental impacts, we consider several variables that correspond to application prioritization and socioeconomic status. We first consider treatment effect heterogeneity based on the contest into which projects were admitted: the smallholder or other interested party contest. Table 2.6 shows that payments through the smallholder contest were more effective than those through the other interested party contest, and further, that payments to the average enrollee in the other interested party contest did not yield any positive tree cover impacts at all. Landowners in the other interested party contest were not eligible for the smallholder contest, either because they had assets exceeding the allowable threshold or a particularly large property (generally greater than 200 hectares). This means that the wealthiest and largest properties were generally cost-ineffective enrollees.

This implies that the decision to separate applicants into the smallholder and other interested party contests was helpful. Applicants to the other interested party contest were much more likely to be rejected. In most years of the competition, smallholders were unlikely to be rejected unless they did not submit proper paperwork or proposed activities that were not eligible to be subsidized. Prioritizing the participation of properties in the smallholder contest, therefore, seems to increase the average effectiveness of payments relative to the case where no targeting exists across contests.

Table 2.6: Heterogeneous tree cover impacts by contest

Outcome var.	Tree Cover		Crop		Grassland	
	(1)	(2)	(3)	(4)	(5)	(6)
Intensity	0.00714**	-0.00758	-0.00163**	-0.00448	-0.00532**	0.00725
	(0.00267)	(0.00555)	(0.00072)	(0.00304)	(0.00243)	(0.0058)
Num.Obs.	133114	50407	133114	50407	133114	50407
Subsample	Smallholder	Other interested	Smallholder	Other interested	Smallholder	Other interested

Standard errors are clustered at the property level.

* p < 0.1, ** p < 0.05, *** p < 0.01

While Table 2.6 shows that the average smallholder generates more tree cover per percent of the property enrolled in the Native Forest Law than the average other interested party, smallholder qualification is quite broad. Many in the smallholder contest had large properties or relatively high levels of wealth (i.e., assets). In order to understand how the effectiveness of payments varied across finer-grain measures of vulnerability, we explore heterogeneity based on comuna-level poverty. Although poverty was not explicitly used by program administrators to target, higher values of $social_i$ are associated with higher comuna-level poverty (Table 2.7), and poverty alleviation is generally the primary goal of targeting payments to marginalized groups. We use the following regression:

$$\begin{aligned} outcome_{it} = &\alpha_0 + \alpha_1 \times intensity_{it} + \alpha_2 \times intensity_{it} \times ln(ComunaPov_i) + \quad (2.3) \\ &\gamma_i + \lambda_t + X_{it} + e_{it} \quad (2.4) \end{aligned}$$

where $ComunaPov_i$ represents the percent of poverty in the Comuna where landowner i resides. Here, α_2 represents the parameter of interest, indicating the association of $ComunaPov_i$ with the treatment effect of the Native Forest Law subsidies. We see that projects located in comunas with higher rates of poverty actually yielded higher treatment effects. While we cannot make claims that increased poverty actually caused these increased treatment effects, it was predictive of increased treatment effects per enrolled hectare in the Native Forest Law subsidy competition, and thus, targeting these marginalized landowners did presumably improve program outcomes.

In order to better interpret the association of comuna-level poverty with program impacts, Figure 2.5 displays marginal effect plots for tree cover. These plots reveal how the treatment effect varies across values of comuna-level poverty. The left panel shows how the tree cover treatment effect varies across the percent of comuna households in poverty, while the right shows how the treatment effect varies across values of its natural log. Again, we see that increased levels of comuna-level poverty are associated with greater program impacts on tree cover. The average percent of poverty within comunas of enrolled properties was 19.80% (blue line), above the country average of 16.11%. While the Native Forest Law subsidies have greater tree cover impacts with increased poverty, the subsidies are predicted to have non-positive effects on tree cover only within particularly low poverty comunas.

Table 2.7: Heterogeneous tree cover impacts by comuna-level poverty

Outcome var.	Tree Cover		Crop		Grassland	
	(1)	(2)	(3)	(4)	(5)	(6)
Intensity	-0.03326**	-0.03952**	0.00297	0.00393	0.03115*	0.03522*
	(0.01548)	(0.01747)	(0.00306)	(0.00313)	(0.01617)	(0.01783)
Intensity x ln(ComunaPov)	0.01631**	0.01952**	-0.00200	-0.00233*	-0.01471**	-0.01696**
	(0.00684)	(0.00790)	(0.00127)	(0.00134)	(0.00694)	(0.00789)
Subsample	Full sample	Smallholder	Full sample	Smallholder	Full sample	Smallholder

Standard errors are clustered at the property level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

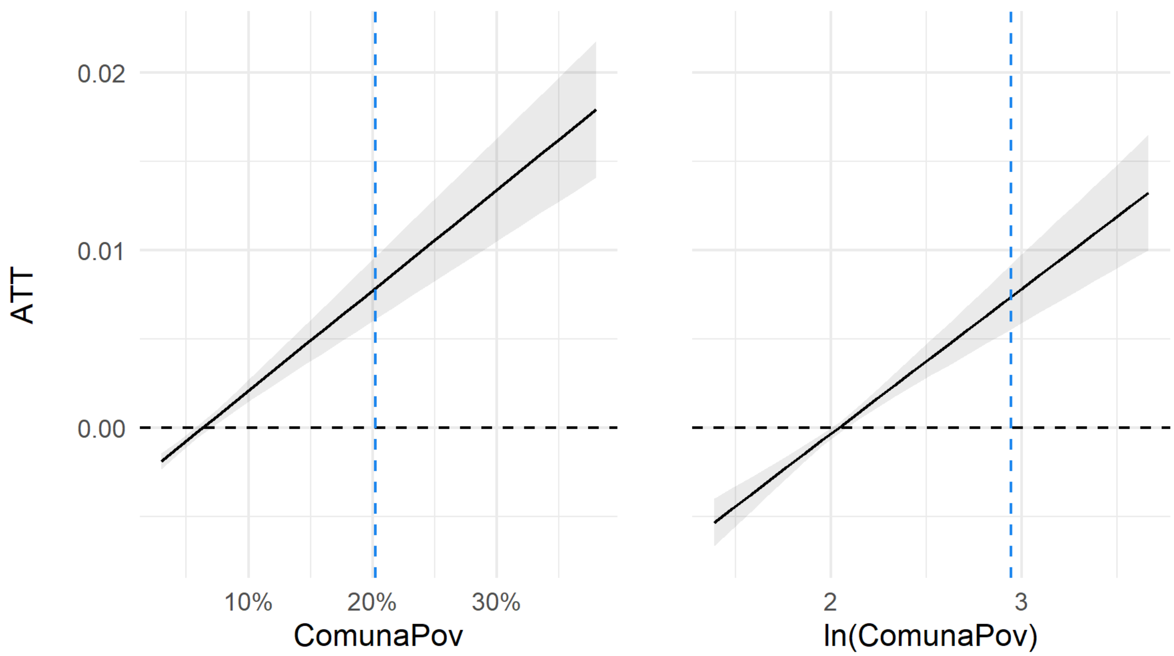

Figure 2.5: Marginal effects plots show how treatment effects vary across different values of comuna-level poverty. Increased comuna poverty is associated with increased treatment effects, although treatment effects are expected to be positive except within particularly low-poverty comunas. The vertical blue dashed line shows mean comuna-level poverty within each panel.

2.5.2 Understanding the role of non-compliance

Non-compliance in the Native Forest Law subsidy competition is high. There are multiple reasons why non-compliance may influence the effectiveness of program administrators' strategy to target for social development in payments for restoration programs. First, targeting payments toward groups that suffer from high levels of non-complance may be costly if compliance is not verified, as is often the case in low-income countries (Alix-Garcia and Wolff, 2014). This is because targeting priority groups would increase the probability of unconditional cash transfers. As discussed, CONAF verified compliance in the Native Forest Law subsidy competition. By requiring both the submission of a management plan as well as third-party verification of activity completion to receive payment, CONAF saved significant resources by never paying landowners who did not complete contracted activities. These compliance checks prevented payment of $28.73 million USD to landowners who never completed their planned forest restoration activities. As seen in Table 2.4, these projects led to no meaningful tree cover impacts, so compliance checks greatly improved overall cost-effectiveness of the program relative to the case where non-compliers were paid.

Second, if targeted characteristics are positively correlated with both program impacts and non-compliance, administrators may want to better understand how to improve compliance amongst these groups. In the Native Forest Law, nearly two-thirds of applicants fail to engage meaningfully with the program after having an application approved in the subsidy competition. These landowners showed interest in participating but ultimately chose to drop out or were unable to meet administrative hurdles. Non-compliers are inherently different than compliers, so even if priority characteristics are associated with increased non-compliance, we cannot claim that they would have produced the same tree cover benefits as similarly prioritized compliers.

Although CONAF avoided unconditional cash transfers, the association of priority characteristics with compliance is informative to our understanding of the effectiveness of the program's design. In order to examine how prioritization of social characteristics in the Native Forest Law scoring system was correlated with compliance, we use regressions of the following form:

$$complied_i = \psi_0 + \psi_1 \times ln(social_i) + \psi_2 \times ln(project_i) + X_i + u_i$$

where $complied_i$ is a dummy variable equal to 1 if landowner i followed through and received payment for successful project completion.

Table 2.8 shows the results of these regressions. Our coefficient of interest is ψ_1, which captures the association of an increase in $social_i$ on compliance, holding the other parts of the project component constant. We see that higher values of $social_i$ are associated with a decreased probability of compliance in the smallholder contest, indicating that landowners who were prioritized by the program were less likely to comply. This association does not hold in the other interested party contest, however. That said, the smallholder contest saw increased levels of compliance on average, possibly due to higher payments and greater use of extensionists. Comuna-level poverty is not clearly associated with reduced compliance, holding $social_i$ and $project_i$ constant. Column 5 shows that extensionists were associated with large increases in compliance probability, perhaps indicating that some of the risk of non-compliance by priority groups can be mitigated.

The results of Table 2.8 suggest a few lessons for the design of payments for restoration programs where non-compliance is high. First, even when priority characteristics are associated with greater levels of non-compliance, administrators can still target these groups without sacrificing cost-effectiveness if compliance is checked. CONAF checked compliance in an inexpensive way, simply requiring submission of a management plan and verification that landowners completed the expected activities. This allowed them to target for social development without increasing unconditional cash transfers. Second, the most cost-effective enrollees conditional on compliance were also the enrollees least likely to follow-through. One possible concern is that ex-post payments may actually reduce compliance if landowners are liquidity constrained (Jack et al., 2022). While untestable, if non-compliers would have provided environmental benefits comparable to similar compliers, efforts to support these landowners early in the program or relieve liquidity constraints could improve overall impacts without sacrificing effectiveness.

Table 2.8: Social score is negatively associated with compliance

Outcome var.	Complied				
	(1)	(2)	(3)	(4)	(5)
ln(Social score)	-0.07505***	-0.07997***	0.00597	-0.08223***	
	(0.01985)	(0.02024)	(0.03622)	(0.02453)	
ln(Project score)	0.06191***	0.06305***	0.06491***	0.08496***	
	(0.01313)	(0.01340)	(0.01346)	(0.01683)	
Smallholder	0.11672***	0.11573***	0.41503***		0.18457***
	(0.01002)	(0.01023)	(0.10594)		(0.04907)
ln(ComunaPov)		0.00146	0.00345	0.00021	0.01620
		(0.01295)	(0.01297)	(0.01673)	(0.01630)
ln(Social score) x Smallholder			-0.12038***		
			(0.04229)		
ln(Subsidy amount)					0.01850***
					(0.00342)
Extensionist					0.12641***
					(0.00842)
ln(ComunaPov) x Smallholder					-0.02007
					(0.01677)
Num.Obs.	13294	12828	12828	8750	12821
R2	0.073	0.075	0.076	0.085	0.090
Subsample	Full sample	Full sample	Full sample	Smallholder	Full sample
Region FE	Yes	Yes	Yes	Yes	Yes

Standard errors are clustered at the property level.

* p < 0.1, ** p < 0.05, *** p < 0.01

2.6 Discussion and conclusion

Prominent initiatives such as the Bonn Challenge, Trillion Trees Initiative, and UN Decade on Ecosystem Restoration hope to address the intertwined challenges of rural poverty, climate change and biodiversity loss through large-scale afforestation and reforestation. Initial national plans indicate that many countries will follow Chile's earlier model for tree cover expansion, relying heavily upon subsidies and plantation forests to achieve their commitments (Lewis et al., 2019). In light of the fact that this model may have negative impacts on native forest extent, biodiversity, and other outcomes, payments for native forest restoration may provide a more sustainable and socially beneficial path forward. Further, this may lead to increased additionality if alternative policies simply subsidize plantation forests that would have been planted anyways.

We analyze the land cover and carbon impacts of payments for native forest restoration through Chile's Native Forest Law competition. To control for selection bias, we match enrolled properties to similar unenrolled properties and use difference-in-differences methods that rely on common trends assumptions. We find that the payments led to native forest expansion and reduced the extent of cropland and grassland on the average beneficiary property. Back of the envelope calculations suggest that payments achieved carbon removals a relatively low carbon price of \$36.78 USD per tonne CO2. These findings indicate that large-scale payments for native forest restoration may be a viable approach to achieve carbon removals as countries search for policies to help scale-up tree cover in line with their commitments.

Targeting for poverty alleviation in payments for avoided deforestation has often led to tradeoffs in terms of environmental efficacy (e.g. Alix-Garcia et al., 2015). The effectiveness of payments for avoided deforestation depends on deforestation risk, and a large literature shows that enrollment of areas with low deforestation pressure tend to generate minimal additional forest cover gains (Cisneros et al., 2022; Alix-Garcia et al., 2019). For example, enrollment of many landowners in low deforestation pressure areas led to low additionality in Costa Rica's national PES program (Hanauer and Canavire-Bacarreza, 2015) and in several PES programs in Brazil (Giudice et al., 2019; Cisneros et al., 2022). Deforestation pressure is often negatively correlated with poverty, and therefore, targeting payments toward these areas often generates little avoided deforestation. Moreover, capital constraints may undermine environmental effectiveness when targeting low-income landowners (Alix-Garcia et al., 2012).

The question of whether this strategy similarly undermines environmental effectiveness in payments for reforestation and restoration has gone unstudied. Land owned by smallholders may not suffer from deforestation risk in the absence of payments, but relatively poorer landowners may not produce new forest without payments. This may be due to credit constraints or a lack of technical capacity that prevent smallholders from scaling up restoration activities. In Chile's Native Forest Law subsidy competition, we find that socially prioritized landowners generated the greatest environmental benefits conditional on compliance. Smallholders in high-poverty comunas actually saw the greatest tree cover gains per enrolled hectare. In contrast, paying larger, wealthier landowners to restore forests generated no additional tree cover or carbon benefits. This may be due to the fact that larger wealthier property owners had scaled-up reforestation even prior to applying for Native Forest Law subsidies. Future work should focus on the degree to which this targeting improved livelihoods, as participation does not necessarily guarantee poverty alleviation (Jayachandran, 2022). Non-compliance was relatively higher amongst prioritized landowners, which could have been exacerbated by ex-post payments to liquidity or credit constrained landowners. Verifying project completion prior to payment did allow administrators to target priority groups without increasing the probability of unconditional cash transfers, which would degrade program cost-effectiveness. These findings indicate that targeting marginalized groups was not detrimental to environmental effectiveness, as has often been seen in payments for avoided deforestation, but was in fact beneficial.

Chapter 3

Education impacts of metropolitan tree cover: evidence from invasive species induced tree loss

3.1 Introduction

There is a growing recognition of the need for research on the economic and social impacts of ecosystem disruption (Ferraro et al., 2019b; Fenichel et al., 2014). A vast literature has sought to quantify the value of ecosystem services (Farber et al., 2002), however, causal evidence detailing how changes in biological features of the environment affect human behavior is scarce (Frank and Sudarshan, 2023). This is particularly true for the relationship between economic outcomes and invasive species induced ecosystem changes. Economists have only recently started to study the pathways through which invasive species affect economic outcomes (Jones, 2020).

Prior research has established environmental drivers of educational outcomes including the impacts of temperature (e.g. Park et al., 2020a) and pollution (e.g. Marcotte,

2017) on learning. Vegetation and greenery has long been associated with improved educational outcomes (e.g. Kweon et al., 2017), however, much of this work is purely correlational. Prior work has struggled to establish causality in this context, because tree cover is often positively associated with income and other factors that drive educational achievement. Because of these challenges, there exists little to no evidence detailing the causal relationship between tree cover and education outcomes outside of laboratory settings. Trees have been shown to mitigate extreme temperatures and air pollution in addition to providing psychological benefits (Turner-Skoff and Cavender, 2019), each of which are mechanisms through which tree cover may affect education outcomes.

In this paper, I recover the causal impacts of the emerald ash borer on both tree cover and education outcomes in the metropolitan Chicago region. To overcome the aforementioned challenges previous studies have faced to establish causality in similar contexts, I take advantage of the idiosyncratic spread and detection of a forest-attacking pest, the emerald ash borer. The emerald ash borer represents one of the most destructive invasive species in the United States. A beetle that exclusively targets ash trees, the ash borer has decimated the primary street tree used in many US cities. This includes metropolitan Chicago, where prior to the arrival of the ash borer, ash were the most common non-invasive tree species in the region's streets and parks (Morton Arboretum, 2020). The staggered spread of the ash borer throughout the region provides changes to the provision of tree cover and associated ecosystem services throughout the region.

Using a difference-in-differences approach robust to general treatment effect heterogeneity, I find that ash borer infestation led to a 1.4% percent reduction in the canopy cover of affected areas. Infestation led to not only increased tree cover loss but also declines in tree cover gain in affected areas. Further, infestation led to relatively poorer standardized test performance at schools within the direct vicinity of a confirmed infestation site. Infestation reduced the share of students that met or exceeded benchmarks by 1.86% at the average exposed school. Low-income students' test performance was affected most, while there was little to no impact on the non-low income group.

The causal interpretation of these results is supported by several factors. First, I provide event study estimates showing the time-varying tree cover and test score impacts. Treatment and control groups showed similar trends in outcomes prior to the arrival of the ash borer, lending credence to the conditional common trends assumption needed for causal interpretation. While ash borer arrival leads to an initial spike in rates of tree cover loss, longer term canopy cover decline comes, in part, from reduced tree cover gain in affected areas. This is consistent with work showing that tree planting budgets and priorities changed following the arrival of the pest (e.g. Hauer and Peterson, 2017). Further, the canopy cover impacts are greatest in areas where ash make up a greater proportion of an area's tree population. There are no impacts of ash borer infestation on school enrollment, something that we wouldn't expect to change.

This work makes several key contributions. First, this paper complements the literature on the economic impacts of invasive species (e.g. Epanchin-Niell, 2017). Recent work in economics has explored the impacts of the emerald ash borer (e.g. Jones, 2023), but to my knowledge, my work is the first to measure the ash borer's impact in a concentrated metropolitan region. It is also the first to examine invasive species' impacts on education outcomes. I also contribute to the literature detailing the environmental drivers of education performance, causally linking invasive species and subsequent changes in vegetation to education outcomes for the first time outside of a laboratory setting.

3.2 Background

3.2.1 Emerald ash borer in the Chicago region

One major threat to metropolitan and urban tree populations is the introduction of insect pests. The emerald ash borer is one such pest, first detected in North America in 2002. While it is one of several introduced pests that has affected US tree cover, it may represent a worst-case scenario (Herms and McCullough, 2014). The ash borer was first detected in North America in 2002 and has been referred to as the most destructive forest pest ever introduced to the United States (Nowak et al., 2016). Ash borer exclusively target ash trees, and infestation is essentially fatal to any infested tree, so the arrival of the pest has meant the death or removal of millions of ash trees across the country.

The Chicago region is the third-largest metropolitan region in the United States, and lies in a region heavily affected by emerald ash borer. The impacts on the region's ash populations has been well recognized. Prior to the arrival of the ash borer, ash trees were the most numerous non-invasive tree species in the region (Morton Arboretum 2020). However, a Chicago region tree census revealed that the area's standing ash population nearly halved between 2010 and 2020, dropping from an estimated 13 million to under 7 million (Morton Arboretum 2020). Of those 7 million remaining standing trees, 4 million are either dead or in decline. Further, as of 2020, more than 30% of ash trees in the region are saplings, likely having regenerated from removed adults. Although many ash trees were replaced with alternative species and canopy cover broadly increased across the region, the overall number of large trees ($>$ 6 inch diameter) dropped as mature ash were replaced by smaller trees.

There exists previous work on impacts of ash borer and other forest attacking pests in the economics literature. For example, Tan (2022) uses staggered county-level ash borer detection to show its impact on mortality, through a channel of increased pollution. Jones and McDermott (2018) show that ash borer arrival led to increased rates of mortality and morbidity, stemming from increased high temperatures. Druckenmiller (2020) instruments for tree mortality in the American West using the temperature threshold at which bark beetles experience winter die-off. They find that beetle-induced tree mortality decreased the value of timber tracts and home values and reduced hazard protection from air pollution, flood risk, and burn area.

3.2.2 Educational responses to ecosystem disturbance and vegetation

Trees provide substantial ecosystem services, particularly in metropolitan settings. These benefits include reduction of traffic pollution, psychological benefits, and moderation of hot temperatures. The association between the amount of trees and improved educational attainment and performance has been well documented. For example, Kweon et al. (2017) find that Washington DC schools with more trees had a higher percentage of proficient or advanced standardized test scores. However, despite vast amounts of correlational research, there exists scant causal evidence outside of the laboratory setting. Children exposed to more green vegetation show enhanced cognitive development and higher scores on cognitive development tests (Dadvand et al., 2015). Green environments, such as open spaces with big trees, are even related to reduced symptoms of ADD and ADHD (Faber Taylor and Kuo, 2009).

These improved outcomes may operate through several different channels. Urban tree cover is known to mitigate traffic-related air pollution (Nowak et al., 2006). Trees may also reduce noise, allowing students to better focus (Gidlöf-Gunnarsson and Öhrström, 2007). There are also psychological benefits associated with increased tree cover in neighborhoods and surrounding schools. Li and Sullivan (2016) found that students who had views of trees and green environment from their classrooms, as compared to being in a room without windows or a room with a view of a brick wall, scored substantially higher on tests measuring attention. Studies consistently show a positive relationship between natural landscapes and enhanced physical activity amongst younger students (Dyment and Bell, 2007).

While causal studies detailing the effects of tree cover on educational outcomes are difficult to find, there exists a significant literature on the education impacts of ecosystem disturbances more broadly. High temperature has consistently been shown to affect learning (e.g. Park, 2022; Park et al., 2020a). For example, Park et al. (2020b) show that hotter school days in the years leading up to the PSAT reduce scores, with extreme heat being particularly damaging. Pollution has also been shown to reduce performance on tests. Wen and Burke (2022) find that wildfire smoke exposure in the year leading up to a standardized test negatively affects performance.

3.3 Data and descriptive statistics

3.3.1 Emerald ash borer survey

While infestations expand over relatively short distances through natural dispersal, infestations may go undetected for long periods. Long-distance spread occurs when infested material such as nursery stock or firewood is moved, potentially spurring new satellite populations. Ash borer infestation is fatal to all ash trees, however, it is difficult to detect until a tree is extensively damaged by the ash borer and begins to show symptoms.

After the first ash borer detections in 2006, the Illinois Department of Agriculture (IDA) initiated survey efforts to determine the extent of ash borer spread. The IDA survey consisted of destructive bark peeling of selected trees. Selected trees were generally 4-8 diameters in width and in areas of easy and clear right-of-way access, with efforts to sample 1 tree per 4 square miles. The Chicago metropolitan region was the priority, but other parts of the state were surveyed as well, particularly if an infestation was suspected. Initially, damage was minimal as the detection method results were mostly negative, but positive finds became more and more prevalent. Ultimately, the state stopped survey efforts in 2015, as ash borer spread had become extensive. Figure 3.1 displays the locations of confirmed ash borer infestations by year throughout the seven-county Chicago metropolitan region.

