



WORKING PAPERS

N° 1703

January 2026

“An Evaluation of Protected Area Policies in the European Union”

Tristan Earle Grupp, Prakash Mishra, Mathias Reynaert
and Arthur A. van Benthem

An Evaluation of Protected Area Policies in the European Union

Tristan Earle Grupp, Prakash Mishra, Mathias Reynaert, Arthur A. van Benthem *

January 27, 2026

Abstract

The European Union designates 26% of its landmass as protected areas, limiting economic development for biodiversity. We use the staggered introduction of protected areas between 1985 and 2019 to study the selection of protected land and the causal effect of protection on vegetation cover and nightlights. We find no meaningful impacts on either outcome across four decades, countries, protection cohorts, or land characteristics. These null effects are consistent with the political economy of EU land protection: weak incentives to internalize biodiversity gains, green-glow motives, and area-based targets shape local siting and stringency choices. In practice, strict protection is applied where development pressure is low—so that protection has little bite—while in high-pressure regions, protection is typically weak, imposing only limited constraints on economic activity.

Keywords: land protection, protected areas, conservation, biodiversity, deforestation, vegetation cover, nightlights, staggered difference-in-differences, Europe

JEL codes: Q23, Q24, Q57, R14

*Earle Grupp: University of Pennsylvania; Mishra: University of Pennsylvania; Reynaert: Toulouse School of Economics; van Benthem: University of Pennsylvania. We thank Adam Streff and Danial Syed for excellent research assistance. Thanks to seminar participants at the AERE and EAERE 2023 summer conferences, CEPR, Cornell, Georgia Tech, Imperial College London, LSE/Imperial/King's Workshop in Environmental Economics, Paris School of Economics, SUNY Albany, Toulouse School of Economics, Tilburg University, Tulane University, UC Berkeley, UC Santa Barbara, and the Universities of Bristol, Cambridge, Delaware, Gothenburg, Mannheim, Maryland, Oxford, Pennsylvania, Wisconsin-Madison, UBC, UPF and Villanova. We thank the editor and four anonymous referees for their excellent suggestions, as well as Robin Burgess, Eli Fenichel, Alex Pfaff, Andrew Plantinga, Santiago Saavedra, Ulrich Wagner, and Matthew Wibbenmeyer for their helpful comments. We gratefully acknowledge financial support from FORMAS grant number 2020-00371. Reynaert acknowledges funding by the European Union (ERC, SPACETIME, grant n° 101077168). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. Reynaert acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program). Van Benthem thanks Penn Global, the Kleinman Center for Energy Policy, the Mack Institute, the Wharton Dean's Research Fund, the Wharton Global Initiatives Research Program, and Analytics at Wharton for generous support.

1 Introduction

How effective are protected area policies at restoring vegetation cover, constraining economic activity, and improving biodiversity more broadly? These questions are central to assessing the pledge of 196 countries to protect 30% of the earth’s land and waters by 2030. This ‘30x30 target’, which was the main outcome of the COP 15 Kunming-Montreal global biodiversity conference and sometimes referred to as the ‘Paris Agreement for Nature’, was ratified in December 2022 (Einhorn 2022). How ambitious is this 30% target? This depends strongly on how policymakers select land to protect and if protection successfully limits economic activity.

We study these questions with a focus on the European Union (EU), the region that is closest to reaching the target. As of 2023, the EU had protected over 26% of its total landmass (Eurostat 2022). As one of the world’s largest coordinated land-protection networks (Oceana Europe 2022), the EU’s flagship Natura 2000 policy may offer insights into how the global 30x30 commitment might perform.

We evaluate the causal effect of Europe’s protected area policies on an important dimension of biodiversity—vegetation cover—and on human presence in protected areas—measured by night-lights. We make three primary contributions. First, we provide unique estimates of the long-term impact of one of the world’s largest land-protection policies over four decades. Second, we combine theory and recent econometric advances to study the gradual staggered introduction of more than 100,000 protected areas over the course of many decades. We contrast how a social planner and the EU’s policy procedure could lead to different protection allocations and empirically test how the effect of protection varies across countries, over time since protection, and across earlier and later protected areas. We also estimate treatment-effect heterogeneity with unprecedented granularity across observable land, soil and climate characteristics. Finally, we provide evidence on the environmental impacts of land protection in advanced economies, which has been virtually non-existent to date.

We assemble a high-resolution remote-sensing dataset that spans the entirety of the European Union from 1985-2020. Our data include key outcome variables (vegetation cover, nightlights) at the 300x300 meter or 1x1 km level, treatment variables (location of protected areas and date of first protection), and a wide range of control variables that measure climate, weather, land, and soil characteristics. We use a continuous measure of vegetation greenness, Landsat Normalized Difference Vegetation Index (NDVI), as it reflects gradual changes in vegetation cover not captured by discrete land-