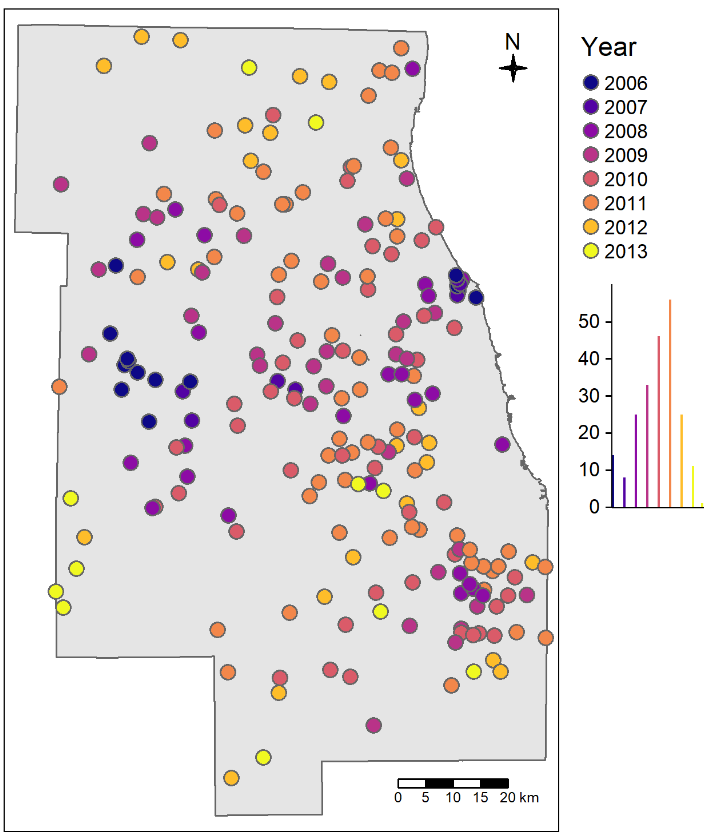

Figure 3.1: This map shows the location and timing of confirmed emerald ash borer infestations within the Chicago metropolitan region based on the IDA and USDA survey efforts.

3.3.2 Chicago metropolitan area tree cover

I use two main data sources to understand how ash borer affected tree cover in the Metropolitan Chicago region. I first use annual 30m resolution tree loss and tree gain maps developed in McCabe et al. (2018). These maps span 2000 to 2015 and importantly, allow me to differentiate changes in tree cover gain from tree cover loss. In order to better understand overall canopy cover trends and impacts, I also use annual 30m maps of canopy cover spanning from 1990 to 2017 developed in Hooper and Kennedy (2018). These yield the random forest probability that a pixel is canopy cover in a given year.

The biggest concern with both of these data products is low resolution. Many ash trees in metropolitan areas are street trees, the loss of which may not manifest in a 30m resolution product. This fact makes the Hooper and Kennedy (2018) product particularly valuable, because while small canopy cover changes may not manifest in a binary canopy cover classification, they are likely reflected in the latent probability metric. Imposing a forest classification threshold would needlessly sacrifice information on possible canopy cover changes, giving this product an advantage over typical binary canopy cover classification products.

3.3.3 Education and test score data

Education data come from the Illinois State Board of Education (ISBE), which release annual reports on the performance of each school in the state on a number of metrics. I geocode locations of each public school in the state of Illinois using addresses provided by ISBE. This allows me to know the location of each school relative to confirmed ash borer infestations. I take advantage of two primary outcomes: 1) attendance rates; and 2) standardized test performance.

The main standardized test I use is the Illinois Standards Achievement Test (ISAT).

The ISAT was instituted for the purpose of identifying failing schools, and students were tested in reading and math from grades 3–8. ISBE reported composite school-level performance on the test between 2003 and 2014, when the ISAT was retired. These data include the percentage of students who meet or exceed standards set by the state for a given year. Many schools also separately report these statistics across low vs. non-low income students for each grade.

ISBE also reports annual descriptive statistics about each school. For example, these data include low-income students, english-as-a-second language students, racial demographics, and enrollment. I use these as controls in my difference-in-differences estimation strategy.

3.3.4 Descriptive statistics

The association between tree cover, income, and test scores has been well documented. Often times, students in relatively higher income areas outperform their lower income counterparts in low-income areas. Additionally, income and tree cover often have a positive relationship in urban and metropolitan settings (Gerrish and Watkins, 2018). Figure 3.2 shows this to be the case in metropolitan Chicago. Canopy cover in the neighborhoods surrounding schools is positively correlated with performance on the ISAT and negatively correlated with the share of the student population that is low-income.

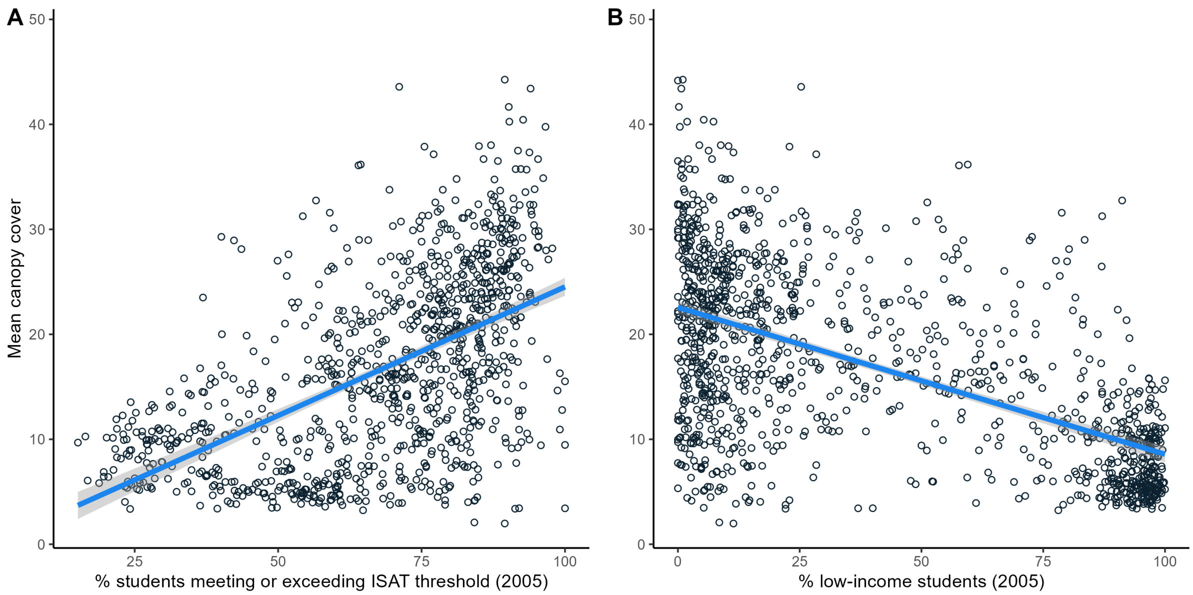

Figure 3.2: This figure shows how canopy cover surrounding schools in the Chicago metropolitan area is correlated with A) the percent of students who met or exceeded the ISAT benchmark; and B) the share of the student population classified as low-income. Both plots show canopy cover and education data for the year 2005, the year prior to the arrival of the emerald ash borer to the region.

In order to learn about patterns of ash borer infestation, I explore how schools with nearby ash borer infestations differ from those without. Figure 3.3 shows that prior to the arrival of the ash borer to the region, areas that would have confirmed infestations tended to have slightly greater levels of canopy cover. The second and third panels show that schools near an infestation site had slightly higher ISAT performance and a lower share of low-income students on average.

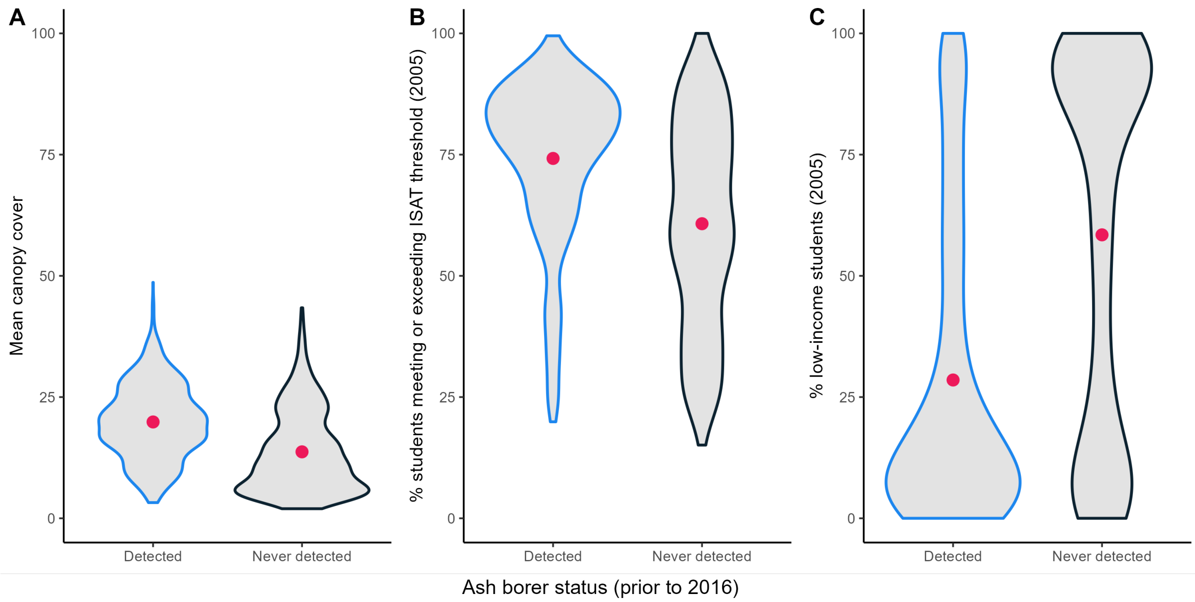

Figure 3.3: This figure shows how A) mean canopy cover; B) the percent of students who met or exceeded the ISAT benchmark; and C) the share of the student population classified as low-income, differed between schools with a confirmed emerald ash borer infestation site within 2 miles of the school prior to 2015 and those without.

3.4 Empirical strategy

3.4.1 Difference-in-differences methods

Given the idiosyncratic nature of ash borer spread through insect flight or accidental transportation, the timing of confirmed detections can be thought of as quasi-random. These quasi-random infestation confirmations provide me with a shock to tree communities and canopy cover. I define treatment status using confirmed infestations from the IDA survey as described previously. An IDA confirmed infestation indicates that not only are trees in the vicinity infested and ultimately likely to die, but that community officials are aware of the need for tree removal and replacement. Because infestations are often hard to detect and trees typically do not die from infestation immediately, a confirmed infestation can be thought of as an exogenous event that spurs new removal and replacement activity.

I use difference-in-differences (DID) approaches to evaluate the contemporaneous and dynamic effects of ash borer infestation. Identification relies on a common trends assumption, which amounts to assuming that outcomes in the treatment group would have followed the same evolution as those in the control and not yet treated groups, had treatment never occurred.

Recent papers have shown that the typical two-way fixed effects estimator may generate biased results in the presence of treatment effect heterogeneity (e.g. Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020). The estimator proposed in Callaway and Sant'Anna (2020) computes each 2x2 group-time treatment effect ($ATT_{g,t}$) individually, before aggregating them with intuitive weights. This estimator also allows causal interpretation to rely on a conditional common trends assumption, meaning that common trends need only hold after conditioning on relevant pre-treatment covariates.

Limitations to the DID approach arise from the distinction between confirmed ash borer infestation and ash borer presence. It is possible, and perhaps even likely, that ash borer are present in the area for some amount of time prior to confirmed infestation. Further, nearby infestations may go unreported if survey efforts are stopped near already confirmed infested areas. If this were the case, some infested areas may be incorrectly classified as untreated. In the context of my DID approaches, this should underestimate the impacts of the ash borer.

3.5 Results

3.5.1 Region tree cover impacts

I begin by examining the impact of ash borer infestation on canopy cover outcomes. This establishes the pathway through which ash borer detection is likely to affect education outcomes. In other words, hypothesized mechanisms by which we expect infestations to affect education outcomes operate first and foremost through a change in tree cover.

There are two main channels through which ash borer detection may result in tree loss. The first is through the ash borer directly, as an ash tree will die between one and four years following infestation, depending on the size and health of the individual ash tree. The second is through intentional removal of infested trees. Many communities declare any confirmed infested tree a public nuisance and require that the tree be removed (e.g. Macomb EAB Readiness Plan, 2007). As such, a confirmed infestation is likely to lead to manual removal of infested or dead trees in the vicinity.

Because removed trees are likely replaced with new trees, some may question why canopy cover loss should be observed on an annual scale. Most replacement trees are unlikely to generate the canopy cover of large healthy trees. A newly planted tree may have trunk diameter as small as two inches and provide little to no canopy cover, while a large healthy ash may grow 60 feet tall and 25 to 40 feet wide (Morton Arboretum). Ash borer detection may also lead to lower levels of tree cover gain. Removal of damaged or dead trees is costly to individuals and communities. Estimates suggest a cost near $1,000 to remove and replant a single tree. More broadly, ash borer had a massive impact on forestry budgets across the United States (Hauer and Peterson, 2017). While budgets in states with confirmed ash borer infestation saw sizable increases in tree removal budgets relative to non-ash borer confirmed states, budgets for tree planting did not change. Because most removed trees are likely replaced, it is plausible that trees that would've been planted in the absence of the pest were never established. In other words, funding allocated for establishing new tree cover instead went toward replacement of ash borer infested trees.

Table 3.1 shows estimates of the ash borer's impact on tree cover outcomes using the difference-in-differences approach described above. The unit of analysis is the grid cell, where a grid cell is considered treatment in the years after an infestation is first confirmed. My preferred grid cell size is 5×5 km, but canopy cover impacts are robust to grid cell sizes anywhere from 1 to 5 km (Section A.3.2). Confirmed infestation led to an estimated 1.4% decline in canopy cover probability within treated grid cells on average. The second and third columns further decompose this effect, exploring how infestation affected rates of tree cover gain and loss. Confirmed infestations led to a loss of nearly half an acre of tree cover per year (a 24.5% increase in rates of tree cover loss) in affected areas. We also see that infestations led to an 18.6% decline in rates of tree cover gain following infestation detection. These results align with perceptions that affected areas saw noticeable declines in mature trees and plantings. Table A.2 in Section A.3.5 shows that areas with a greater share of ash trees as the total tree population saw greater canopy cover declines, further supporting ash borer infestation as the driver of canopy cover loss in affected areas.

3.5.2 School-level impacts

While the above section documented the effects of ash borer infestations on the tree population in the broader Chicago metropolitan region, the main goal of this paper is to identify effects of the ash borer on education outcomes. In order to understand how ash borer induced tree cover changes influenced school-level outcomes, I use the individual school as the unit of analysis. I control for factors including racial, language, and income

Table 3.1: This table shows difference-in-differences estimates of the impact of ash borer infestation on tree cover outcomes across the Chicago metropolitan region. All estimates are based on the Callway and Sant'anna (2020) estimator and use both not-yet-treated and never-treated grid cells in the control group.

	Outcome		
	Canopy	Loss (Acres/year)	Gain (Acres/year)
ATT	-0.292***	0.467*	-1.449**
	(0.088)	(0.253)	(0.582)
Pre-treat mean	20.800	1.904	10.787
N grid cells	466	466	466

Note:
* $p<0.1$, ** $p<0.05$, *** $p<0.01$; standard errors clustered at grid level

characteristics related to the student population and school enrollment. A school is considered treated in the years after an infestation is detected within 3.22km (2 miles) of the school. This distance was selected based on guidance from the USDA Animal Plant Health and Inspection Service. In the emerald ash borer program manual (USDA, 2020), surveyors are given the following guidance after confirmed detection of ash borer infestation: "After detecting adult emerald ash borer in traps or finding an infested tree, conduct a visual survey until symptomatic trees are no longer found. Continue visual survey for a distance of two miles beyond the initial trap capture or infested tree detection".

Estimates of the ash borer's tree cover impacts within the 3.22km (2 miles) surrounding schools are displayed in Table 3.2. We see that confirmed infestations again lead to a decline in canopy cover, an increase in rates of tree cover loss, and a decrease in rates of tree cover gain, this time in the direct vicinity of schools. These results indicate that students at schools exposed to ash borer infestations observe salient tree cover impacts in the vicinity of their school. These tree cover impacts likely affect the provision of ecosystem services such as pollution and high-temperature mitigation to students at exposed schools.

Table 3.2: This table shows difference-in-differences estimates of the impact of ash borer infestation on tree cover outcomes within 3.22km (2 miles) of the school. All estimates are based on the Callway and Sant'anna (2020) estimator and use both not-yet-treated and never-treated schools in the control group.

	Outcome		
	Canopy	Loss (Acres/year)	Gain (Acres/year)
ATT	-0.112***	0.431***	-1.872***
	(0.042)	(0.158)	(0.412)
Pre-treat mean	20.365	2.346	10.076
N schools	2232	2232	2232

Note:
* $p<0.1$, ** $p<0.05$, *** $p<0.01$; standard errors clustered at school level

Table 3.2 now shows estimates of confirmed ash borer infestation on several education outcomes of interest. The first column indicates that confirmed ash borer infestation within 3.22km (2 miles) reduces the percentage of students at the school that meet or exceed the ISAT benchmark on average. This amounts to a 1.86% reduction in the share of students meeting or exceeding the benchmark for the characteristic school. The second column indicates a similar impact on all standardized tests more broadly. The primary additional test included in this metric is the PSAT. For comparison, Park et al. (2020b) find that a 1°F hotter school year reduces learning by 1 percent, with impacts disproportionately concentrated among low-income students. While the outcome in Park et al. (2020b) is not identical to mine in this setting, the results suggest that ash borer infestation may have an impact comparable to 1°F hotter temperatures.

While overall attendance does not appear to be affected by ash borer infestation, attendance rates among low-income students do appear to decline slightly. Low-income student attendance rates decline by 0.3%. Lastly, as a falsification test, I examine whether infestation affects enrollment contemporaneously. The final column of Table 3.3 indicates that enrollment is unaffected by confirmed infestation.

Table 3.3: This table shows difference-in-differences estimates of the impact of ash borer infestation on school-level education outcomes. All estimates are based on the Callway and Sant'anna (2020) estimator and use both not-yet-treated and never-treated schools in the control group.

	Outcome				
	ISAT composite	All tests	Attendance rate	Low-income attend.	Enrollment
ATT	-1.455***	-1.437***	-0.1	-0.283**	3.967
	(0.301)	(0.316)	(0.069)	(0.121)	(3.743)
Pre-treat mean	78.216	74.879	94.958	94.240	496.872
N schools	2232	2232	2232	2232	2232

Note:

* $p<0.1$, ** $p<0.05$, *** $p<0.01$; standard errors clustered at school level

Event study estimates in Figure 3.4 show how treatment effects vary through time for school-level A) Canopy cover; and B) composite ISAT performance. In both plots, confirmed infestation does not have any effect on either outcome prior to the actual date of confirmation. This lends credence to the conditional common trends assumption on which causal interpretation of the difference-in-differences estimates relies. The canopy cover impacts appear to stabilize 5 years post-infestation confirmation. Test score impacts on the other hand, are sustained through the 7 year period post infestation confirmation. This suggests that the impacts are not due to factors associated with detection, such as noise caused by initial tree removal or increased presence of arborists. Instead, sustained increases in temperature, pollution, or pesticide use are more plausible mechanisms.

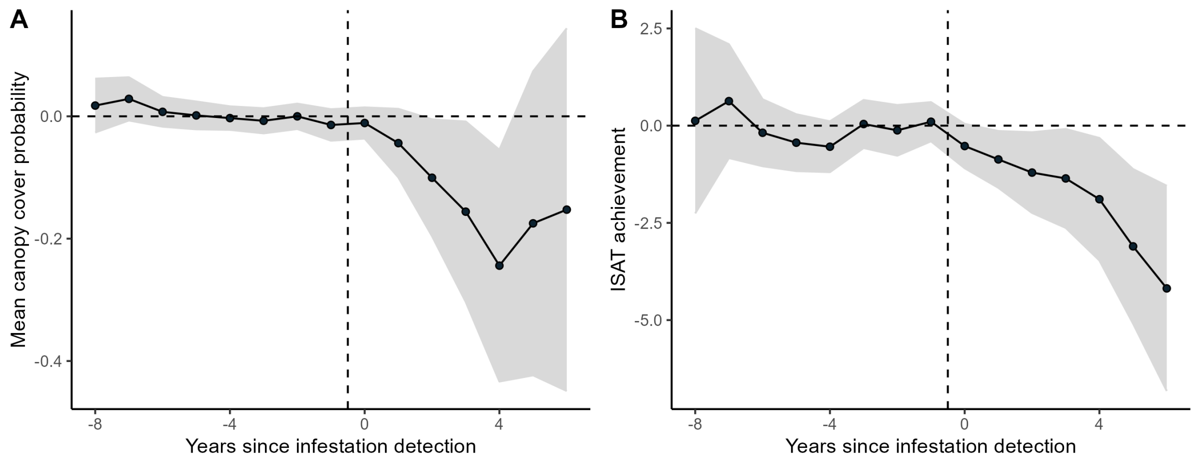

Figure 3.4: This figure shows how estimates of infestation impacts vary across event time for two outcomes: A) canopy cover; and B) the percent of students who met or exceeded the ISAT benchmark. We see declines in both outcomes after confirmed infestation detection within 3.22km (2 miles) of a school. There are no effects of infestation prior to the actual date of detection.

3.6 Heterogeneity

3.6.1 Impacts concentrated on low-income elementary students

ISAT performance is also provided by ISBE stratified by grade level, subject, low versus non-low income, and benchmark categorizations. Figure 3.5 shows difference-in-differences estimates across subsets of the data along these dimensions, where each of the four plots represent the percent of students performing within one of the following benchmarks: A) Academic warning; B) Below standards; C) Meets standards; and D) exceeds standards.

The most notable pattern to emerge from Figure 3.5 is the disparity along the low versus non-low income dimension. Impacts are consistently seen across each of the four benchmark measures for low-income students (highlighted in blue) at exposed schools. In contrast, non-low income students are not clearly affected at all.

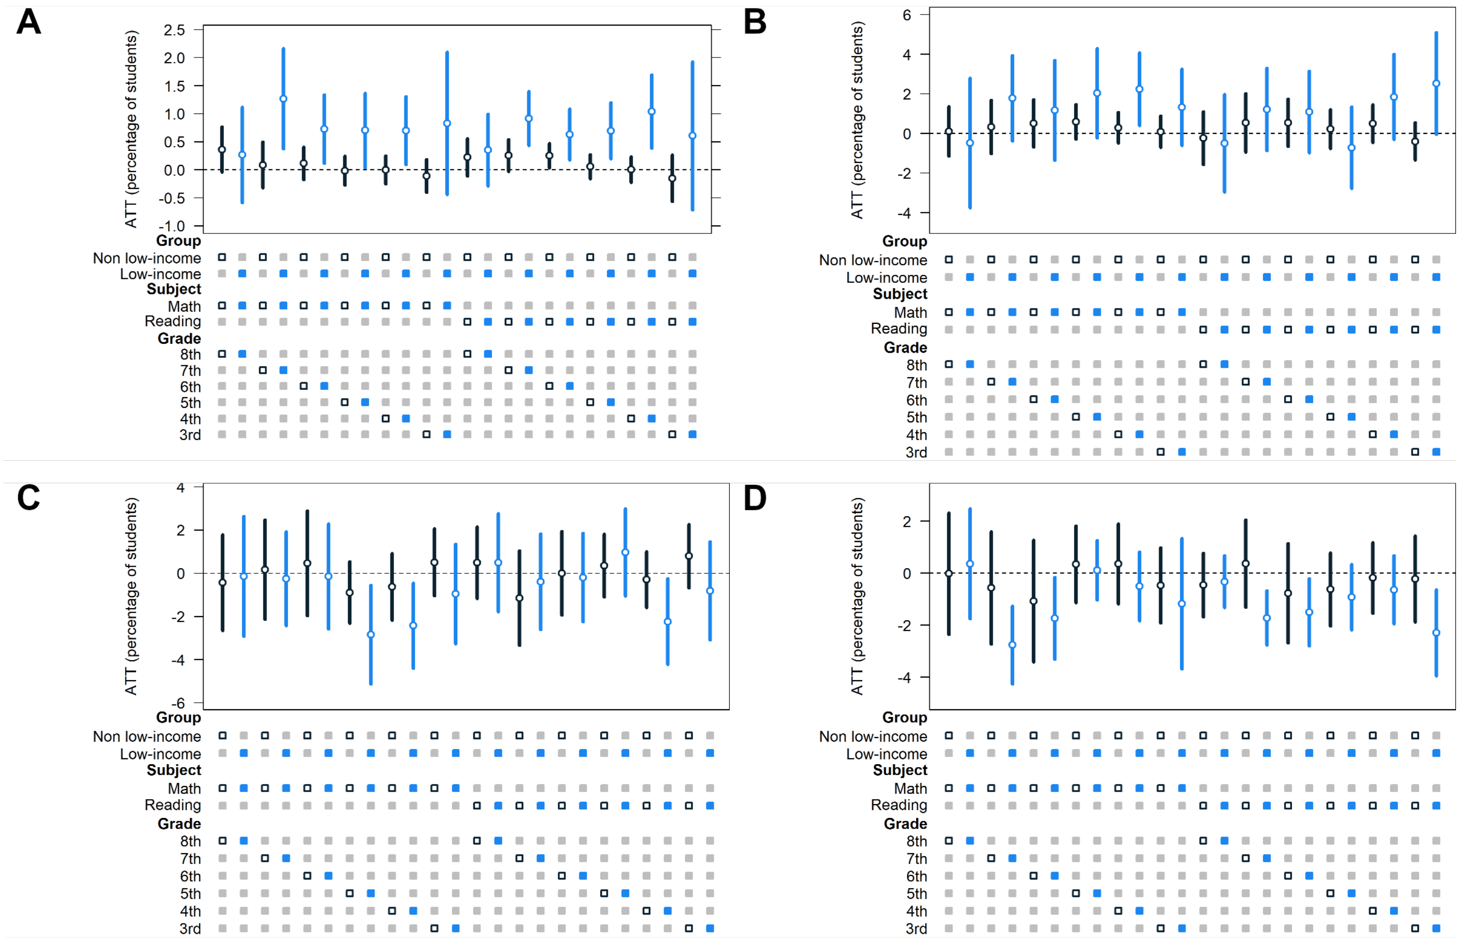

Figure 3.5: Difference-in-differences estimates of the impact of ash borer infestation on ISAT performance. Outcomes represent: A) percent of students with academic warning; B) percent of students below benchmark; C) percent of students meeting benchmark; and D) percent of students exceeding benchmark.

3.7 Conclusion

This study estimates the education impacts of a quasiexperimental shock to tree cover induced by the invasive emerald ash borer. I first address the causal channel through which student outcomes are impacted by showing that ash borer altered tree cover dynamics in affected areas. Consistent evidence suggests that infestations led to declines in overall canopy cover and that this was due to both increased rates of tree loss and declines in tree gain. I then demonstrate that infestation led to reduced test performance at schools in the immediate vicinity of infestations. Further, these education impacts are concentrated almost entirely within the low-income student population at impacted schools

This work presented here contributes to our understanding of the social impacts of human-induced environmental change, particularly the introduction of invasive species. I also make important advances to the literature on the benefits of vegetation and trees to human well-being. Many studies have sought to establish the link between tree cover and education outcomes, however, there exists scant causal evidence. The staggered and idiosyncratic spread of the ash borer provides a setting in which causality can be plausibly established. As such, my work is some of the first to use observational data to causally address the tree cover-education link.

Appendix A

Appendix

A.1 Chapter 1 appendix

A.1.1 Setup for analytical results

Let y_{it} be the binary outcome of interest for individual unit i at time t. We assume that researchers have access to outcome data pre-treatment ($t = 1$) and post-treatment ($t = 2$). Some units ($D_i = 1$) are exposed to a policy treatment in the second time period ($t_0 = 2$ denotes the time of first treatment for treated points). Let $W_{it} = 1$ if unit i is treated before time t and $W_{it} = 0$ otherwise. Using the potential outcome notation, denote $y_{it}(0)$ the outcome of unit i at time t if it does not receive treatment by time t and $y_{it}(1)$ the outcome for the same unit if it receives treatment.

Thus, the realized outcome for unit i at time t is

$$y_{it} = W_{it}y_{it}(1) + (1 - W_{it})y_{it}(0)$$

.

The parameter of interest, the ATT is defined:

$$ATT = E[y_{it}(1) - y_{it}(0)|D_i = 1]$$

Make the following common trends assumption, under which we evaluate these methods:

$$E[y_{i2}(0) - y_{i1}(0)|D_i = 1] = E[y_{i2}(0) - y_{i1}(0)|D_i = 0]$$

Lastly, define C_i as the first year in which an irreversible event of interest (e.g., deforestation) is realized for individual unit i and suppose y_{it} is not observable when $t > C_i$.

A.1.2 TWFE regression models with point fixed effects do not identify ATT

In settings with a binary and unrepeatable outcome variable, the commonly used unit-level TWFE model yields the post-treatment difference in outcomes (single difference), rather than the desired ATT.

We define the observed outcome y^o_{it}:

$$y^o_{it} = \begin{cases} 1 & t = C_i \\ 0 & t < C_i \\ - & t > C_i \end{cases} \tag{A.1}$$

, where $y^o_{it} = -$ indicates that the outcome for pixel i has been dropped from the panel in time t.

Lastly, define the traditional individual unit level TWFE regression:

$$y^o_{it} = \alpha + \beta_{TWFE} \text{ x } W_{it} + \gamma_i + \lambda_t + u_{it}$$

, where

- γ_i indicate point fixed effects

- λ_t indicate year fixed effects

In the 2x2 case, we can write

$$y^o_{i1} = \alpha + \gamma_i + u_{i1}$$

and

$$y^o_{i2} = \begin{cases} \alpha + \beta_{TWFE} \text{ x } D_i + \gamma_i + \eta_{t=2} + u_{i2} & y^o_{i1} = 0 \\ - & y^o_{i1} \neq 0 \end{cases}$$

, where $\eta_{t=2}$, an indicator for the post-treatment period, subsumes λ_t. Note that we substituted W_{it} for D_i, since the two are equivalent post-treatment.

In the 2x2 case, the TWFE estimator is equivalent to the first differences estimator, and yields:

$$y^o_{i2} - y^o_{i1} = \begin{cases} (\alpha + \beta_{TWFE} \text{ x } D_i + \gamma_i + \eta_{t=2} + u_{i2}) - (\alpha + \gamma_i + u_{i1}) & y^o_{i1} = 0 \\ - & y^o_{i1} \neq 0 \end{cases}$$

Focusing on the first case, where $y_{i1}^o = 0$

$$\begin{aligned} y_{i2}^o - y_{i1}^o &= (\alpha + \beta_{TWFE} \text{ x } D_i + \gamma_i + \eta_{t=2} + u_{i2}) - (\alpha + \gamma_i + u_{i1}) \\ &= \beta_{TWFE} \text{ x } D_i + \eta_{t=2} + \Delta u_i \end{aligned}$$

The general expression can be restated as:

$$y_{i2}^o - y_{i1}^o = \begin{cases} \beta_{TWFE} \text{ x } D_i + \eta_{t=2} + \Delta u_i & y_{i1}^o = 0 \\ - & y_{i1}^o \neq 0 \end{cases}$$

With binary treatment (D_i), $\hat{\beta}_{TWFE}$, the regression's estimate of β_{TWFE} can be expressed as the double difference in mean outcomes across treated / untreated units, and across the two time periods:

$$\hat{\beta} = \frac{1}{n_{i:D_i=1}} \sum_{i:D_i=1} y_{i2}^o - \frac{1}{n_{i:D_i=1}} \sum_{i:D_i=1} y_{i1}^o - (\frac{1}{n_{i:D_i=0}} \sum_{i:D_i=0} y_{i2}^o - \frac{1}{n_{i:D_i=0}} \sum_{i:D_i=0} y_{i1}^o)$$

However, this is only valid when $y_{i1}^o = 0$. As a result, we can restate as:

$$\begin{aligned} \hat{\beta}_{TWFE} &= \frac{1}{n_{i:D_i=1}} \sum_{i:D_i=1} y_{i2}^o - 0 - (\frac{1}{n_{i:D_i=0}} \sum_{i:D_i=0} y_{i2}^o - 0) \\ &= \frac{1}{n_{i:D_i=1}} \sum_{i:D_i=1} y_{i2}^o - \frac{1}{n_{i:D_i=0}} \sum_{i:D_i=0} y_{i2}^o \end{aligned}$$

Thus far, we have shown that $\hat{\beta}_{TWFE}$ is equal to the ex-post difference in means between treatment and control units.

We now examine what this means for estimating the parameter of interest, the ATT.

Applying the potential outcomes notation to indicate whether we see the treated or untreated outcome:

$$\hat{\beta}_{TWFE} = \frac{1}{n_{i:D_i=1}} \sum_{i:D_i=1} y^o_{i2}(1) - \frac{1}{n_{i:D_i=0}} \sum_{i:D_i=0} y^o_{i2}(0)$$

Adding and subtracting $\frac{1}{n_{i:D_i=1}} \sum_{i:D_i=1} y^o_{i2}(0)$ gives:

$$\begin{aligned} \hat{\beta}_{TWFE} = & \frac{1}{n_{i:D_i=1}} \sum_{i:D_i=1} y^o_{i2}(1) - y^o_{i2}(0) \\ & + \frac{1}{n_{i:D_i=1}} \sum_{i:D_i=1} y^o_{i2}(0) - \frac{1}{n_{i:D_i=0}} \sum_{i:D_i=0} y^o_{i2}(0) \end{aligned}$$

Taking the expectation gives:

$$E[\hat{\beta}_{TWFE}] = ATT + E[y_{i2}(0)|D_i = 1] - E[y_{i2}(0)|D_i = 0]$$

, where the expectation of the y^o_{ivt}s is equal to that of the y_{ivt}s if they are i.id..

$$\beta_{TWFE} = ATT + E[y_{i2}(0)|D_i = 1] - E[y_{i2}(0)|D_i = 0]$$

Now, imposing the common trends assumption, substituting for $E[y_{i2}(0)|D_i = 1]$, and simplifying:

$$\beta_{TWFE} = ATT + E[y_{i1}(0)|D_i = 1] - E[y_{i1}(0)|D_i = 0]$$

A.1.3 Cox PH DID identifies $HRTT$ when proportional trends assumption holds

Consider the cox proportional hazards model of the censored y_{it} regressed on the treatment dummy, D_i, the post dummy, $\mathbb{1}\{t \geq t_0\}$, and their interaction:

$$h(t) = \delta_0(t) exp(\alpha_0 + \alpha_1 D_i + \alpha_2 \mathbb{1}\{t \geq t_0\} + \beta_{coxDID} \text{ x } D_i \mathbb{1}\{t \geq t_0\} + \epsilon_{it})$$

, where $h(t)$ is the hazard rate of deforestation, t years into the study period; and $\delta_0(t)$ is the baseline hazard function.

The exponentiated coefficient on the interaction between two binary variables, D_i and $\mathbb{1}\{t \geq t_0\}$, $exp(\beta_{coxDID})$, is expressed as the ratio of the two pre-post hazard rate ratios across the two groups:

$$exp(\beta_{coxDID}) = \frac{E[y_{i2}|D_i = 1]/E[y_{i1}|D_i = 1]}{E[y_{i2}|D_i = 0]/E[y_{i1}|D_i = 0]} \tag{A.2}$$

Introducing potential outcomes and simplifying:

$$exp(\beta_{coxDID}) = \frac{E[y_{i2}(1)|D_i = 1]E[y_{i1}(0)|D_i = 0]}{E[y_{i2}(0)|D_i = 0]E[y_{i1}(0)|D_i = 1]} \tag{A.3}$$

Now, operating under Assumption 1 (Proportional Trends) and substituting for the right-hand side of (7):

$$exp(\beta_{coxDID}) = \frac{E[y_{i2}(1)|D_i = 1]E[y_{i1}(0)|D_i = 1]}{E[y_{i2}(0)|D_i = 1]E[y_{i1}(0)|D_i = 1]} \tag{A.4}$$

, showing that under proportional trends, $exp(\beta_{coxDID}) = HRRT$

A.1.4 Keeping pixels in periods after they are first deforested is not a viable solution

Remotely sensed metrics of deforestation at the pixel level are often subject to the dynamics of forest disturbance and regrowth. After a deforestation event occurs, the deforested area is unlikely to revert to forest cover within the study period, as it takes several years for forest to regenerate to a detectable level. Further, many data products do not allow for the monitoring of forest regrowth. In the panel therefore, it is likely that in the periods after a pixel is first realized as deforested, subsequent observations of the pixel will also observe the pixel as deforested.

The logic for dropping binary pixels after they first become deforested is as follows. A forested pixel switches from its assigned value of 0 to a value of 1 following a discrete deforestation event. Keeping the deforested pixel in the panel beyond the first period in which it was observed as deforested may imply that it has actively been deforested in each subsequent time period. In fact, no new deforestation event has ocurred, but the area simply remains deforested from the prior event. These pixels, therefore, contribute positively towards the deforestation rate in each period they are left in the panel. As such, the coefficient cannot recover the ATT.

Define an alternative observed outcome y_{it}^{alt}, where rather than dropping units from the panel, the outcome is imputed as a 1 when $t > C_0$.

$$y_{it}^{alt} = \begin{cases} 0 & t < C_0 \\ 1 & t \geq C_0 \end{cases} \tag{A.5}$$

Now, consider the difference-in-differences estimand with respect to outcome y_{it}^{alt}:

$$\psi = E[y_{i2}^{alt}|D_i = 1] - E[y_{i1}^{alt}|D_i = 1] - (E[y_{i2}^{alt}|D_i = 0] - E[y_{i1}^{alt}|D_i = 0])$$

Because y_{it}^{alt} is binary, psi can be re-expressed using probabilities:

$$\psi = P(y_{i2}^{alt} = 1|D_i = 1) - P(y_{i1}^{alt} = 1|D_i = 1) - (P(y_{i2}^{alt} = 1|D_i = 0) - P(y_{i1}^{alt} = 1|D_i = 0))$$

Because we observe $y_{it}^{alt} = 1$ anytime after C_0, we can express the probability that $y_{it}^{alt} = 1$ as a function of the probability that C_0 occurred prior to time t. Then, the terms with y_{i2}^{alt}s are re-expressed.

$$\begin{aligned} \psi = &P(y_{i2} = 1|D_i = 1) \cup P(y_{i1} = 1|D_i = 1) - P(y_{i1} = 1|D_i = 1) - \\ &(P(y_{i2} = 1|D_i = 0) \cup (P(y_{i1} = 1|D_i = 0) - P(y_{i1} = 1|D_i = 0)) \end{aligned}$$

Applying the potential outcomes notation:

$$\begin{aligned}\psi =&P(y_{i2}(1) = 1|D_i = 1) \cup P(y_{i1}(0) = 1|D_i = 1) - P(y_{i1}(0) = 1|D_i = 1)- \\ &(P(y_{i2}(0) = 1|D_i = 0) \cup (P(y_{i1}(0) = 1|D_i = 0) - P(y_{i1}(0) = 1|D_i = 0))\end{aligned}$$

The bias that results from using y_{it}^{alt} is expressed as the difference between the ATT and ψ. After re-expressing the ATT in terms of probability:

$$\begin{aligned}\psi - ATT =&P(y_{i2}(1) = 1|D_i = 1) \cup P(y_{i1}(0) = 1|D_i = 1) - P(y_{i1}(0) = 1|D_i = 1)- \\ &(P(y_{i2}(0) = 1|D_i = 0) \cup (P(y_{i1}(0) = 1|D_i = 0) - P(y_{i1}(0) = 1|D_i = 0)) \\ &-[P(y_{i2}(1) = 1|D_i = 1) - P(y_{i2}(0) = 1|D_i = 1)]\end{aligned}$$

Under the common trends assumption and substituting for the unobserved term, $P(y_{i2}(0) = 1|D_i = 1)$, yields

$$\begin{aligned}\psi - ATT =&P(y_{i2}(1) = 1|D_i = 1) \cup P(y_{i1}(0) = 1|D_i = 1) - P(y_{i1}(0) = 1|D_i = 1)- \\ &(P(y_{i2}(0) = 1|D_i = 0) \cup (P(y_{i1}(0) = 1|D_i = 0) - P(y_{i1}(0) = 1|D_i = 0)) \\ &- [P(y_{i2}(1) = 1|D_i = 1) - (P(y_{i2}(0) = 1|D_i = 0) - P(y_{i1}(0) = 1|D_i = 0) \\ &+ P(y_{i1}(0) = 1|D_i = 1))] \\ =&P(y_{i2}(1) = 1|D_i = 1) \cap P(y_{i1}(0) = 1|D_i = 1) + P(y_{i2}(1) = 1|D_i = 1)- \\ &(P(y_{i2}(0) = 1|D_i = 0) \cap (P(y_{i1}(0) = 1|D_i = 0) + P(y_{i2}(0) = 1|D_i = 0)) \\ &- [P(y_{i2}(1) = 1|D_i = 1) - (P(y_{i2}(0) = 1|D_i = 0) - P(y_{i1}(0) = 1|D_i = 0) \\ &+ P(y_{i1}(0) = 1|D_i = 1))]\end{aligned}$$

, where the second equality uses the definition of union operator.

Simplifying now yields:

$$\begin{aligned}\psi - ATT =& P(y_{i2}(1) = 1|D_i = 1) \cap P(y_{i1}(0) = 1|D_i = 1) - \\ & (P(y_{i2}(0) = 1|D_i = 0) \cap (P(y_{i1}(0) = 1|D_i = 0) \\ & -[P(y_{i1}(0) = 1|D_i = 1) - P(y_{i1}(0) = 1|D_i = 0)]\end{aligned}$$

A.1.5 Analytical expression of non-random sample selection bias in two-period two-group setting

Consider again, the observed outcome, y^o_{it}. We begin with the DID estimand in the two-group, two-period case:

$$\begin{aligned}\beta_{DID} =& E[y^o_{ivt}|t \geq t_0, D_i = 1] - E[y^o_{ivt}|t < t_0, D_i = 1] \\ & - (E[y^o_{ivt}|t \geq t_0, D_i = 0] - E[y^o_{ivt}|t < t_0, D_i = 0])\end{aligned}$$

Now the bias generated due to non-random sample selection can be represented as the difference between this estimand and the ATT:

$$\begin{aligned}\beta_{DID} - ATT =& E[y^o_{ivt}|t \geq t_0, D_i = 1] - E[y^o_{ivt}|t < t_0, D_i = 1] \\ & - (E[y^o_{ivt}|t \geq t_0, D_i = 0] - E[y^o_{ivt}|t < t_0, D_i = 0]) \\ & - (E[y_{ivt}(1)|t \geq t_0, D_i = 1] - E[y_{ivt}(0)|t \geq t_0, D_i = 1])\end{aligned}$$

In the first period, the expectation of y^o_{ivt} is the same as that of y_{ivt}, giving:

$$\begin{aligned}\beta_{DID} - ATT =& E[y^o_{ivt}|t \geq t_0, D_i = 1] - E[y_{ivt}|t < t_0, D_i = 1] \\ & - (E[y^o_{ivt}|t \geq t_0, D_i = 0] - E[y_{ivt}|t < t_0, D_i = 0]) \\ & - (E[y_{ivt}(1)|t \geq t_0, D_i = 1] - E[y_{ivt}(0)|t \geq t_0, D_i = 1])\end{aligned}$$

Applying potential outcomes:

$$\begin{aligned}\beta_{DID} - ATT =& E[y^o_{ivt}(1)|t \geq t_0, D_i = 1] - E[y_{ivt}(0)|t < t_0, D_i = 1] \\ & - (E[y^o_{ivt}(0)|t \geq t_0, D_i = 0] - E[y_{ivt}(0)|t < t_0, D_i = 0]) \\ & - (E[y_{ivt}(1)|t \geq t_0, D_i = 1] - E[y_{ivt}(0)|t \geq t_0, D_i = 1])\end{aligned}$$

Applying our common trends assumption:

$$\begin{aligned}\beta_{DID} - ATT =& E[y^o_{ivt}(1)|t \geq t_0, D_i = 1] - E[y_{ivt}(0)|t < t_0, D_i = 0] \\ & - (E[y^o_{ivt}(0)|t \geq t_0, D_i = 0] - E[y_{ivt}(0)|t < t_0, D_i = 0]) \\ & - (E[y_{ivt}(1)|t \geq t_0, D_i = 1] - E[y_{ivt}(0)|t \geq t_0, D_i = 0])\end{aligned}$$

Simplifying:

$$\begin{aligned}\beta_{DID} - ATT =& E[y^o_{ivt}(1)|t \geq t_0, D_i = 1] - E[y^o_{ivt}(0)|t \geq t_0, D_i = 0] \\ & - (E[y_{ivt}(1)|t \geq t_0, D_i = 1] - E[y_{ivt}(0)|t \geq t_0, D_i = 0])\end{aligned}$$

Extending to our simulations:

In the context of our monte carlo simulations, this can be extended:

$$\begin{aligned}\beta_{DID} - ATT =& P[y^o_{ivt}(1) = 1 | t \geq t_0, D_i = 1] - P[y^o_{ivt}(0) = 1 | t \geq t_0, D_i = 0] \\ & - (P[y_{ivt}(1) = 1 | t \geq t_0, D_i = 1] - P[y_{ivt}(0) = 1 | t \geq t_0, D_i = 0])\end{aligned}$$

$$\begin{aligned}\beta_{DID} - ATT =& P[(y_{ivt}(1) = 1 | t \geq t_0, D_i = 1) | (y_{ivt}(0) = 0 | t < t_0, D_i = 1)] \\ & - P[(y_{ivt}(0) = 1 | t \geq t_0, D_i = 0) | (y_{ivt}(0) = 0 | t < t_0, D_i = 0)] \\ & - (P[y_{ivt}(1) = 1 | t \geq t_0, D_i = 1] - P[y_{ivt}(0) = 1 | t \geq t_0, D_i = 0])\end{aligned}$$

$$\begin{aligned}\beta_{DID} - ATT =& P[(y^*_{ivt}(1) > 0 | t \geq t_0, D_i = 1) | (y^*_{ivt}(0) \leq 0 | t < t_0, D_i = 1)] \\ & - P[(y^*_{ivt}(0) > 0 | t \geq t_0, D_i = 0) | (y^*_{ivt}(0) \leq 0 | t < t_0, D_i = 0)] \\ & - (P[y^*_{ivt}(1) > 0 | t \geq t_0, D_i = 1] - P[y^*_{ivt}(0) > 0 | t \geq t_0, D_i = 0])\end{aligned}$$

Here, we let $t \in \{1, 2\}$ denote the first and second periods, respectively:

$$
\begin{aligned}
\beta_{DID} - ATT = & P[(\beta_0 + \beta_1 + \beta_{2,1} + \beta_3 + \alpha_i + u_{i2} + \rho_v > 0 | t \geq t_0, D_i = 1) \\
& |(\beta_0 + \beta_1 + \alpha_i + u_{i1} + \rho_v \leq 0 | t < t_0, D_i = 1)] \\
& - P[(\beta_0 + \beta_{2,0} + \alpha_i + u_{i2} + \rho_v > 0 | t \geq t_0, D_i = 0) \\
& |(\beta_0 + \alpha_i + u_{i1} + \rho_v \leq 0 | t < t_0, D_i = 0)] \\
& - (P[\beta_0 + \beta_1 + \beta_{2,1} + \beta_3 + \alpha_i + u_{i2} + \rho_v > 0 | t \geq t_0, D_i = 1] \\
& - P[\beta_0 + \beta_{2,0} + \alpha_i + u_{i2} + \rho_v > 0 | t \geq t_0, D_i = 0])
\end{aligned}
$$

A.1.6 Further Monte Carlo evidence that individual unit-level TWFE is equivalent to coefficient from same regression on dataset without pixels deforested pre-treatment

Table A.1 shows coefficient estimates from the Monte Carlo setup described in the main text on altered datasets. It demonstrates that the coefficient of interest from Regression 2 is numerically equivalent to that from the same regression on a dataset where all pixels deforested in the first period are dropped completely. The estimated coefficient is not numerically equivalent to the ex-post difference in means, although this is true in the 2x2 case. This exercise provides further evidence that this commonly used TWFE regression does not use the pre-treatment variation in deforestation at all, which is necessary to recover the ATT in this setting.

Table A.1: TWFE with pixel fixed effects is numerically equivalent to TWFE on dataset with all pixels deforested pre-treatment dropped completely from the dataset

Model	**Bias**	**RMSE**	**0.25 to 0.75 quantile**
DID	0.0001034	0.0015268	-0.00088 , 0.0011
ex-post difference in means	0.0300441	0.0300609	0.02935 , 0.03074
TWFE	0.0359502	0.0359744	0.03501 , 0.03684
TWFE on dataset dropping deforested pixels prior to treatment	0.0359502	0.0359744	0.03501 , 0.03684

A.1.7 Initial Monte Carlo parameter to β coefficient mapping

The following five parameters and their definitions inform the simulation parameterizations.

$$
\begin{aligned}
baseline_0 &= E[y_{it}(0)|t < t_0, D_i = 0] \\
baseline_1 &= E[y_{it}(0)|t < t_0, D_i = 1] \\
trend_0 &= E[y_{it}(0)|t \geq t_0, D_i = 0] - E[y_{it}(0)|t < t_0, D_i = 0] \\
trend_1 &= E[y_{it}(0)|t \geq t_0, D_i = 1] - E[y_{it}(0)|t < t_0, D_i = 1] \\
ATT &= E[y_{it}(1) - y_{it}(0)|t \geq t_0, D_i = 1]
\end{aligned}
$$

Note the following constraints on the parameters:

$$
\begin{aligned}
E[y_{it}(0)|t \geq t_0, D_i = 0] \geq 0 \\
E[y_{it}(1)|t \geq t_0, D_i = 1] \geq 0
\end{aligned}
$$

The parameters can be expressed as follows:

$$
\begin{aligned}
ATT =& E[y_{it}(1) - y_{it}(0)|t \geq t_0, D_i = 1] \\
=& E[y_{it}(1)|t \geq t_0, D_i = 1] - E[y_{it}(0)|t \geq t_0, D_i = 1] \\
=& P(y_{it}(1) = 1|t \geq t_0, D_i = 1) - P(y_{it}(0) = 1|t \geq t_0, D_i = 1) \\
=& P(y^*_{it}(1) > 0|t \geq t_0, D_i = 1) - P(y^*_{it}(0) > 0|t \geq t_0, D_i = 1) \\
=& P(\beta_0 + \beta_1 + \beta_{2,1} + \beta_3 + \alpha_i + u_{it} > 0) - P(\beta_0 + \beta_1 + \beta_{2,1} + \alpha_i + u_{it} > 0) \\
=& P(-\alpha_i - u_{it} < \beta_0 + \beta_1 + \beta_{2,1} + \beta_3) - P(-\alpha_i - u_{it} < \beta_0 + \beta_1 + \beta_{2,1}) \\
=& F(\beta_0 + \beta_1 + \beta_{2,1} + \beta_3) - F(\beta_0 + \beta_1 + \beta_{2,1})
\end{aligned}
$$

$$
\begin{aligned}
trend_0 =& E[y_{it}(0)|t \geq t_0, D_i = 0] - E[y_{it}(0)|t < t_0, D_i = 0] \\
=& P(y_{it}(0) = 1|t \geq t_0, D_i = 0) - P(y_{it}(0) = 1|t < t_0, D_i = 0) \\
=& P(y^*_{it}(0) > 0|t \geq t_0, D_i = 0|y^*_{it}(0) < 0|t < t_0, D_i = 0) \\
& - P(y^*_{it}(0) > 0|t < t_0, D_i = 0) \\
=& \frac{(1 - P(y^*_{it}(0) > 0|t < t_0, D_i = 0))P(y^*_{it}(0) > 0|t \geq t_0, D_i = 0)}{(1 - P(y^*_{it}(0) > 0|t < t_0, D_i = 0))} \\
& - P(y^*_{it}(0) > 0|t < t_0, D_i = 0) \\
=& P(-\alpha_i - u_{it} < \beta_0 + \beta_{2,0}) - P(-\alpha_i - u_{it} < \beta_0) \\
=& F(\beta_0 + \beta_{2,0}) - F(\beta_0)
\end{aligned}
$$

$$
\begin{aligned}
trend_1 =& E[y_{it}(0)|t \geq t_0, D_i = 1] - E[y_{it}(0)|t < t_0, D_i = 1] \\
=& P(y_{it}(0) = 1|t \geq t_0, D_i = 1) - P(y_{it}(0) = 1|t < t_0, D_i = 1) \\
=& P(y^*_{it}(0) > 0|t \geq t_0, D_i = 1 \cap y^*_{it}(0) < 0|t < t_0, D_i = 1) \\
& - P(y^*_{it}(0) > 0|t < t_0, D_i = 1) \\
=& P(-\alpha_i - u_{it} < \beta_0 + \beta_1 + \beta_{2,1}) - P(-\alpha_i - u_{it} < \beta_0 + \beta_1) \\
=& F(\beta_0 + \beta_1 + \beta_{2,1}) - F(\beta_0 + \beta_1)
\end{aligned}
$$

$$
\begin{aligned}
baseline_0 =& E[y_{it}(0)|t < t_0, D_i = 0] \\
=& P(y_{it}(0) = 1|t < t_0, D_i = 0) \\
=& P(y^*_{it}(0) > 0|t < t_0, D_i = 0) \\
=& P(-\alpha_i - u_{it} < \beta_0) \\
=& F(\beta_0)
\end{aligned}
$$

$$
\begin{aligned}
baseline_1 =& E[y_{it}(0)|t < t_0, D_i = 1] \\
=& P(y_{it}(0) = 1|t < t_0, D_i = 1) \\
=& P(y^*_{it}(0) > 0|t < t_0, D_i = 1) \\
=& P(-\alpha_i - u_{it} < \beta_0 + \beta_1) \\
=& F(\beta_0 + \beta_1)
\end{aligned}
$$

, Where $F()$ is the CDF of a $N(0, \sigma_a^2 + \sigma_u^2 + \sigma_p^2)$

Now solving for the β coefficients:

solving for β_0

$$baseline_0 = F(\beta_0)$$
$$\Leftrightarrow$$
$$\beta_0 = F^{-1}(baseline_0)$$

solving for β_1

$$baseline_1 = F(\beta_0 + \beta_1)$$
$$\Leftrightarrow$$
$$\beta_1 = F^{-1}(baseline_1) - \beta_0$$

solving for $\beta_{2,0}$

$$trend = F(\beta_0 + \beta_{2,0}) - F(\beta_0)$$
$$\Leftrightarrow$$
$$trend + baseline_0 = F(\beta_0 + \beta_{2,0})$$
$$\Leftrightarrow$$
$$F^{-1}(trend + baseline_0) = \beta_0 + \beta_{2,0}$$
$$\Leftrightarrow$$
$$\beta_{2,0} = F^{-1}(trend + baseline_0) - \beta_0$$

solving for $\beta_{2,1}$

$$trend = F(\beta_0 + \beta_1 + \beta_{2,1}) - F(\beta_0 + \beta_1)$$

$$\Leftrightarrow$$

$$trend + baseline_1 = F(\beta_0 + \beta_1 + \beta_{2,1})$$

$$\Leftrightarrow$$

$$F^{-1}(trend + baseline_1) = \beta_0 + +\beta_1 + \beta_{2,1}$$

$$\Leftrightarrow$$

$$\beta_{2,1} = F^{-1}(trend + baseline_1) - \beta_0 - \beta_1$$

solving for β_3

$$ATT = F(\beta_0 + \beta_1 + \beta_{2,1} + \beta_3) - F(\beta_0 + \beta_1 + \beta_{2,1})$$

$$\Leftrightarrow$$

$$ATT + F(\beta_0 + \beta_1 + \beta_{2,1}) = F(\beta_0 + \beta_1 + \beta_{2,1} + \beta_3)$$

$$\Leftrightarrow$$

$$F^{-1}(ATT + F(\beta_0 + \beta_1 + \beta_{2,1})) = \beta_0 + \beta_1 + \beta_{2,1} + \beta_3$$

$$\Leftrightarrow$$

$$\beta_3 = F^{-1}(ATT + F(\beta_0 + \beta_1 + \beta_{2,1})) - (\beta_0 + \beta_1 + \beta_{2,1})$$

when treatment effects are correlated with property size

$$
\begin{aligned}
ATT &= E(\beta_0 + \beta_1 + \beta_{2,1} + \beta_3) - E(\beta_0 + \beta_1 + \beta_{2,1}) \\
&= P(-\alpha_i - u_{it} - \beta_3 < \beta_0 + \beta_1 + \beta_{2,1} + \mu) - P(-\alpha_i - u_{it} < \beta_0 + \beta_1 + \beta_{2,1}) \\
&= G(\beta_0 + \beta_1 + \beta_{2,1} + \mu) - F(\beta_0 + \beta_1 + \beta_{2,1})
\end{aligned}
$$

, where $\beta_3 \sim N(\mu, \sigma_{te}^2)$ and $G()$ is the CDF of a $N(0, \sigma_a^2 + \sigma_u^2 + \sigma_p^2 + \sigma_{te}^2)$ and

$$ATT = G(\beta_0 + \beta_1 + \beta_{2,1}) - F(\beta_0 + \beta_1 + \beta_{2,1})$$

$$\Leftrightarrow$$

$$ATT + F(\beta_0 + \beta_1 + \beta_{2,1}) = G(\beta_0 + \beta_1 + \beta_{2,1} + \mu)$$

$$\Leftrightarrow$$

$$G^{-1}(ATT + F(\beta_0 + \beta_1 + \beta_{2,1})) = \beta_0 + \beta_1 + \beta_{2,1} + \mu$$

$$\Leftrightarrow$$

$$\mu = G^{-1}(ATT + F(\beta_0 + \beta_1 + \beta_{2,1})) - (\beta_0 + \beta_1 + \beta_{2,1})$$

A.1.8 Full summary figure from all specifications and values of σ_p

Using Figure A.1 to compare across all specifications and varying σ_p, we see that RMSE tends to increase across all specifications as σ_p increases. The pixel-level TWFE specifications with spatially aggregated unit fixed effects tend to have the lowest RMSE whenever σ_p is nonzero. In contrast, the specification with the property as the unit of analysis and pixel-level DID tend to have the highest RMSE.

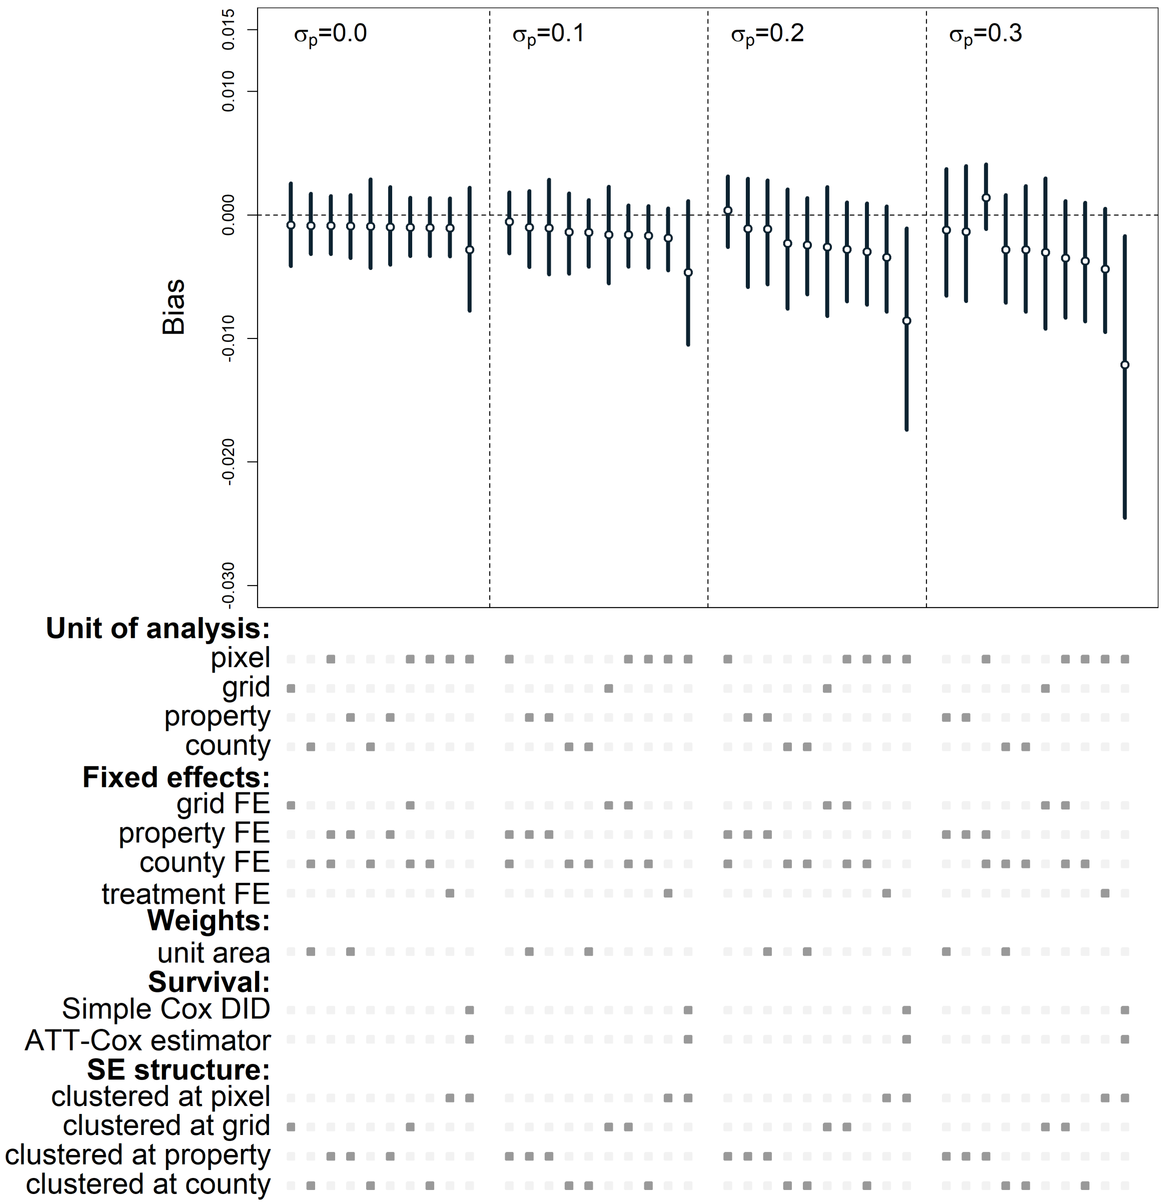

Figure A.1: Bias, RMSE, and coverage for all candidate models with varying values of σ_p.

A.1.9 Alternate landscape parameterizations

Figures A.2 and A.3 show model performance for two alternate landscape parameterizations in the presence of property unobservables ($\sigma_p = 0.3$), ordered from least to most bias. Figure A.2 considers our initial parameterization, but switches the pre-treatment deforestation rates for the two groups. This leaves a pre-treatment deforestation rate of 5% in the control area and 2% in the intervention area. Figure A.3 considers the initial parameterization, but instead makes the ATT positive (i.e. $ATT = 0.01$ rather than -0.01).

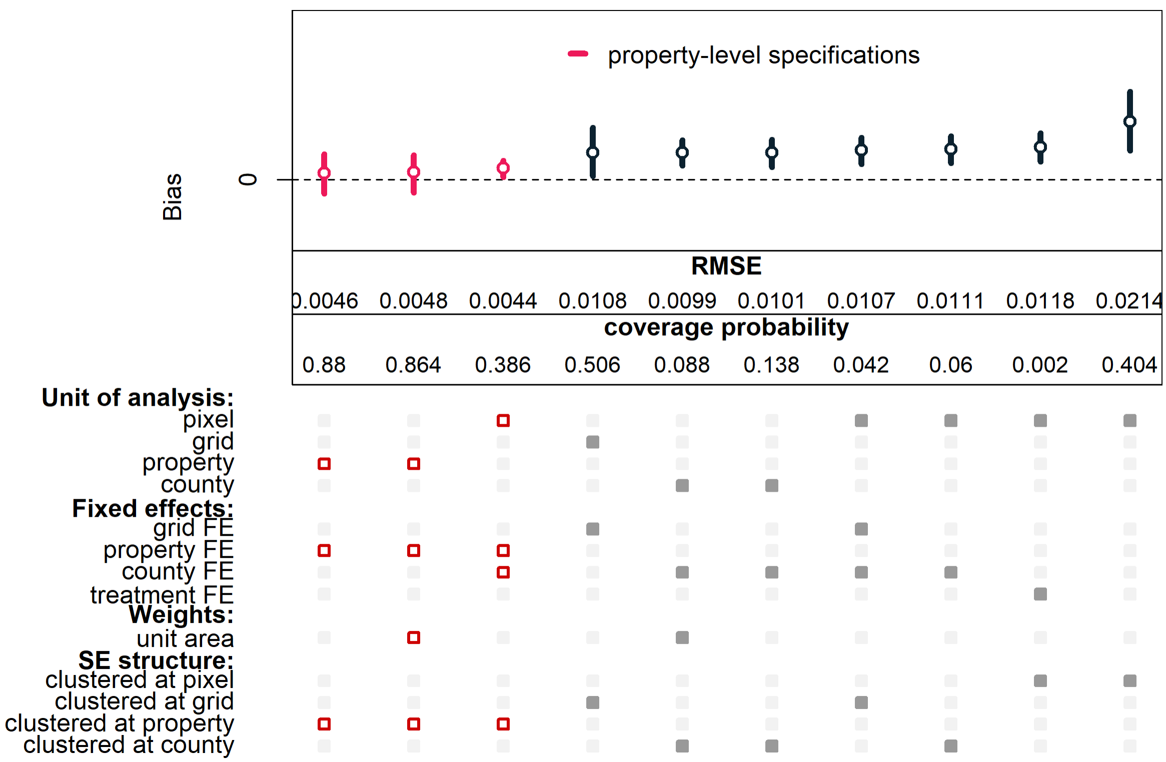

Figure A.2: Model performance under altrnate parameterization 1. Property-level specification still outperform others in presence of property unobservables with alternative landscape parameterizations.

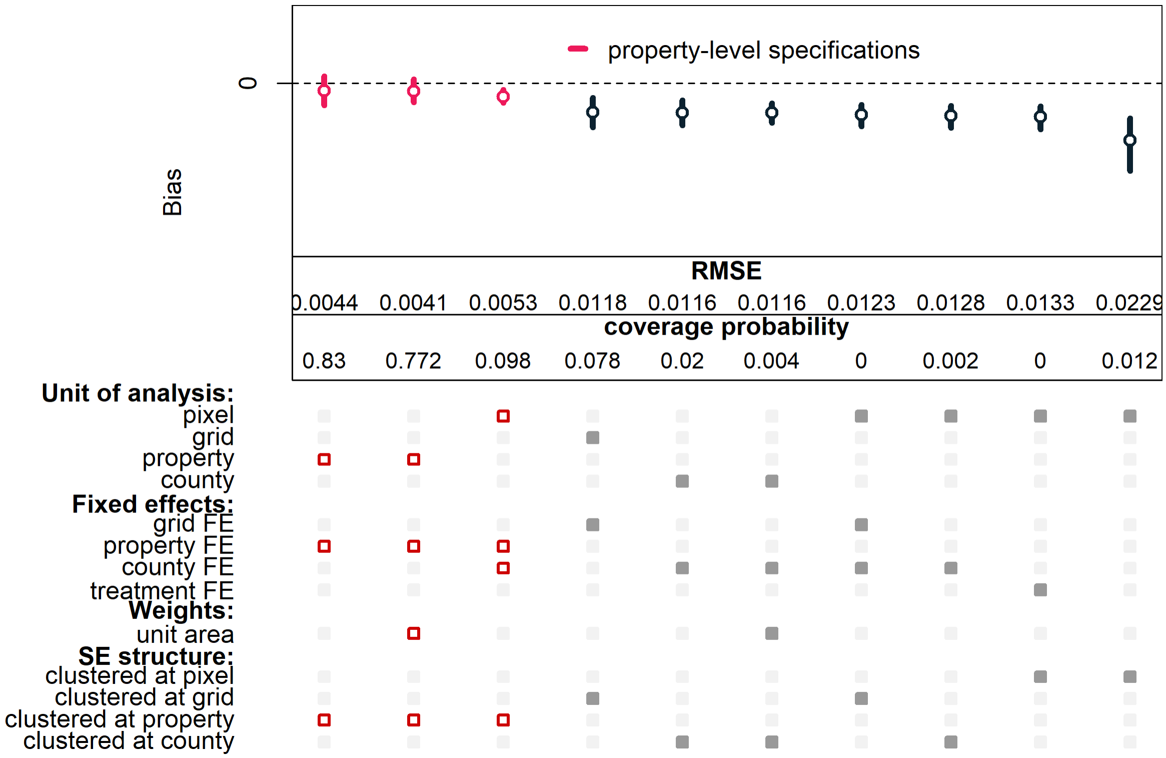

Figure A.3: Model performance under alternate parameterization 2. Property-level specification still outperform others in presence of property unobservables with alternative landscape parameterizations.

A.1.10 DGP for multiple groups and variation in treatment timing

The following parameters and their definitions inform the simulation parameterizations.

$$\begin{aligned}
baseline_a &= E[y_{it}(0)|t < t_0, G_i = a] \\
baseline_b &= E[y_{it}(0)|t < t_0, G_i = b] \\
baseline_c &= E[y_{it}(0)|t < t_0, G_i = c] \\
trend_1 &= E[y_{it}(0)|t = 1, G_i = g] - E[y_{it}(0)|t = 0, G_i = g] \\
trend_2 &= E[y_{it}(0)|t = 2, G_i = g] - E[y_{it}(0)|t = 1, G_i = g] \\
trend_3 &= E[y_{it}(0)|t = 3, G_i = g] - E[y_{it}(0)|t = 2, G_i = g] \\
trend_4 &= E[y_{it}(0)|t = 4, G_i = g] - E[y_{it}(0)|t = 3, G_i = g] \\
ATT &= E[y_{it}(1) - y_{it}(0)|t \geq t_0, G_i = g]
\end{aligned}$$

Here, three groups, $g \in \{a, b, c\}$ have different baseline deforestation rates, and all three groups would experience the same trends in the absence of treatment. Group a experiences treatment in time 2, group b experiences treatment in time 3, and group c is never treated. The ATT is equal across the two treated groups and there are no dynamic effects.

The DGP for each observation can be written as follows:

Group a:

$$\begin{aligned}
y_{it}^* =& \beta_{0,a} 1\{t = 0\} + \beta_{1,a} 1\{t = 1\} + \beta_{2,a} 1\{t = 2\} + \beta_{3,a} 1\{t = 3\} + \beta_{4,a} 1\{t = 4\} \\
&+ \tau_a 1\{t \geq 2\} + \alpha_i + u_{it}
\end{aligned}$$

Group b:

$$
\begin{aligned}
y_{it}^{*} =& \beta_{0,b}1\{t=0\} + \beta_{1,b}1\{t=1\} + \beta_{2,b}1\{t=2\} + \beta_{3,b}1\{t=3\} + \beta_{4,b}1\{t=4\} \\
& + \tau_b 1\{t \geq 3\} + \alpha_i + u_{it}
\end{aligned}
$$

Group c:

$$
\begin{aligned}
y_{it}^{*} =& \beta_{0,c}1\{t=0\} + \beta_{1,c}1\{t=1\} + \beta_{2,c}1\{t=2\} + \beta_{3,c}1\{t=3\} + \beta_{4,c}1\{t=4\} \\
& + \alpha_i + u_{it}
\end{aligned}
$$

, where the β and τ coefficients are calculated as follows:

$$
\begin{aligned}
\beta_{0,a} &= F^{-1}(baseline_a) \\
\beta_{1,a} &= F^{-1}(trend_1 + baseline_a) - \beta_{0,a} \\
\beta_{2,a} &= F^{-1}(trend_2 + F(\beta_{0,a} + \beta_{1,a})) - \beta_{0,a} - \beta_{1,a} \\
\beta_{3,a} &= F^{-1}(trend_3 + F(\beta_{0,a} + \beta_{1,a} + \beta_{2,a})) - \beta_{0,a} - \beta_{1,a} - \beta_{2,a} \\
\beta_{4,a} &= F^{-1}(trend_4 + F(\beta_{0,a} + \beta_{1,a} + \beta_{2,a} + \beta_{3,a})) - \beta_{0,a} - \beta_{1,a} - \beta_{2,a} - \beta_{3,a} \\
\tau_a &= F^{-1}(ATT + F(\beta_{0,a} + \beta_{1,a} + \beta_{2,a})) - \beta_{0,a} - \beta_{1,a} - \beta_{2,a}
\end{aligned}
$$

$$
\begin{aligned}
\beta_{0,b} &= F^{-1}(baseline_b)\\
\beta_{1,b} &= F^{-1}(trend_1 + baseline_b) - \beta_{0,b}\\
\beta_{2,b} &= F^{-1}(trend_2 + F(\beta_{0,b} + \beta_{1,b})) - \beta_{0,b} - \beta_{1,b}\\
\beta_{3,b} &= F^{-1}(trend_3 + F(\beta_{0,b} + \beta_{1,b} + \beta_{2,b})) - \beta_{0,b} - \beta_{1,b} - \beta_{2,b}\\
\beta_{4,b} &= F^{-1}(trend_4 + F(\beta_{0,b} + \beta_{1,b} + \beta_{2,b} + \beta_{3,b})) - \beta_{0,b} - \beta_{1,b} - \beta_{2,b} - \beta_{3,b}\\
\tau_b &= F^{-1}(ATT + F(\beta_{0,b} + \beta_{1,b} + \beta_{2,b} + \beta_{3,b})) - \beta_{0,b} - \beta_{1,b} - \beta_{2,b} - \beta_{3,b}
\end{aligned}
$$

$$
\begin{aligned}
\beta_{0,c} &= F^{-1}(baseline_c)\\
\beta_{1,c} &= F^{-1}(trend_1 + baseline_c) - \beta_{0,c}\\
\beta_{2,c} &= F^{-1}(trend_2 + F(\beta_{0,c} + \beta_{1,c})) - \beta_{0,c} - \beta_{1,c}\\
\beta_{3,c} &= F^{-1}(trend_3 + F(\beta_{0,c} + \beta_{1,c} + \beta_{2,c})) - \beta_{0,c} - \beta_{1,c} - \beta_{2,c}\\
\beta_{4,c} &= F^{-1}(trend_4 + F(\beta_{0,c} + \beta_{1,c} + \beta_{2,c} + \beta_{3,c})) - \beta_{0,c} - \beta_{1,c} - \beta_{2,c} - \beta_{3,c}
\end{aligned}
$$

, Where $F()$ is the CDF of a $N(0, \sigma_a^2 + \sigma_u^2)$

parameterization for heterogeneous treatment effects example

The following parameters and their definitions inform the simulation parameterizations.

$$
\begin{aligned}
baseline_a &= E[y_{it}(0)|t < t_0, G_i = a] \\
baseline_b &= E[y_{it}(0)|t < t_0, G_i = b] \\
baseline_c &= E[y_{it}(0)|t < t_0, G_i = c] \\
trend_1 &= E[y_{it}(0)|t = 1, G_i = g] - E[y_{it}(0)|t = 0, G_i = g] \\
trend_2 &= E[y_{it}(0)|t = 2, G_i = g] - E[y_{it}(0)|t = 1, G_i = g] \\
trend_3 &= E[y_{it}(0)|t = 3, G_i = g] - E[y_{it}(0)|t = 2, G_i = g] \\
trend_4 &= E[y_{it}(0)|t = 4, G_i = g] - E[y_{it}(0)|t = 3, G_i = g] \\
ATT_{0,a} &= E[y_{it}(1) - y_{it}(0)|t = 2, G_i = a] \\
ATT_{1,a} &= E[y_{it}(1) - y_{it}(0)|t = 3, G_i = a] \\
ATT_{2,a} &= E[y_{it}(1) - y_{it}(0)|t = 4, G_i = a] \\
ATT_{0,b} &= E[y_{it}(1) - y_{it}(0)|t = 3, G_i = b] \\
ATT_{1,b} &= E[y_{it}(1) - y_{it}(0)|t = 4, G_i = b]
\end{aligned}
$$

Here, three groups, $g \in \{a, b, c\}$ have different baseline deforestation rates, and all three groups would experience the same trends in the absence of treatment. Group a experiences treatment in time 2, group b experiences treatment in time 3, and group c is never treated. The ATT is equal across the two treated groups and there are no dynamic effects.

The DGP for each observation can be written as follows:

Group a:

$$\begin{aligned} y_{it} =& \beta_{0,a} 1\{t=0\} + \beta_{1,a} 1\{t=1\} + (\beta_{2,a} + \tau_{0,a}) 1\{t=2\} \\ &+ (\beta_{3,a} + \tau_{1,a}) 1\{t=3\} + (\beta_{4,a} + \tau_{2,a}) 1\{t=4\} + \alpha_i + u_{it} \end{aligned}$$

Group b:

$$\begin{aligned} y_{it} =& \beta_{0,b} 1\{t=0\} + \beta_{1,b} 1\{t=1\} + \beta_{2,b} 1\{t=2\} \\ &+ (\beta_{3,b} + \tau_{0,b}) 1\{t=3\} + (\beta_{4,b} + \tau_{1,b}) 1\{t=4\} + \alpha_i + u_{it} \end{aligned}$$

Group c:

$$\begin{aligned} y_{it} =& \beta_{0,c} 1\{t=0\} + \beta_{1,c} 1\{t=1\} + \beta_{2,c} 1\{t=2\} \\ &+ \beta_{3,c} 1\{t=3\} + \beta_{4,c} 1\{t=4\} + \alpha_i + u_{it} \end{aligned}$$

, where the β and tau coefficients are calculated as follows:

$$\beta_{0,a} = F^{-1}(baseline_a)$$

$$\beta_{1,a} = F^{-1}(trend_1 + baseline_a) - \beta_{0,a}$$

$$\beta_{2,a} = F^{-1}(trend_2 + F(\beta_{0,a} + \beta_{1,a})) - \beta_{0,a} - \beta_{1,a}$$

$$\beta_{3,a} = F^{-1}(trend_3 + F(\beta_{0,a} + \beta_{1,a} + \beta_{2,a})) - \beta_{0,a} - \beta_{1,a} - \beta_{2,a}$$

$$\beta_{4,a} = F^{-1}(trend_4 + F(\beta_{0,a} + \beta_{1,a} + \beta_{2,a} + \beta_{3,a})) - \beta_{0,a} - \beta_{1,a} - \beta_{2,a} - \beta_{3,a}$$

$$\tau_{0,a} = F^{-1}(ATT_{0,a} + F(\beta_{0,a} + \beta_{1,a} + \beta_{2,a})) - \beta_{0,a} - \beta_{1,a} - \beta_{2,a}$$

$$\tau_{1,a} = F^{-1}(ATT_{1,a} + F(\beta_{0,a} + \beta_{1,a} + \beta_{2,a} + \beta_{3,a} + \tau_{0,a})) - \beta_{0,a} - \beta_{1,a} - \beta_{2,a} - \beta_{3,a} - \tau_{0,a}$$

$$\tau_{2,a} = F^{-1}(ATT_{2,a} + F(\beta_{0,a} + \beta_{1,a} + \beta_{2,a} + \beta_{3,a} + \tau_{0,a} + \tau_{1,a})) - \beta_{0,a} - \beta_{1,a} - \beta_{2,a} - \beta_{3,a} - \tau_{0,a} - \tau_{1,a}$$

$$\beta_{0,b} = F^{-1}(baseline_b)$$

$$\beta_{1,b} = F^{-1}(trend_1 + baseline_b) - \beta_{0,b}$$

$$\beta_{2,b} = F^{-1}(trend_2 + F(\beta_{0,b} + \beta_{1,b})) - \beta_{0,b} - \beta_{1,b}$$

$$\beta_{3,b} = F^{-1}(trend_3 + F(\beta_{0,b} + \beta_{1,b} + \beta_{2,b})) - \beta_{0,b} - \beta_{1,b} - \beta_{2,b}$$

$$\beta_{4,b} = F^{-1}(trend_4 + F(\beta_{0,b} + \beta_{1,b} + \beta_{2,b} + \beta_{3,b})) - \beta_{0,b} - \beta_{1,b} - \beta_{2,b} - \beta_{3,b}$$

$$\tau_{b,0} = F^{-1}(ATT_{0,b} + F(\beta_{0,b} + \beta_{1,b} + \beta_{2,b} + \beta_{3,b})) - \beta_{0,b} - \beta_{1,b} - \beta_{2,b} - \beta_{3,b}$$

$$\tau_{b,1} = F^{-1}(ATT_{1,b} + F(\beta_{0,b} + \beta_{1,b} + \beta_{2,b} + \beta_{3,b} + \tau_{b,0})) - \beta_{0,b} - \beta_{1,b} - \beta_{2,b} - \beta_{3,b} - \tau_{b,0}$$

$$\beta_{0,c} = F^{-1}(baseline_c)$$

$$\beta_{1,c} = F^{-1}(trend_1 + baseline_c) - \beta_{0,c}$$

$$\beta_{2,c} = F^{-1}(trend_2 + F(\beta_{0,c} + \beta_{1,c})) - \beta_{0,c} - \beta_{1,c}$$

$$\beta_{3,c} = F^{-1}(trend_3 + F(\beta_{0,c} + \beta_{1,c} + \beta_{2,c})) - \beta_{0,c} - \beta_{1,c} - \beta_{2,c}$$

$$\beta_{4,c} = F^{-1}(trend_4 + F(\beta_{0,c} + \beta_{1,c} + \beta_{2,c} + \beta_{3,c})) - \beta_{0,c} - \beta_{1,c} - \beta_{2,c} - \beta_{3,c}$$

, Where $F()$ is the CDF of a $N(0, \sigma_a^2 + \sigma_u^2)$

A.2 Chapter 2 appendix

A.2.1 Descriptive statistics and figures

Property size distributions

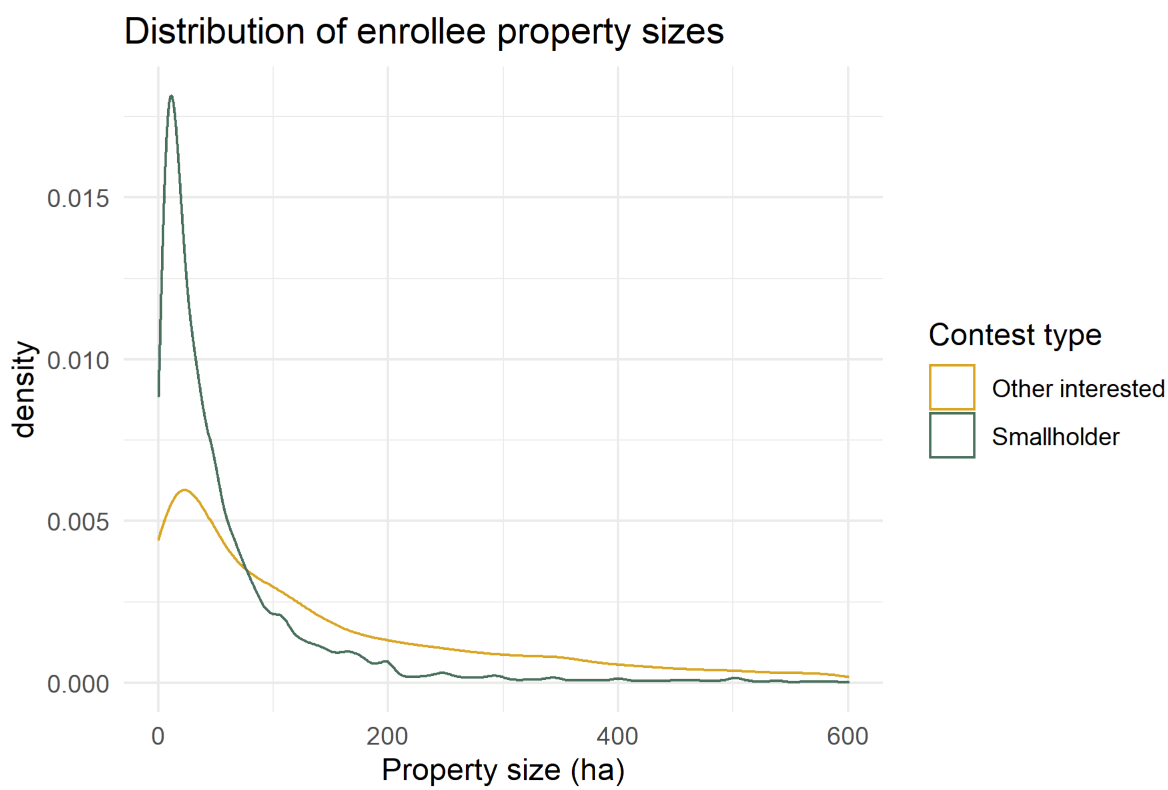

Figure A.4: Distribution of property sizes amongst enrollees in both contests.

Proportion of property subsidized

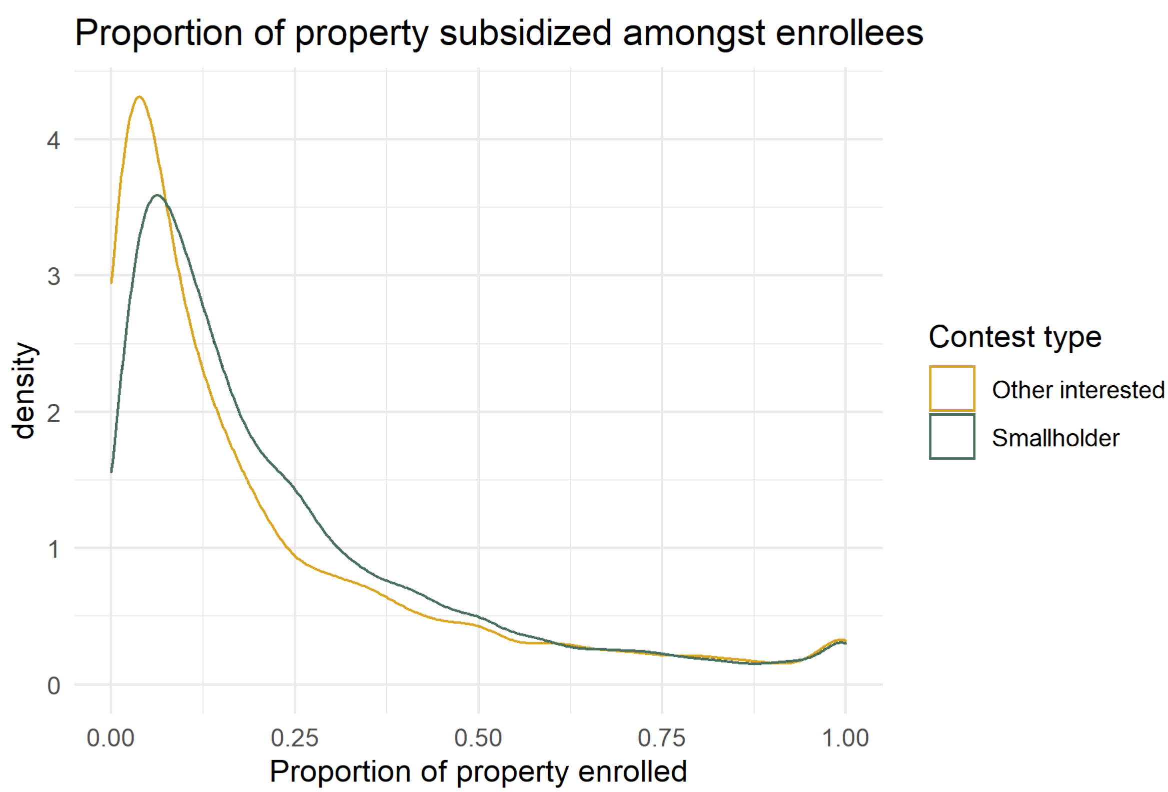

Figure A.5: Distribution of proportion of property enrolled amongst enrollees in both contests.

Beneficiaries vs. non-compliers

Smallholder vs. other interested party contest

A.2.2 Smallholder eligibility rules

Smallholder:

Such is understood to be those who meet all of the following requirements:

- The person who has title to one or more rural properties whose surface area does not exceed 200 hectares, or 500 hectares when they are located between regions I and IV , including the XV; or 800 hectares for properties located in the comuna of Lonquimay, in Region IX; in the province of Palena, in the X Region; or in Region XI and XII, and

- That their assets do not exceed the equivalent of 3,500 unidades de fomento; and

- That his income comes mainly from agricultural or forestry exploitation and that he directly works the land, on his property or on another third-party property.

They are also smallholders:

- Agricultural communities regulated by decree with force of law No. 5, of the Ministry of Agriculture, of 1968,

- Indigenous communities governed by Law No. 19,253,

- Communities over common property resulting from the Agrarian Reform process,

- Societies dry land constituted in accordance with Article 1 of Decree Law No. 2,247 of 1978, and

- The companies referred to in Article 6 of Law No. 19,118, provided that at least 60% of the capital stock of such companies are in the hands of the original partners or persons who have the status of small forest owners, as certified by the Agricultural and Livestock Service.

A.2.3 Pre-processing

A.2.4 Common trends assumption

Raw trends

As mentioned in the main text, for β_1 from Equation 1 to yield the causal effect of the program, we need to rely on a common trends assumption. In order to evaluate the plausibility of common trends in this setting, we examine the raw tree cover trends of the complying enrollees relative to the matched control group. Figure A.7 shows that the matched group had slightly higher levels of tree cover, but the trends are comparable prior

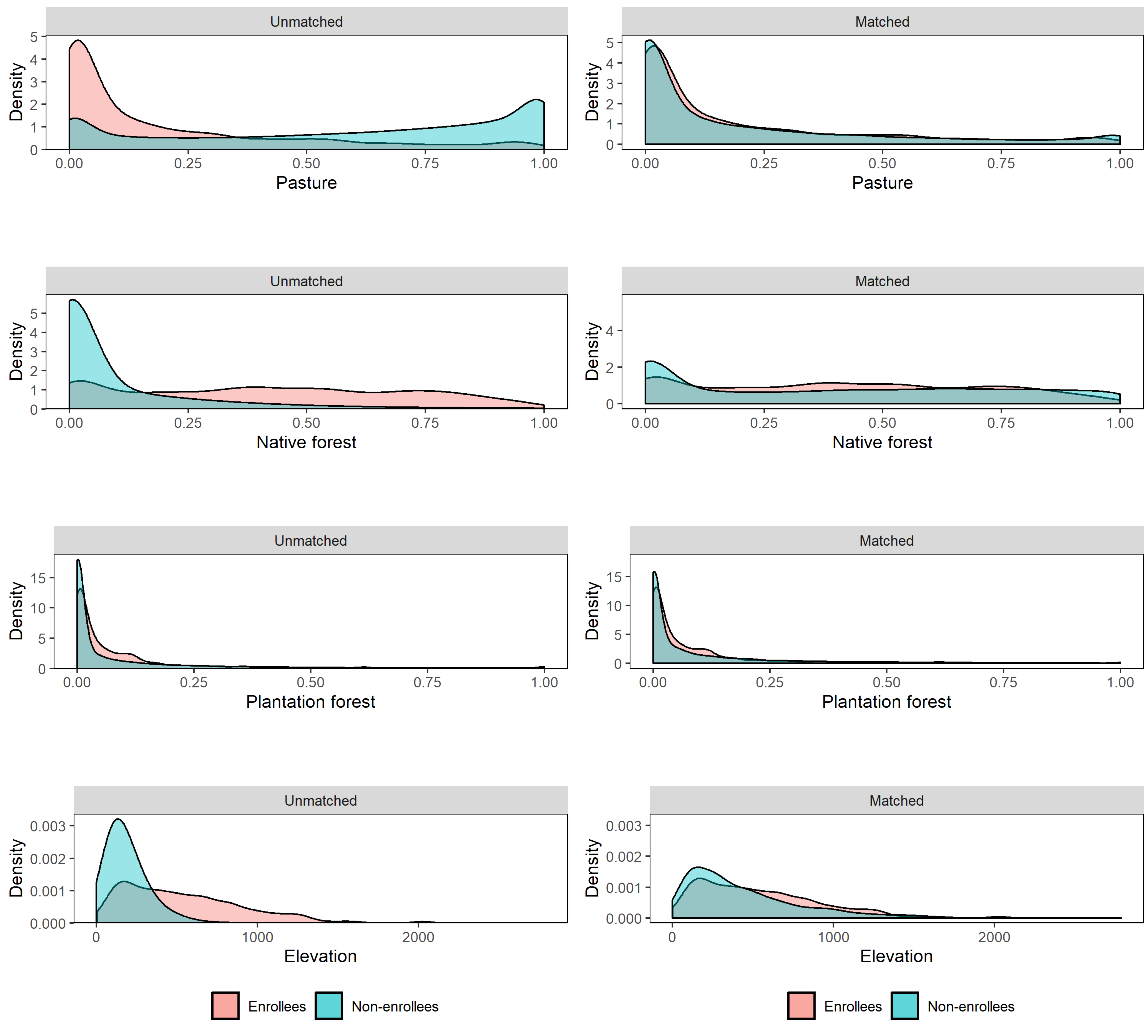

Figure A.6: This figure shows how matching affected distribution balance for selected covariates. The left panel shows distributions of covariates for treated and unenrolled properties. The right panel shows distributions of covariates for treated and matched control properties.

to the existence of the Native Forest Law. This is made particularly clear in the right panel, which adjusts for 2008 differences in tree cover to better evaluate the pre-trends prior to the implementation of the program in 2009.

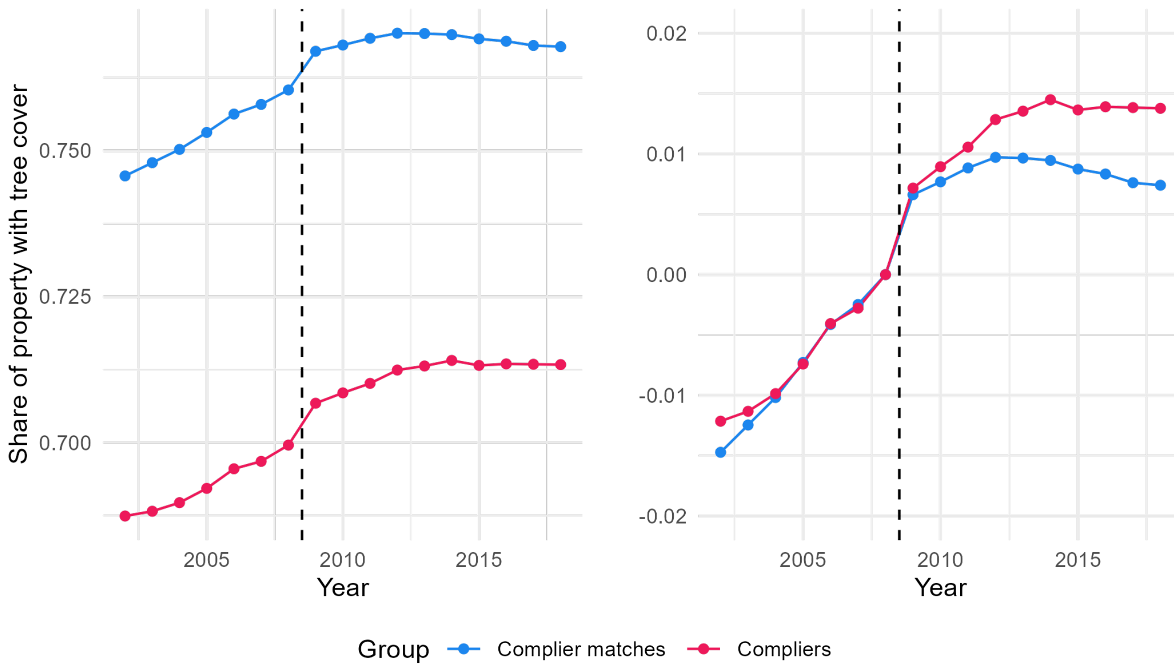

Figure A.7: Raw trends in tree cover shares between enrolled properties and matched control properties. The right panel de-means the trends relative to 2008, the year prior to the first Native Forest Law subsidy contest.

Unconditional pre-trend event study

We also use an event study to examine the plausibility of common trends. The event study accounts for staggered treatment in a way that examining raw trends cannot. No covariates are included, meaning that unconditional common trends are evaluated. Figure A.8 displays pre-treatment pseudo-ATT estimates based on the event study.

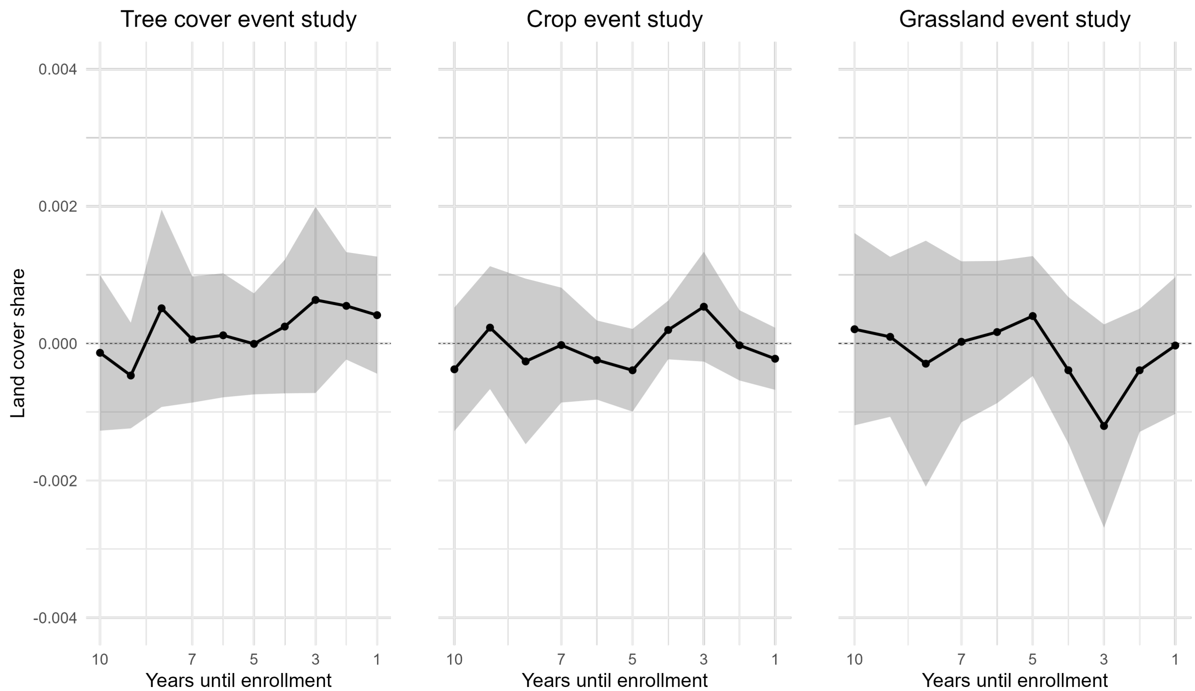

Figure A.8: This figure shows pre-treatment event time treatment effects of the Native Forest Law subsidy contest for beneficiary properties based on unconditional common trends. The three panels show three land cover types: Tree cover, Crop, and Grassland.

A.2.5 Event-study estimates

Callaway and Sant'anna estimator

Recent papers have shown that the typical two-way fixed effects estimator may generate biased results in the presence of treatment effect heterogeneity (e.g. Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020). This could be particularly important in our case, given that there are over 150 cohort-time cells. The estimator proposed in Callaway and Sant'Anna (2020) computes each 2x2 cohort-time treatment effect ($ATT_{g,t}$) individually, before aggregating them with intuitive weights.

The estimand for each of the $ATT_{g,t}$s is as follows:

$$ATT_{g,t} = E[outcome_{it}(1) - outcome_{it}(0)|G_i = g, t \geq t_o]$$

Each $ATT_{g,t}$ then represents the treatment effect for cohort g in time t. To generate the $ATT_{g,t}$s, we first subset the data to only contain observations at time t and $g-1$, from units with either $G_i = g$ or that are in the control group. For example, for the $ATT_{2015,2019}$, we subset to only the 2015 cohort and control group for the years 2014 and 2019. Then using only the observations from this subset, we calculate $ATT_{g,t}$ using the doubly robust difference-in-differences estimator developed in Sant'Anna and Zhao (2020). This involves first estimating a propensity score using a logit model and allows for common trends to hold only after conditioning on pre-treatment covariates. With this method, we can identify the $ATT_{g,t}$s if either (but not necessarily both) the propensity score or outcome regression is correctly specified (Sant'Anna and Zhao, 2020).

We focus on event study measures of the ATT. Within each event time window, we aggregate the $ATT_{g,t}$s with weights corresponding to group size.

$$ATT_{es}(e) = \sum_{g \in G} \mathbb{1}\{g + e \leq \mathcal{T}\} P(G_i = g | G_i + e \leq \mathcal{T}) ATT_{g,g+e}$$

This is the average effect of participating in the treatment e time periods after a characteristic property is enrolled in the program across all cohorts that are ever observed to have participated in the treatment for exactly e time periods. The year a property enrolls in the program is denoted by $e = 0$.

Balanced event study

Callaway and Sant'Anna (2020) discuss the fact that interpretation of $ATT_{es}(e)$ may be complicated by compositional changes through time. In our case, this may impact the interpretation of dynamic treatment effects if changes in the cohort composition through event time create the appearance of increasing effects through time. To determine whether this may affect interpretation of treatment effects through time, we also estimate

$ATT_{bal}(e, e')$:

$$ATT_{bal}(e, e') = \sum_{g \in G} \mathbb{1}\{g + e' \leq \mathcal{T}\} P(G_i = g | G_i + e' \leq \mathcal{T}) ATT_{g,g+e}$$

The definition of $ATT_{bal}(e, e')$ is very similar to $ATT_{es}(e)$ except that it calculates the average group-time average treatment effect for units whose event time is equal to e and who are observed to participate in the treatment for at least e' periods. In our context, we set $e' = 8$, meaning that only properties observed at least 8 years after the initial enrollment year are included. Differences in $ATT_{bal}(e, e')$ across different values of e, therefore, cannot be due to differences in the composition of groups at different values of e (Callaway and Sant'Anna, 2020).

Event study estimates based on only never-treated units

We explore whether exclusion of not yet treated cohorts changes our estimates. Figure A.9 shows that using only the matched control group as the control yields comparable event study estimates.

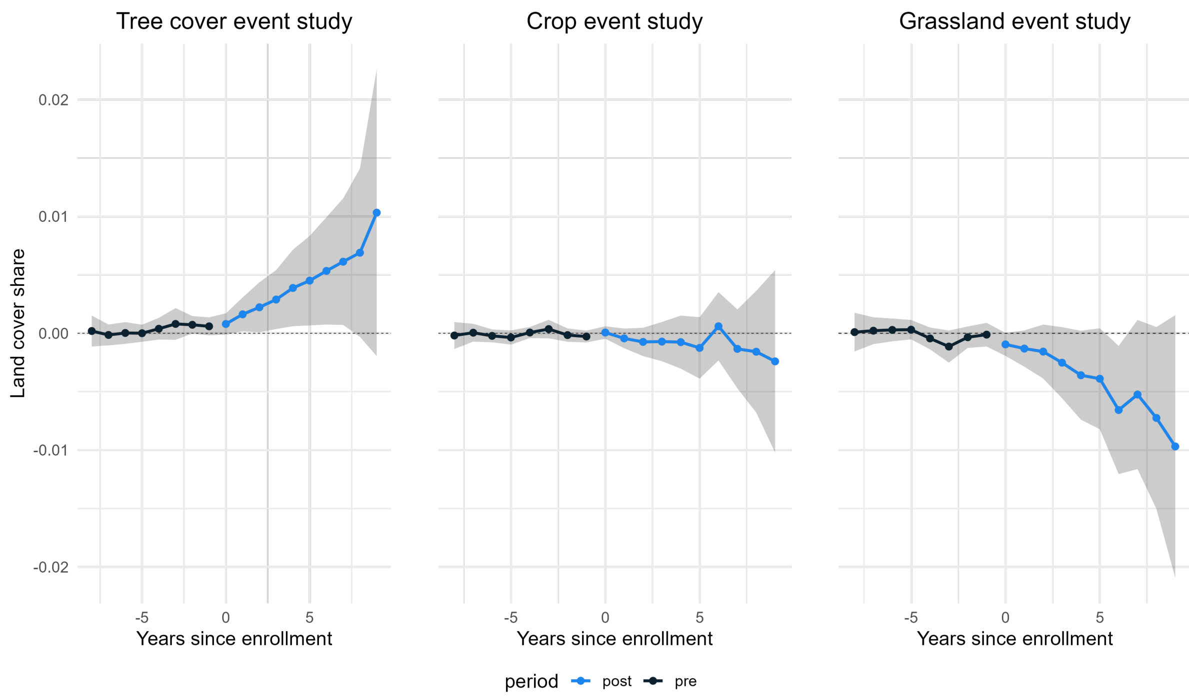

Figure A.9: This figure shows event time treatment effects of the Native Forest Law subsidy contest for beneficiary properties including only the matched control group in the control group. The three panels show event time treatment effects based on binary treatment for three land cover types: Tree cover, Crop, and Grassland.

A.2.6 Cohort treatment effects

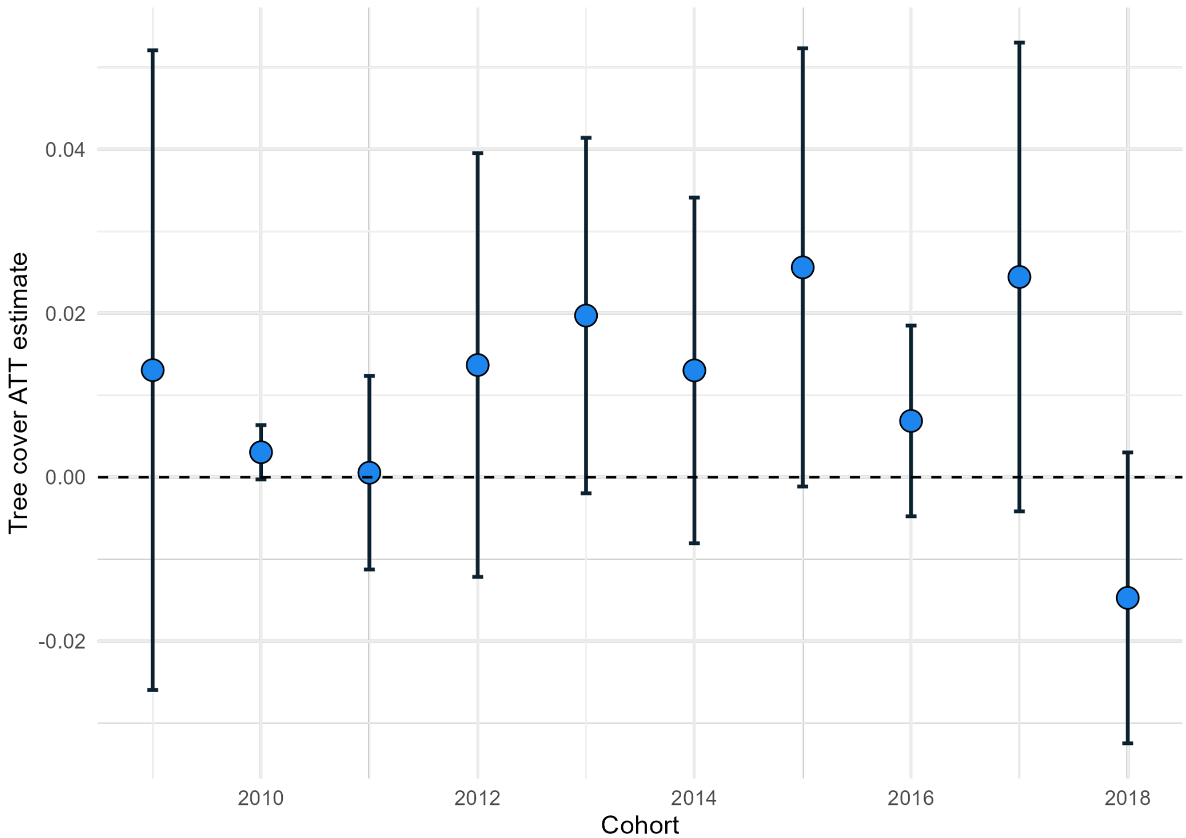

Figure A.10: This figure shows tree cover treatment effects of the Native Forest Law subsidy contest for beneficiary properties based on Equation 1, separated by cohort.

A.2.7 Rejected applicant histogram

During the annual Native Forest Law contest, applicants submit a management plan detailing the specifics of the project to be considered for an award. Judges score each application based on a number of criteria including the size of the property, project extent, specific activities to be performed, and the cost of the project. After scores are tallied, awards are dispersed in order of project score. Awards are given to the smallholder contest first, and subsequently to the other interested party contest. Thus, in years when the contest's funds run out, other interested parties generally go unfunded. In years in which the contest does not exceed the funding threshold for both groups, a second smallholder-only contest is held for any additional project applicants. These contests generate unawarded smallholders. Projects can become rejected either by scoring below the threshold that receives funding or because of unapproved proposed activities in the application itself. Thus there are two ways to get rejected.

Figure A.11 shows the distribution of rejected applicants by contest across different contest years. One thing to note is that many of the rejected applicants are able to adjust their application, reapply, and enroll in subsequent years. We see that the two contests see different trends in the number of rejected applicants through time.

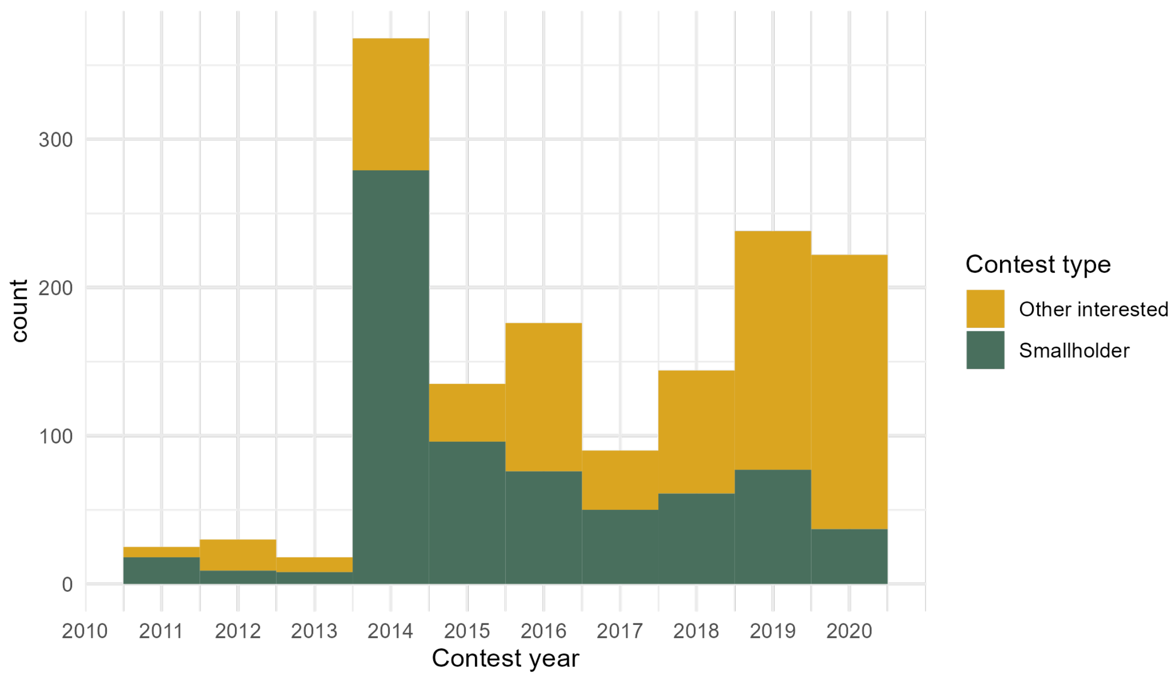

Figure A.11: This figure shows rejected applicants by contest across years.

A.3 Chapter 3 appendix

A.3.1 Time varying treatment effects of ash borer infestation on tree cover across Chicago metropolitan region

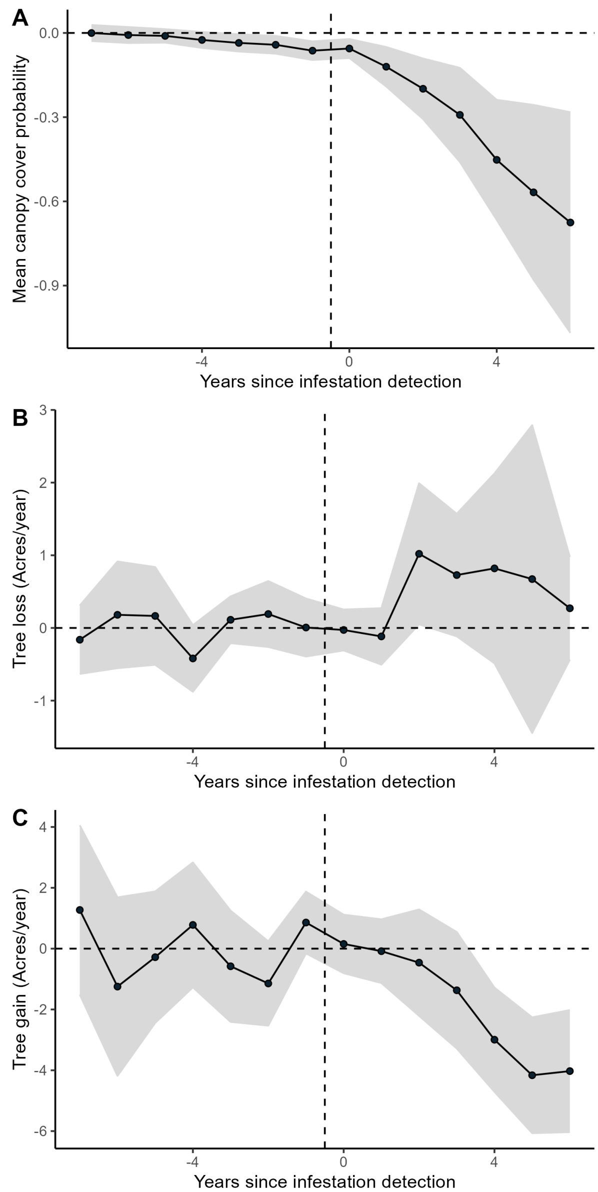

Figure A.12: This figure shows how estimates of infestation impacts vary across event time for tree cover outcomes with treatment defined at the 5 x 5km grid cell. The tree cover outcomes represent: A) canopy cover; B) rate of tree cover loss in acres per year; and C) rate of tree cover gain in acres per year

A.3.2 Canopy cover effect robust to alternate grid cell size and treatment assignment

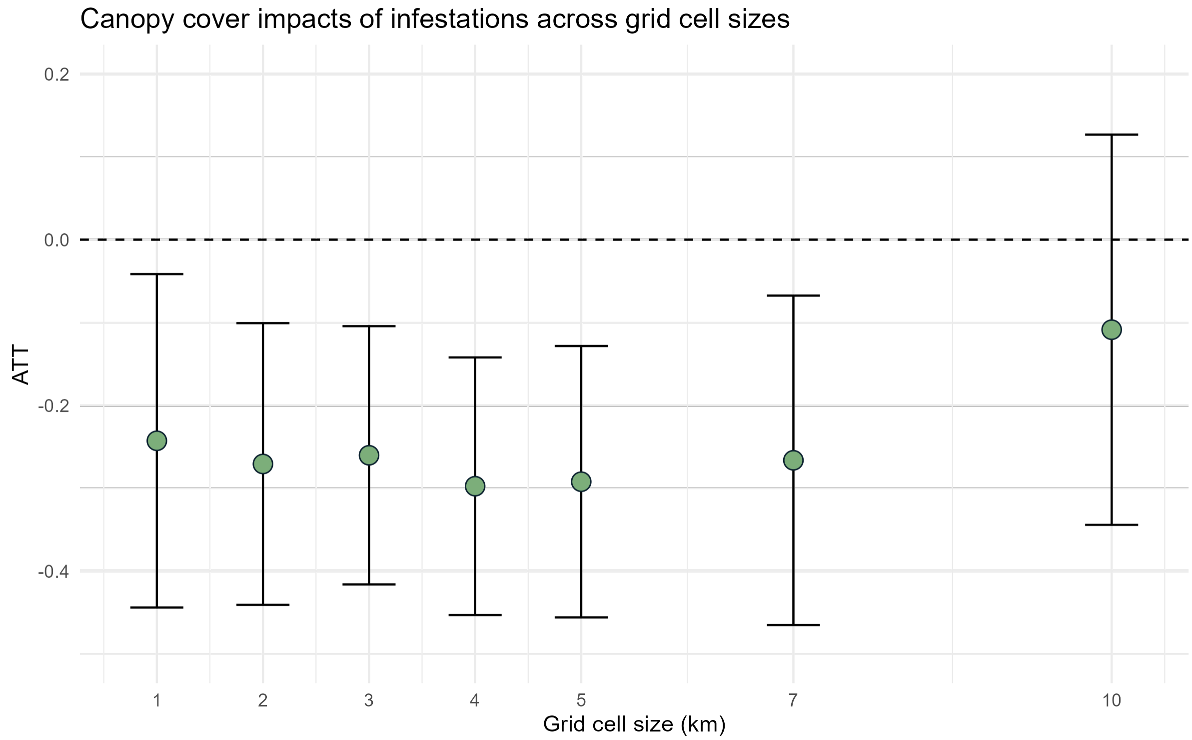

Figure A.13: This specification chart shows difference-in-differences estimates of the impact of infestation on canopy cover using a variety of grid cell size definitions.

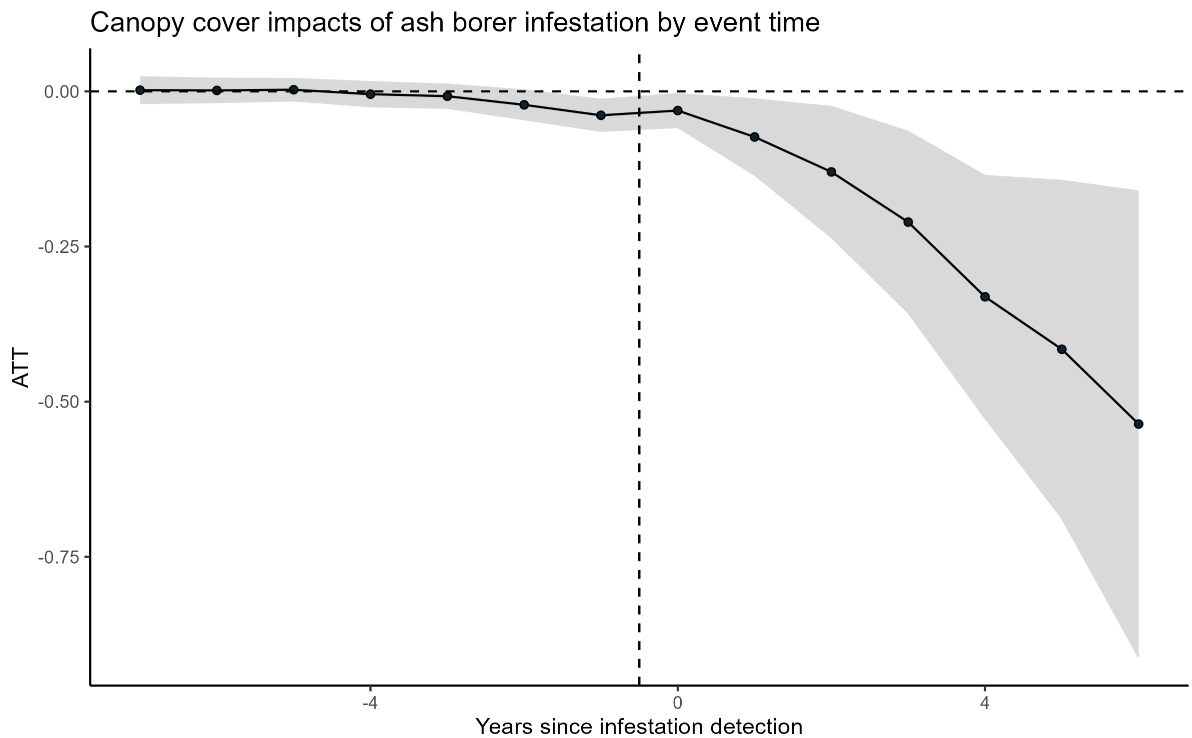

Figure A.14: This figure shows estimates of canopy cover infestation impacts across event time when treatment is assigned at the city level. This is done in order to ensure the estimates are robust in the case that surveyors avoid areas without a confirmed infestation because one already exists in the city.

A.3.3 Event study estimator description

A.3.4 Event-study estimates

The estimand for each of the $ATT_{g,t}$s is as follows:

$$ATT_{g,t} = E[outcome_{it}(1) - outcome_{it}(0)|G_i = g, t \geq t_o]$$

Each $ATT_{g,t}$ then represents the treatment effect for group g in time t. To generate the $ATT_{g,t}$s, we first subset the data to only contain observations at time t and $g-1$, from units with either $G_i = g$ or that are in the control group. For example, for the $ATT_{2010,2013}$, we subset to only the group first treated in 2010 and control group (includes not-yet-treated groups) for the years 2009 and 2013. Then using only the observations from this subset, we calculate $ATT_{g,t}$ using the doubly robust difference-in-differences estimator developed in Sant'Anna and Zhao (2020). This involves first estimating a propensity score using a logit model and allows for common trends to hold only after conditioning on pre-treatment covariates. With this method, we can identify the $ATT_{g,t}$s if either (but not necessarily both) the propensity score or outcome regression is correctly specified (Sant'Anna and Zhao, 2020).

Callaway and Sant'anna estimator

Here, we focus on event study measures of the ATT. Within each event time window, we aggregate the $ATT_{g,t}$s with weights corresponding to group size.

$$ATT_{es}(e) = \sum_{g \in G} \mathbb{1}\{g + e \leq \mathcal{T}\} P(G_i = g | G_i + e \leq \mathcal{T}) ATT_{g,g+e}$$

This is the average effect of exposure to treatment e time periods after a grid cell or school (depending on section/unit of analysis) is exposed to ash borer infestation and that has been exposed for exactly e time periods. The year in which the infestation is first detected is denoted $e = 0$.

For high values of e, very few schools may be included in the calculation of $ATT_{es}(e)$. This may lead to unrepresentative treatment effect estimates. I therefore, display estimates of $ATT_{es}(e)$ for event time windows which contain at least 10% of the total ever-treated schools. Figure A.15 shows the event time windows included and excluded in the event study plots.

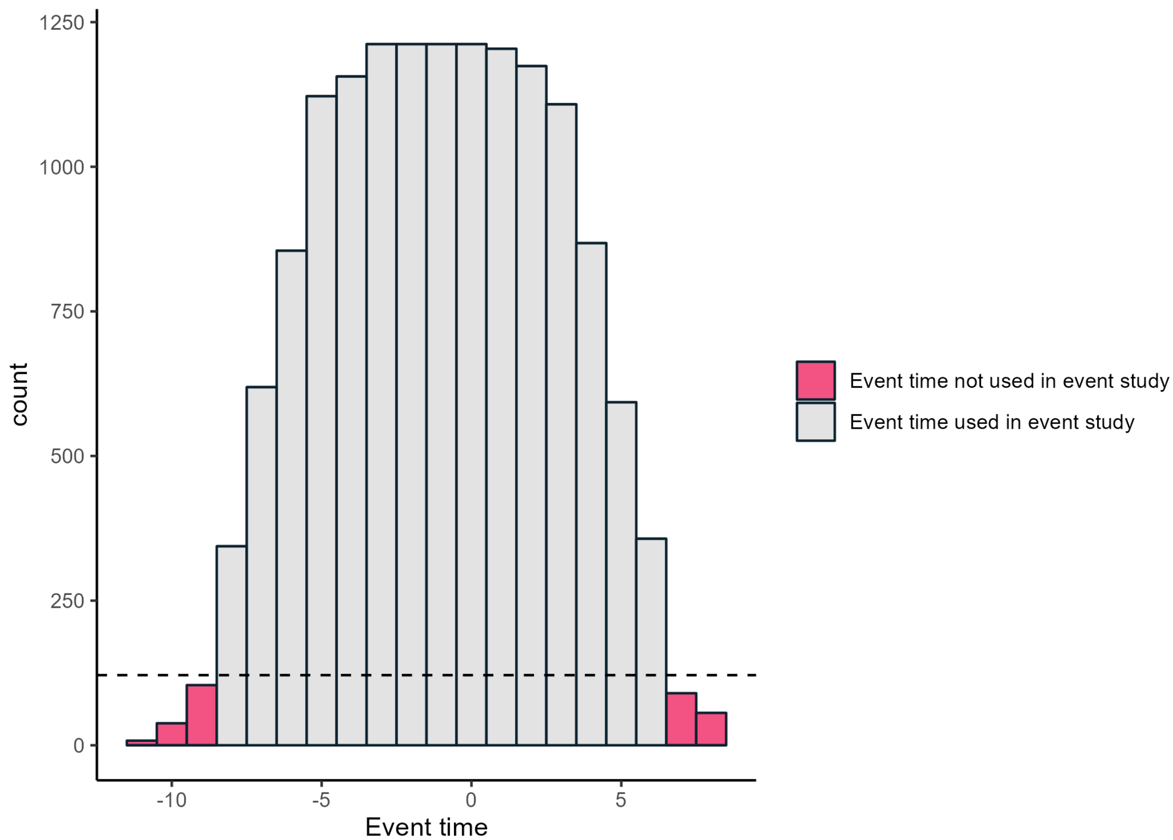

Figure A.15: This figure shows which event times are included in the event study plots throughout the paper for school-level outcomes. Event times that represent less than 10% of the total number of ever-treated schools are excluded from the plots.

A.3.5 Ash tree presence drives canopy cover declines

In order to understand how various factors affect the loss of canopy cover following ash borer infestation, I use two-way fixed effect regressions of the following form:

$$canopy_{it} = \beta_0 \times infestation_{it} + \beta_1 \times infestation_{it} \times Z_i + \gamma_i + \lambda_t + u_{it} \tag{A.6}$$

where Z_i represents the covariate of interest. β_1 is then the coefficient of interest and represents how Z_i moderates or exacerbates the canopy cover loss after infestation detection.

Table A.2 shows that as the number of ash trees per acre increases, the expected canopy cover loss from the ash borer increases as well. This is in line with the fact that only ash trees should be affected by the ash borer directly.

Table A.2: Heterogeneous canopy cover impacts of ash borer infestation

	(1)	(2)	(3)
Infestation	-2.4273***	-16.6830***	-13.8495***
	(0.5898)	(4.6411)	(4.5719)
Infestation x ln(Ash per acre)	-0.2763**		-0.4197***
	(0.0993)		(0.1021)
Infestation x ln(Canopy baseline)	0.8522***		1.0596***
	(0.2176)		(0.2190)
Infestation x ln(Med. income)		1.4612***	0.9855**
		(0.4152)	(0.4173)
Num.Obs.	7056	4576	4560
R2	0.993	0.993	0.993

* p < 0.1, ** p < 0.05, *** p < 0.01
Standard errors are clustered at the grid level.

Bibliography

Amanda Y. Agan and Michael D. Makowsky. The minimum wage, EITC, and criminal recidivism. *NBER WORKING PAPER SERIES*, Working Paper 25116, 2018.

J. M. Alix-Garcia, E. N. Shapiro, and K. R. E. Sims. Forest Conservation and Slippage: Evidence from Mexico's National Payments for Ecosystem Services Program. *Land Economics*, 88(4):613–638, November 2012. ISSN 0023-7639, 1543-8325. doi: 10.3368/le.88.4.613. URL `http://le.uwpress.org/cgi/doi/10.3368/le.88.4.613`.

Jennifer Alix-Garcia and Holly K. Gibbs. Forest conservation effects of Brazil's zero deforestation cattle agreements undermined by leakage. *Global Environmental Change*, 47:201–217, November 2017. ISSN 0959-3780. doi: 10.1016/j.gloenvcha.2017.08.009.

Jennifer Alix-Garcia and Daniel Millimet. Remotely Incorrect? Accounting for Nonclassical Measurement Error in Satellite Data on Deforestation. *Journal of the Association of Environmental and Resource Economists*, page 723723, December 2022. ISSN 2333-5955, 2333-5963. doi: 10.1086/723723. URL `https://www.journals.uchicago.edu/doi/10.1086/723723`.

Jennifer Alix-Garcia and Hendrik Wolff. Payment for Ecosystem Services from Forests. *Annual Review of Resource Economics*, 6(1):361–380, November 2014. ISSN 1941-1340, 1941-1359. doi: 10.1146/annurev-resource-100913-012524. URL `http://www.annualreviews.org/doi/10.1146/annurev-resource-100913-012524`.

Jennifer Alix-Garcia, Lisa L. Rausch, Jessica L'Roe, Holly K. Gibbs, and Jacob Munger. Avoided Deforestation Linked to Environmental Registration of Properties in the Brazilian Amazon: Environmental registration in the Amazon. *Conservation Letters*, 11(3):e12414, May 2018. ISSN 1755263X. doi: 10.1111/conl.12414.

Jennifer M. Alix-Garcia, Katharine R. E. Sims, and Patricia Yañez-Pagans. 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, November 2015. ISSN 1945-7731, 1945-774X. doi: 10.1257/pol.20130139. URL `http://pubs.aeaweb.org/doi/10.1257/pol.20130139`.

Jennifer M. Alix-Garcia, Katharine R.E. Sims, Victor Hugo Orozco-Olvera, Laura Costica, Jorge David Fernandez Medina, Sofia Romo-Monroy, and Stefano Pagiola. *Can*

Environmental Cash Transfers Reduce Deforestation and Improve Social Outcomes? A Regression Discontinuity Analysis of Mexico's National Program (2011–2014). World Bank, Washington, DC, January 2019. doi: 10.1596/1813-9450-8707. URL `http://hdl.handle.net/10986/31175`.

Christa M. Anderson, Gregory P. Asner, William Llactayo, and Eric F. Lambin. Overlapping land allocations reduce deforestation in Peru. *Land Use Policy*, 79:174–178, December 2018. ISSN 02648377. doi: 10.1016/j.landusepol.2018.08.002.

Claudio Araujo, Catherine Araujo Bonjean, Jean-Louis Combes, Pascale Combes Motel, and Eustaquio J. Reis. Property rights and deforestation in the Brazilian Amazon. *Ecological Economics*, 68(8-9):2461–2468, June 2009. ISSN 09218009. doi: 10.1016/j.ecolecon.2008.12.015.

R. A. Arriagada, P. J. Ferraro, E. O. Sills, S. K. Pattanayak, and S. Cordero-Sancho. Do Payments for Environmental Services Affect Forest Cover? A Farm-Level Evaluation from Costa Rica. *Land Economics*, 88(2):382–399, May 2012. ISSN 0023-7639, 1543-8325. doi: 10.3368/le.88.2.382.

KG Austin, JS Baker, BL Sohngen, CM Wade, A Daigneault, SB Ohrel, S Ragnauth, and A Bean. The economic costs of planting, preserving, and managing the world's forests to mitigate climate change. *Nature communications*, 11(1):1–9, 2020.

Andre Fernandes Tomon Avelino, Kathy Baylis, and Jordi Honey-Rosés. Goldilocks and the Raster Grid: Selecting Scale when Evaluating Conservation Programs. *PLOS ONE*, 11(12):e0167945, December 2016. ISSN 1932-6203. doi: 10.1371/journal.pone.0167945.

Christian Baehr, Ariel BenYishay, and Bradley Parks. Linking Local Infrastructure Development and Deforestation: Evidence from Satellite and Administrative Data. *Journal of the Association of Environmental and Resource Economists*, 8(2):375–409, March 2021. ISSN 2333-5955, 2333-5963. doi: 10.1086/712800.

Kathy Baylis, Jordi Honey, and M Isabel Ramírez. Conserving Forests: Mandates, Management or Money? *Selected Paper prepared for presentation at the Agricultural & Applied Economics Association's 2012 AAEA Annual Meeting*, page 17, August 2012.

Kathy Baylis, Jordi Honey-Rosés, Jan Börner, Esteve Corbera, Driss Ezzine-de Blas, Paul J. Ferraro, Renaud Lapeyre, U. Martin Persson, Alex Pfaff, and Sven Wunder. Mainstreaming Impact Evaluation in Nature Conservation. *Conservation Letters*, 9 (1):58–64, 2016. ISSN 1755-263X. doi: 10.1111/conl.12180.

Ariel BenYishay, Silke Heuser, Daniel Runfola, and Rachel Trichler. Indigenous land rights and deforestation: Evidence from the Brazilian Amazon. *Journal of Environmental Economics and Management*, 86:29–47, November 2017. ISSN 00950696. doi: 10.1016/j.jeem.2017.07.008.

Allen Blackman. Evaluating forest conservation policies in developing countries using remote sensing data: An introduction and practical guide. *Forest Policy and Economics*, 34:1–16, September 2013. ISSN 1389-9341. doi: 10.1016/j.forpol.2013.04.006.

Allen Blackman. Strict versus mixed-use protected areas: Guatemala's Maya Biosphere Reserve. *Ecological Economics*, 112:14–24, April 2015. ISSN 09218009. doi: 10.1016/j.ecolecon.2015.01.009.

Allen Blackman, Leonardo Corral, Eirivelthon Santos Lima, and Gregory P. Asner. Titling indigenous communities protects forests in the Peruvian Amazon. *Proceedings of the National Academy of Sciences*, 114(16):4123–4128, April 2017. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1603290114.

Allen Blackman, Leonard Goff, and Marisol Rivera Planter. Does eco-certification stem tropical deforestation? Forest Stewardship Council certification in Mexico. *Journal of Environmental Economics and Management*, 89:306–333, May 2018. ISSN 0095-0696. doi: 10.1016/j.jeem.2018.04.005.

Bryan Bollinger, Kenneth Gillingham, A. Justin Kirkpatrick, and Steven Sexton. Visibility and Peer Influence in Durable Good Adoption. *Marketing Science*, 41(3):453–476, May 2022. ISSN 0732-2399, 1526-548X. doi: 10.1287/mksc.2021.1306. URL `http://pubsonline.informs.org/doi/10.1287/mksc.2021.1306`.

Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting Event Study Designs: Robust and Efficient Estimation. *arXiv:2108.12419 [econ]*, August 2021. URL `http://arxiv.org/abs/2108.12419`. arXiv: 2108.12419.

Brett A. Bryan, Rebecca K. Runting, Tim Capon, Michael P. Perring, Shaun C. Cunningham, Marit E. Kragt, Martin Nolan, Elizabeth A. Law, Anna R. Renwick, Sue Eber, Rochelle Christian, and Kerrie A. Wilson. Designer policy for carbon and biodiversity co-benefits under global change. *Nature Climate Change*, 6(3):301–305, March 2016. ISSN 1758-678X, 1758-6798. doi: 10.1038/nclimate2874. URL `http://www.nature.com/articles/nclimate2874`.

Jonah Busch and Kalifi Ferretti-Gallon. What Drives Deforestation and What Stops It? A Meta-Analysis. *Review of Environmental Economics and Policy*, 11(1):3–23, January 2017. ISSN 1750-6816, 1750-6824. doi: 10.1093/reep/rew013. URL `https://www.journals.uchicago.edu/doi/10.1093/reep/rew013`.

Jonah Busch, Kalifi Ferretti-Gallon, Jens Engelmann, Max Wright, Kemen G. Austin, Fred Stolle, Svetlana Turubanova, Peter V. Potapov, Belinda Margono, Matthew C. Hansen, and Alessandro Baccini. Reductions in emissions from deforestation from Indonesia's moratorium on new oil palm, timber, and logging concessions. *Proceedings of the National Academy of Sciences*, 112(5):1328–1333, February 2015. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1412514112.

Jonah Busch, Jens Engelmann, Susan C. Cook-Patton, Bronson W. Griscom, Timm Kroeger, Hugh Possingham, and Priya Shyamsundar. Potential for low-cost carbon dioxide removal through tropical reforestation. *Nature Climate Change*, 9(6):463–466, June 2019. ISSN 1758-678X, 1758-6798. doi: 10.1038/s41558-019-0485-x. URL http://www.nature.com/articles/s41558-019-0485-x.

Van Butsic, David J. Lewis, Volker C. Radeloff, Matthias Baumann, and Tobias Kuemmerle. Quasi-experimental methods enable stronger inferences from observational data in ecology. *Basic and Applied Ecology*, 19:1–10, March 2017a. ISSN 1439-1791. doi: 10.1016/j.baae.2017.01.005.

Van Butsic, Catalina Munteanu, Patrick Griffiths, Jan Knorn, Volker C. Radeloff, Juraj Lieskovský, Daniel Mueller, and Tobias Kuemmerle. The effect of protected areas on forest disturbance in the Carpathian Mountains 1985-2010: Carpathian Protected Areas. *Conservation Biology*, 31(3):570–580, June 2017b. ISSN 08888892. doi: 10.1111/cobi.12835.

Jan Börner, Kathy Baylis, Esteve Corbera, Driss Ezzine-de Blas, Jordi Honey-Rosés, U. Martin Persson, and Sven Wunder. The Effectiveness of Payments for Environmental Services. *World Development*, 96:359–374, August 2017. ISSN 0305750X. doi: 10.1016/j.worlddev.2017.03.020. URL https://linkinghub.elsevier.com/retrieve/pii/S0305750X17300827.

Jan Börner, Dario Schulz, Sven Wunder, and Alexander Pfaff. The Effectiveness of Forest Conservation Policies and Programs. *Annual Review of Resource Economics*, 12(1):45–64, October 2020. ISSN 1941-1340, 1941-1359. doi: 10.1146/annurev-resource-110119-025703. URL https://www.annualreviews.org/doi/10.1146/annurev-resource-110119-025703.

Brantly Callaway and Pedro H.C. Sant'Anna. Difference-in-Differences with multiple time periods. *Journal of Econometrics*, page S0304407620303948, December 2020. ISSN 03044076. doi: 10.1016/j.jeconom.2020.12.001. URL https://linkinghub.elsevier.com/retrieve/pii/S0304407620303948.

Brantly Callaway, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. Difference-in-Differences with a Continuous Treatment, July 2021. URL http://arxiv.org/abs/2107.02637. arXiv:2107.02637 [econ].

Kimberly M. Carlson, Robert Heilmayr, Holly K. Gibbs, Praveen Noojipady, David N. Burns, Douglas C. Morton, Nathalie F. Walker, Gary D. Paoli, and Claire Kremen. Effect of oil palm sustainability certification on deforestation and fire in Indonesia. *Proceedings of the National Academy of Sciences*, 115(1):121–126, January 2018. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1704728114.

Robin Chazdon and Pedro Brancalion. Restoring forests as a means to many ends. *Science*, 365(6448):24–25, 2019.

Elías Cisneros, Jan Börner, Stefano Pagiola, and Sven Wunder. Impacts of conservation incentives in protected areas: The case of Bolsa Floresta, Brazil. *Journal of Environmental Economics and Management*, 111:102572, January 2022. ISSN 00950696. doi: 10.1016/j.jeem.2021.102572. URL https://linkinghub.elsevier.com/retrieve/pii/S0095069621001200.

Roger Alex Clapp. Regions of refuge and the agrarian question: Peasant agriculture and plantation forestry in Chilean araucanía. *World Development*, 26(4):571–589, April 1998. ISSN 0305750X. doi: 10.1016/S0305-750X(98)00010-2. URL https://linkinghub.elsevier.com/retrieve/pii/S0305750X98000102.

Eric A. Coleman, Bill Schultz, Vijay Ramprasad, Harry Fischer, Pushpendra Rana, Anthony M. Filippi, Burak Güneralp, Andong Ma, Claudia Rodriguez Solorzano, Vijay Guleria, Rajesh Rana, and Forrest Fleischman. Limited effects of tree planting on forest canopy cover and rural livelihoods in Northern India. *Nature Sustainability*, 4(11): 997–1004, November 2021. ISSN 2398-9629. doi: 10.1038/s41893-021-00761-z. URL https://www.nature.com/articles/s41893-021-00761-z.

CONAF. Evaluación del Funcionamiento del Fondo de Conservación, Recuperación y Manejo Sustentable del Bosque Nativo Ley N°20.283. Technical Report Licitación: 633-38-LE19, Corporación Nacional Forestal de Chile, 2020. URL https://concursolbn.conaf.cl/ayuda/2022/784_Resumen_ejecutivo-Evaluacion_del_fondo_conservacion_2020.pdf.

Payam Dadvand, Mark J. Nieuwenhuijsen, Mikel Esnaola, Joan Forns, Xavier Basagaña, Mar Alvarez-Pedrerol, Ioar Rivas, Mónica López-Vicente, Montserrat De Castro Pascual, Jason Su, Michael Jerrett, Xavier Querol, and Jordi Sunyer. Green spaces and cognitive development in primary schoolchildren. *Proceedings of the National Academy of Sciences*, 112(26):7937–7942, June 2015. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1503402112. URL https://pnas.org/doi/full/10.1073/pnas.1503402112.

Clément de Chaisemartin and Xavier D'Haultfœuille. Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9):2964–2996, September 2020. ISSN 0002-8282. doi: 10.1257/aer.20181169.

Clement de Chaisemartin and Xavier D'Haultfœuille. Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey. page 32, 2022.

Clément de Chaisemartin and Xavier D'Haultfœuille. Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9):2964–2996, September 2020. ISSN 0002-8282. doi: 10.1257/aer.20181169. URL https://pubs.aeaweb.org/doi/10.1257/aer.20181169.

Hannah Druckenmiller. Estimating an Economic and Social Value for Healthy Forests: Evidence from Tree Mortality in the American West. page 52, 2020.

J. E. Dyment and A. C. Bell. Grounds for movement: green school grounds as sites for promoting physical activity. *Health Education Research*, 23(6):952–962, October 2007. ISSN 0268-1153, 1465-3648. doi: 10.1093/her/cym059. URL https://academic.oup.com/her/article-lookup/doi/10.1093/her/cym059.

Ryan B Edwards, Walter P Falcon, Gracia Hadiwidjaja, Matthew M Higgins, Rosamond L Naylor, and Sudarno Sumarto. Fight fire with finance: A randomized field experiment to curtail land-clearing fire in Indonesia. page 62, 2020.

Frank Emmert-Streib and Matthias Dehmer. Introduction to Survival Analysis in Practice. *Machine Learning and Knowledge Extraction*, 1(3):1013–1038, September 2019. ISSN 2504-4990. doi: 10.3390/make1030058. URL https://www.mdpi.com/2504-4990/1/3/58.

Rebecca S. Epanchin-Niell. Economics of invasive species policy and management. *Biological Invasions*, 19(11):3333–3354, November 2017. ISSN 1387-3547, 1573-1464. doi: 10.1007/s10530-017-1406-4. URL http://link.springer.com/10.1007/s10530-017-1406-4.

F. España, R. Arriagada, O. Melo, and W. Foster. Forest plantation subsidies: Impact evaluation of the Chilean case. *Forest Policy and Economics*, 137:102696, April 2022. ISSN 13899341. doi: 10.1016/j.forpol.2022.102696. URL https://linkinghub.elsevier.com/retrieve/pii/S1389934122000089.

Andrea Faber Taylor and Frances E. Kuo. Children With Attention Deficits Concentrate Better After Walk in the Park. *Journal of Attention Disorders*, 12(5):402–409, March 2009. ISSN 1087-0547, 1557-1246. doi: 10.1177/1087054708323000. URL http://journals.sagepub.com/doi/10.1177/1087054708323000.

Stephen C. Farber, Robert Costanza, and Matthew A. Wilson. Economic and ecological concepts for valuing ecosystem services. *Ecological Economics*, 41(3):375–392, June 2002. ISSN 09218009. doi: 10.1016/S0921-8009(02)00088-5. URL https://linkinghub.elsevier.com/retrieve/pii/S0921800902000885.

Eli P. Fenichel, Timothy J. Richards, and David W. Shanafelt. The Control of Invasive Species on Private Property with Neighbor-to-Neighbor Spillovers. *Environmental and Resource Economics*, 59(2):231–255, October 2014. ISSN 0924-6460, 1573-1502. doi: 10.1007/s10640-013-9726-z. URL http://link.springer.com/10.1007/s10640-013-9726-z.

P. J. Ferraro and M. M. Hanauer. Quantifying causal mechanisms to determine how protected areas affect poverty through changes in ecosystem services and infrastructure. *Proceedings of the National Academy of Sciences*, 111(11):4332–4337, March 2014. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1307712111.

Paul J. Ferraro, James N. Sanchirico, and Martin D. Smith. Causal inference in coupled human and natural systems. *Proceedings of the National Academy of Sciences*, 116(12): 5311–5318, March 2019a. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1805563115.

Paul J. Ferraro, James N. Sanchirico, and Martin D. Smith. Causal inference in coupled human and natural systems. *Proceedings of the National Academy of Sciences*, 116(12): 5311–5318, March 2019b. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1805563115. URL `https://pnas.org/doi/full/10.1073/pnas.1805563115`.

Eyal Frank and Anant Sudarshan. The Social Costs of Keystone Species Collapse: Evidence from the Decline of Vultures in India. *SSRN Electronic Journal*, 2023. ISSN 1556-5068. doi: 10.2139/ssrn.4318579. URL `https://www.ssrn.com/abstract=4318579`.

Jed Friedman and Norbert Schady. How Many Infants Likely Died In Africa As A Result Of The 2008-2009 Global Financial Crisis?: Excess Infant Mortality In Africa Due To The Global Financial Crisis. *Health Economics*, 22(5):611–622, May 2013. ISSN 10579230. doi: 10.1002/hec.2818. URL `http://doi.wiley.com/10.1002/hec.2818`.

John Gardner. Two-stage differences in differences. *Working Paper*, page 34, April 2021.

Ed Gerrish and Shannon Lea Watkins. The relationship between urban forests and income: A meta-analysis. *Landscape and Urban Planning*, 170:293–308, February 2018. ISSN 01692046. doi: 10.1016/j.landurbplan.2017.09.005. URL `https://linkinghub.elsevier.com/retrieve/pii/S0169204617302062`.

John Gibson, Susan Olivia, Geua Boe-Gibson, and Chao Li. Which night lights data should we use in economics, and where? *Journal of Development Economics*, 149: 102602, 2021. ISSN 03043878. doi: 10.1016/j.jdeveco.2020.102602. URL `https://linkinghub.elsevier.com/retrieve/pii/S0304387820301772`.

Leah Gichuki, Rens Brouwer, Jonathan Davies, Adriana Vidal, Mirjam Kuzee, Chris Magero, Sven Walter, Pedro Lara, Christiana Oragbade, and Ben Gilbey. *Reviving land and restoring landscapes: policy convergence between forest landscape restoration and land degradation neutrality.* IUCN, International Union for Conservation of Nature, August 2019. ISBN 978-2-8317-1991-7. doi: 10.2305/IUCN.CH.2019.11.en. URL `https://portals.iucn.org/library/node/48515`.

Anita Gidlöf-Gunnarsson and Evy Öhrström. Noise and well-being in urban residential environments: The potential role of perceived availability to nearby green areas. *Landscape and Urban Planning*, 83(2-3):115–126, November 2007. ISSN 01692046.

doi: 10.1016/j.landurbplan.2007.03.003. URL https://linkinghub.elsevier.com/retrieve/pii/S0169204607000722.

Renzo Giudice, Jan Börner, Sven Wunder, and Elias Cisneros. Selection biases and spillovers from collective conservation incentives in the Peruvian Amazon. *Environmental Research Letters*, 14(4):045004, April 2019. ISSN 1748-9326. doi: 10.1088/1748-9326/aafc83. URL https://iopscience.iop.org/article/10.1088/1748-9326/aafc83.

Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, page S0304407621001445, June 2021. ISSN 03044076. doi: 10.1016/j.jeconom.2021.03.014.

Jordan Graesser, Radost Stanimirova, Katelyn Tarrio, Esteban J. Copati, José N. Volante, Santiago R. Verón, Santiago Banchero, Hernan Elena, Diego de Abelleyra, and Mark A. Friedl. Temporally-Consistent Annual Land Cover from Landsat Time Series in the Southern Cone of South America. *Remote Sensing*, 14(16):4005, August 2022. ISSN 2072-4292. doi: 10.3390/rs14164005. URL https://www.mdpi.com/2072-4292/14/16/4005.

Merlin M. Hanauer and Gustavo Canavire-Bacarreza. Implications of heterogeneous impacts of protected areas on deforestation and poverty. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1681):20140272, November 2015. ISSN 0962-8436, 1471-2970. doi: 10.1098/rstb.2014.0272. URL https://royalsocietypublishing.org/doi/10.1098/rstb.2014.0272.

M. C. Hansen, P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, A. Tyukavina, D. Thau, S. V. Stehman, S. J. Goetz, T. R. Loveland, A. Kommareddy, A. Egorov, L. Chini, C. O. Justice, and J. R. G. Townshend. High-Resolution Global Maps of 21st-Century Forest Cover Change. *Science*, 342(6160):850–853, November 2013. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1244693.

Richard J. Hauer and Ward D. Peterson. Effects of emerald ash borer on municipal forestry budgets. *Landscape and Urban Planning*, 157:98–105, January 2017. ISSN 01692046. doi: 10.1016/j.landurbplan.2016.05.023. URL https://linkinghub.elsevier.com/retrieve/pii/S0169204616300901.

Robert Heilmayr and Eric F. Lambin. Impacts of nonstate, market-driven governance on Chilean forests. *Proceedings of the National Academy of Sciences*, 113(11):2910–2915, March 2016. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1600394113.

Robert Heilmayr, Cristian Echeverría, and Eric F. Lambin. Impacts of Chilean forest subsidies on forest cover, carbon and biodiversity. *Nature Sustainability*, 3(9):701–709, September 2020a. ISSN 2398-9629. doi: 10.1038/s41893-020-0547-0. URL http://www.nature.com/articles/s41893-020-0547-0.

Robert Heilmayr, Lisa L. Rausch, Jacob Munger, and Holly K. Gibbs. Brazil's amazon soy moratorium reduced deforestation. *Nature Food*, 1(12):801–810, December 2020b. doi: 10.1038/s43016-020-00194-5. URL `https://doi.org/10.1038/s43016-020-00194-5`.

Daniel A. Herms and Deborah G. McCullough. Emerald Ash Borer Invasion of North America: History, Biology, Ecology, Impacts, and Management. *Annual Review of Entomology*, 59(1):13–30, January 2014. ISSN 0066-4170, 1545-4487. doi: 10.1146/annurev-ento-011613-162051. URL `https://www.annualreviews.org/doi/10.1146/annurev-ento-011613-162051`.

Diego Herrera, Alexander Pfaff, and Juan Robalino. Impacts of protected areas vary with the level of government: Comparing avoided deforestation across agencies in the Brazilian Amazon. *Proceedings of the National Academy of Sciences*, 116(30):14916–14925, July 2019. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1802877116.

Margaret B. Holland, Kelly W. Jones, Lisa Naughton-Treves, José-Luis Freire, Manuel Morales, and Luis Suárez. Titling land to conserve forests: The case of Cuyabeno Reserve in Ecuador. *Global Environmental Change*, 44:27–38, May 2017. ISSN 09593780. doi: 10.1016/j.gloenvcha.2017.02.004.

Paul W. Holland. Statistics and Causal Inference. *Journal of the American Statistical Association*, 81(396):945–960, December 1986. ISSN 0162-1459, 1537-274X. doi: 10.1080/01621459.1986.10478354.

Sam Hooper and Robert E. Kennedy. A spatial ensemble approach for broad-area mapping of land surface properties. *Remote Sensing of Environment*, 210:473–489, June 2018. ISSN 00344257. doi: 10.1016/j.rse.2018.03.032. URL `https://linkinghub.elsevier.com/retrieve/pii/S0034425718301330`.

Kosuke Imai and In Song Kim. On the use of two-way fixed effects regression models for causal inference with panel data. *Political Analysis*, 29(3):405–415, 2021. doi: 10.1017/pan.2020.33.

B. K. Jack, C. Kousky, and K. R. E. Sims. Designing payments for ecosystem services: Lessons from previous experience with incentive-based mechanisms. *Proceedings of the National Academy of Sciences*, 105(28):9465–9470, July 2008. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.0705503104. URL `http://www.pnas.org/cgi/doi/10.1073/pnas.0705503104`.

B. Kelsey Jack. Private Information and the Allocation of Land Use Subsidies in Malawi. *American Economic Journal: Applied Economics*, 5(3):113–135, July 2013. ISSN 1945-7782, 1945-7790. doi: 10.1257/app.5.3.113. URL `http://pubs.aeaweb.org/doi/10.1257/app.5.3.113`.

B. Kelsey Jack and Elsa Cardona Santos. The leakage and livelihood impacts of PES contracts: A targeting experiment in Malawi. *Land Use Policy*, 63:645–658, April 2017. ISSN 02648377. doi: 10.1016/j.landusepol.2016.03.028. URL `https://linkinghub.elsevier.com/retrieve/pii/S026483771630268X`.

B. Kelsey Jack and Seema Jayachandran. Self-selection into payments for ecosystem services programs. *Proceedings of the National Academy of Sciences*, 116(12):5326–5333, March 2019. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1802868115. URL `http://www.pnas.org/lookup/doi/10.1073/pnas.1802868115`.

B Kelsey Jack, Seema Jayachandran, Namrata Kala, and Rohini Pande. Money (not) to burn: Payments for ecosystem services to reduce crop residue burning. 2022.

Meha Jain. The Benefits and Pitfalls of Using Satellite Data for Causal Inference. *Review of Environmental Economics and Policy*, 14(1):157–169, January 2020. ISSN 1750-6816, 1750-6824. doi: 10.1093/reep/rez023.

Seema Jayachandran. The Inherent Trade-Off Between the Environmental and Anti-Poverty Goals of Payments for Ecosystem Services. page 20, 2022.

Seema Jayachandran, Joost de Laat, Eric F. Lambin, Charlotte Y. Stanton, Robin Audy, and Nancy E. Thomas. Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation. *Science*, 357(6348):267–273, July 2017. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.aan0568.

Benjamin A. Jones. Invasive Species Control, Agricultural Pesticide Use, and Infant Health Outcomes. *Land Economics*, 96(2):149–170, March 2020. ISSN 0023-7639, 1543-8325. doi: 10.3368/le.96.2.149. URL `http://le.uwpress.org/lookup/doi/10.3368/le.96.2.149`.

Benjamin A. Jones. Can invasive species lead to sedentary behavior? The time use and obesity impacts of a forest-attacking pest. *Journal of Environmental Economics and Management*, 119:102800, May 2023. ISSN 00950696. doi: 10.1016/j.jeem.2023.102800. URL `https://linkinghub.elsevier.com/retrieve/pii/S0095069623000189`.

Benjamin A. Jones and Shana M. McDermott. Health Impacts of Invasive Species Through an Altered Natural Environment: Assessing Air Pollution Sinks as a Causal Pathway. *Environmental and Resource Economics*, 71(1):23–43, September 2018. ISSN 0924-6460, 1573-1502. doi: 10.1007/s10640-017-0135-6. URL `http://link.springer.com/10.1007/s10640-017-0135-6`.

Kelly W. Jones and David J. Lewis. Estimating the Counterfactual Impact of Conservation Programs on Land Cover Outcomes: The Role of Matching and Panel Regression Techniques. *PLOS ONE*, 10(10):e0141380, October 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0141380.

Kelly W. Jones, Margaret B. Holland, Lisa Naughton-Treves, Manuel Morales, Luis Suarez, and Kayla Keenan. Forest conservation incentives and deforestation in the Ecuadorian Amazon. *Environmental Conservation*, 44(1):56–65, March 2017. ISSN 0376-8929, 1469-4387. doi: 10.1017/S0376892916000308.

Suzi Kerr, Shuguang Liu, Alexander S.P. Pfaff, and R.Flint Hughes. Carbon dynamics and land-use choices: Building a regional-scale multidisciplinary model. *Journal of Environmental Management*, 69(1):25–37, September 2003. ISSN 03014797. doi: 10.1016/S0301-4797(03)00106-3.

Nicolas Koch. Agricultural Productivity and Forest Conservation: Evidence from the Brazilian Amazon. *American Journal of Agricultural Economics*, page 22, 2019.

Byoung-Suk Kweon, Christopher D. Ellis, Junga Lee, and Kim Jacobs. The link between school environments and student academic performance. *Urban Forestry & Urban Greening*, 23:35–43, April 2017. ISSN 16188667. doi: 10.1016/j.ufug.2017.02.002. URL https://linkinghub.elsevier.com/retrieve/pii/S1618866716300140.

Ashley E. Larsen, Kyle Meng, and Bruce E. Kendall. Causal analysis in control–impact ecological studies with observational data. *Methods in Ecology and Evolution*, 10(7): 924–934, 2019. ISSN 2041-210X. doi: 10.1111/2041-210X.13190.

Simon L Lewis, Charlotte E Wheeler, Edward TA Mitchard, and Alexander Koch. Restoring natural forests is the best way to remove atmospheric carbon, 2019.

Dongying Li and William C. Sullivan. Impact of views to school landscapes on recovery from stress and mental fatigue. *Landscape and Urban Planning*, 148:149–158, April 2016. ISSN 01692046. doi: 10.1016/j.landurbplan.2015.12.015. URL https://linkinghub.elsevier.com/retrieve/pii/S0169204615002571.

L. Lipper, T. Sakuyama, R. Stringer, and D. Zilberman. *Payment for Environmental Services in Agricultural Landscapes: Economic Policies and Poverty Reduction in Developing Countries*. Natural Resource Management and Policy. Springer New York, 2009. ISBN 978-0-387-72971-8. URL https://books.google.com/books?id=JkXzgldmkvsC.

Jeremy Luallen, Jared Edgerton, and Deirdre Rabideau. A Quasi-Experimental Evaluation of the Impact of Public Assistance on Prisoner Recidivism. *Journal of Quantitative Criminology*, 34(3):741–773, September 2018. ISSN 0748-4518, 1573-7799. doi: 10.1007/s10940-017-9353-x. URL http://link.springer.com/10.1007/s10940-017-9353-x.

Macomb EAB Readiness Plan, 2007.

Dave E. Marcotte. Something in the air? Air quality and children's educational outcomes. *Economics of Education Review*, 56:141–151, February 2017. ISSN 02727757. doi: 10.1016/j.econedurev.2016.12.003. URL https://linkinghub.elsevier.com/retrieve/pii/S0272775716303703.

Valérie Masson-Delmotte, Panmao Zhai, Hans-Otto Pörtner, Debra Roberts, Jim Skea, Priyadarshi R Shukla, Anna Pirani, Wilfran Moufouma-Okia, Clotilde Péan, Roz Pidcock, et al. Global warming of 1.5 c. *An IPCC Special Report on the impacts of global warming of*, 1(5), 2018.

Giovanni Mastrobuoni and Paolo Pinotti. Legal Status and the Criminal Activity of Immigrants. *American Economic Journal: Applied Economics*, 7(2):175–206, April 2015. ISSN 1945-7782, 1945-7790. doi: 10.1257/app.20140039. URL https://pubs.aeaweb.org/doi/10.1257/app.20140039.

Jennifer D. McCabe, He Yin, Jennyffer Cruz, Volker Radeloff, Anna Pidgeon, David N. Bonter, and Benjamin Zuckerberg. Prey abundance and urbanization influence the establishment of avian predators in a metropolitan landscape. *Proceedings of the Royal Society B: Biological Sciences*, 285(1890):20182120, November 2018. ISSN 0962-8452, 1471-2954. doi: 10.1098/rspb.2018.2120. URL https://royalsocietypublishing.org/doi/10.1098/rspb.2018.2120.

Daniela A. Miteva, Subhrendu K. Pattanayak, and Paul J. Ferraro. Evaluation of biodiversity policy instruments: What works and what doesn't? *Oxford Review of Economic Policy*, 28(1):69–92, March 2012. ISSN 0266-903X, 1460-2121. doi: 10.1093/oxrep/grs009.

Morton Arboretum. 2020 Chicago Region Tree Census. Technical report, 2020. URL https://mortonarb.org/science/tree-census/.

Christoph Nolte, Beatriz Gobbi, Van Butsic, and Eric F Lambin. Decentralized Land Use Zoning Reduces Large-scale Deforestation in a Major Agricultural Frontier. *Ecological Economics*, page 11, 2017.

David J. Nowak, Daniel E. Crane, and Jack C. Stevens. Air pollution removal by urban trees and shrubs in the United States. *Urban Forestry & Urban Greening*, 4(3-4): 115–123, April 2006. ISSN 16188667. doi: 10.1016/j.ufug.2006.01.007. URL https://linkinghub.elsevier.com/retrieve/pii/S1618866706000173.

David J. Nowak, Robert E. Hoehn, Allison R. Bodine, Eric J. Greenfield, and Jarlath O'Neil-Dunne. Urban forest structure, ecosystem services and change in Syracuse, NY. *Urban Ecosystems*, 19(4):1455–1477, December 2016. ISSN 1083-8155, 1573-1642. doi: 10.1007/s11252-013-0326-z. URL http://link.springer.com/10.1007/s11252-013-0326-z.

Paulina Oliva, B. Kelsey Jack, Samuel Bell, Elizabeth Mettetal, and Christopher Severen. Technology Adoption under Uncertainty: Take-Up and Subsequent Investment in Zambia. *The Review of Economics and Statistics*, 102(3):617–632, June 2020. ISSN 0034-6535, 1530-9142. doi: 10.1162/rest_a_00823. URL https://direct.mit.edu/rest/article/102/3/617-632/96771.

Stefano Pagiola, Ana R. Rios, and Agustin Arcenas. Can the poor participate in payments for environmental services? Lessons from the Silvopastoral Project in Nicaragua. *Environment and Development Economics*, 13(3):299–325, June 2008. ISSN 1355-770X, 1469-4395. doi: 10.1017/S1355770X08004270. URL https://www.cambridge.org/core/product/identifier/S1355770X08004270/type/journal_article.

Stephanie Panlasigui. Impacts of certification, uncertified concessions, and protected areas on forest loss in Cameroon, 2000 to 2013. *Biological Conservation*, page 7, 2018.

R. Jisung Park. Hot Temperature and High-Stakes Performance. *Journal of Human Resources*, 57(2):400–434, March 2022. ISSN 0022-166X, 1548-8004. doi: 10.3368/jhr.57.2.0618-9535R3. URL http://jhr.uwpress.org/lookup/doi/10.3368/jhr.57.2.0618-9535R3.

R. Jisung Park, A. Patrick Behrer, and Joshua Goodman. Learning is inhibited by heat exposure, both internationally and within the United States. *Nature Human Behaviour*, 5(1):19–27, October 2020a. ISSN 2397-3374. doi: 10.1038/s41562-020-00959-9. URL https://www.nature.com/articles/s41562-020-00959-9.

R. Jisung Park, Joshua Goodman, Michael Hurwitz, and Jonathan Smith. Heat and Learning. *American Economic Journal: Economic Policy*, 12(2):306–339, May 2020b. ISSN 1945-7731, 1945-774X. doi: 10.1257/pol.20180612. URL https://pubs.aeaweb.org/doi/10.1257/pol.20180612.

Alexander Pfaff, Suzi Kerr, Leslie Lipper, Romina Cavatassi, Benjamin Davis, Joanna Hendy, and G. Arturo Sanchez-Azofeifa. Will buying tropical forest carbon benefit the poor? Evidence from Costa Rica. *Land Use Policy*, 24(3):600–610, July 2007. ISSN 02648377. doi: 10.1016/j.landusepol.2006.01.003. URL https://linkinghub.elsevier.com/retrieve/pii/S026483770600069X.

Alexander Pfaff, Juan Robalino, Eirivelthon Lima, Catalina Sandoval, and Luis Diego Herrera. Governance, Location and Avoided Deforestation from Protected Areas: Greater Restrictions Can Have Lower Impact, Due to Differences in Location. *World Development*, 55:7–20, March 2014. ISSN 0305750X. doi: 10.1016/j.worlddev.2013.01.011. URL https://linkinghub.elsevier.com/retrieve/pii/S0305750X1300017X.

Alexander Pfaff, Juan Robalino, Diego Herrera, and Catalina Sandoval. Protected Areas' Impacts on Brazilian Amazon Deforestation: Examining Conservation – Development

Interactions to Inform Planning. *PLOS ONE*, 10(7):e0129460, July 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0129460.

Alexander S. P. Pfaff and G. Arturo Sanchez-Azofeifa. Deforestation pressure and biological reserve planning: A conceptual approach and an illustrative application for Costa Rica. *Resource and Energy Economics*, 26(2):237–254, June 2004. ISSN 0928-7655. doi: 10.1016/j.reseneeco.2003.11.009.

Alexander S.P. Pfaff. What Drives Deforestation in the Brazilian Amazon? *Journal of Environmental Economics and Management*, 37(1):26–43, January 1999. ISSN 00950696. doi: 10.1006/jeem.1998.1056. URL `https://linkinghub.elsevier.com/retrieve/pii/S0095069698910567`.

Jonathan Proctor, Tamma Carleton, and Sandy Sum. Parameter Recovery Using Remotely Sensed Variables. *NBER WORKING PAPER SERIES*, Working Paper 30861, 2023.

Eli Puterman, Jordan Weiss, Benjamin A. Hives, Alison Gemmill, Deborah Karasek, Wendy Berry Mendes, and David H. Rehkopf. Predicting mortality from 57 economic, behavioral, social, and psychological factors. *Proceedings of the National Academy of Sciences*, 117(28):16273–16282, July 2020. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1918455117. URL `https://pnas.org/doi/full/10.1073/pnas.1918455117`.

Jimena Rico-Straffon, Zhenhua Wang, Stephanie Panlasigui, Colby J. Loucks, Jennifer Swenson, and Alexander Pfaff. Forest concessions and eco-certifications in the Peruvian Amazon: Deforestation impacts of logging rights and logging restrictions. *Journal of Environmental Economics and Management*, 118:102780, March 2023. ISSN 00950696. doi: 10.1016/j.jeem.2022.102780. URL `https://linkinghub.elsevier.com/retrieve/pii/S0095069622001334`.

J. Robalino and A. Pfaff. Ecopayments and Deforestation in Costa Rica: A Nationwide Analysis of PSA's Initial Years. *Land Economics*, 89(3):432–448, August 2013. ISSN 0023-7639, 1543-8325. doi: 10.3368/le.89.3.432.

Juan Robalino, Alexander Pfaff, and Laura Villalobos. Deforestation spillovers from Costa Rican protected areas. Working Paper 201502, Universidad de Costa Rica, 2015.

Jonathan Roth. Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends. *American Economic Review: Insights*, 4(3):305–322, September 2022. ISSN 2640-205X, 2640-2068. doi: 10.1257/aeri.20210236. URL `https://pubs.aeaweb.org/doi/10.1257/aeri.20210236`.

Jonathan Roth, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe. What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature, January 2022. URL http://arxiv.org/abs/2201.01194. arXiv:2201.01194 [econ, stat].

Vilane G. Sales, Eric Strobl, and Robert J.R. Elliott. Cloud cover and its impact on brazil's deforestation satellite monitoring program: Evidence from the cerrado biome of the brazilian legal amazon. *Applied Geography*, 140:102651, March 2022. doi: 10.1016/j.apgeog.2022.102651. URL https://doi.org/10.1016/j.apgeog.2022.102651.

Pedro H. C. Sant'Anna and Jun B. Zhao. Doubly Robust Difference-in-Differences Estimators. *arXiv:1812.01723 [econ]*, May 2020. URL http://arxiv.org/abs/1812.01723. arXiv: 1812.01723.

Payal Shah and Kathy Baylis. Evaluating Heterogeneous Conservation Effects of Forest Protection in Indonesia. *PLOS ONE*, 10(6):e0124872, June 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0124872.

Katharine R.E. Sims and Jennifer M. Alix-Garcia. Parks versus PES: Evaluating direct and incentive-based land conservation in Mexico. *Journal of Environmental Economics and Management*, 86:8–28, November 2017. ISSN 00950696. doi: 10.1016/j.jeem.2016.11.010.

Karyn Tabor, Kelly W. Jones, Jennifer Hewson, Andriambolantsoa Rasolohery, Andoniaina Rambeloson, Tokihenintsoa Andrianjohaninarivo, and Celia A. Harvey. Evaluating the effectiveness of conservation and development investments in reducing deforestation and fires in Ankeniheny-Zahemena Corridor, Madagascar. *PLOS ONE*, 12 (12):e0190119, December 2017. ISSN 1932-6203. doi: 10.1371/journal.pone.0190119.

Bing Yang Tan. Save a Tree and Save a Life: Estimating the Health Benefits of Urban Forests. *Environmental and Resource Economics*, 82(3):657–680, July 2022. ISSN 0924-6460, 1573-1502. doi: 10.1007/s10640-022-00677-y. URL https://link.springer.com/10.1007/s10640-022-00677-y.

Jan Tinbergen. *On the theory of economic policy*. North-Holland Publishing Company, Amsterdam, 1952. ISBN 978-0-7204-3130-8. URL hdl.handle.net/1765/15884.

Jessica B. Turner-Skoff and Nicole Cavender. The benefits of trees for livable and sustainable communities. *PLANTS, PEOPLE, PLANET*, 1(4):323–335, October 2019. ISSN 2572-2611, 2572-2611. doi: 10.1002/ppp3.39. URL https://onlinelibrary.wiley.com/doi/10.1002/ppp3.39.

E. Uchida, J. Xu, and S. Rozelle. Grain for Green: Cost-Effectiveness and Sustainability of China's Conservation Set-Aside Program. *Land Economics*, 81(2):247–264, May 2005. ISSN 0023-7639, 1543-8325. doi: 10.3368/le.81.2.247. URL http://le.uwpress.org/cgi/doi/10.3368/le.81.2.247.

USDA. Usda aphis emerald ash borer program manual. Technical report, 2020. URL https://mortonarb.org/science/tree-census/.

Jeff Wen and Marshall Burke. Lower test scores from wildfire smoke exposure. *Nature Sustainability*, 5(11):947–955, September 2022. ISSN 2398-9629. doi: 10.1038/s41893-022-00956-y. URL https://www.nature.com/articles/s41893-022-00956-y.

K. J. Wendland, M. Baumann, D. J. Lewis, A. Sieber, and V. C. Radeloff. Protected Area Effectiveness in European Russia: A Postmatching Panel Data Analysis. *Land Economics*, 91(1):149–168, February 2015. ISSN 0023-7639, 1543-8325. doi: 10.3368/le.91.1.149.

David R. Williams, Andrew Balmford, and David S. Wilcove. The past and future role of conservation science in saving biodiversity. *Conservation Letters*, n/a(n/a):e12720, 2020. ISSN 1755-263X. doi: 10.1111/conl.12720.

Jeffrey M. Wooldridge. Econometric Analysis of Cross Section and Panel Data, second edition, 2010.

Sven Wunder. Payments for environmental services and the poor: concepts and preliminary evidence. *Environment and Development Economics*, 13(3):279–297, June 2008. ISSN 1355-770X, 1469-4395. doi: 10.1017/S1355770X08004282. URL https://www.cambridge.org/core/product/identifier/S1355770X08004282/type/journal_article.

David Zilberman, Leslie Lipper, and Nancy Mccarthy. When could payments for environmental services benefit the poor? *Environment and Development Economics*, 13(3):255–278, June 2008. ISSN 1355-770X, 1469-4395. doi: 10.1017/S1355770X08004294. URL https://www.cambridge.org/core/product/identifier/S1355770X08004294/type/journal_article.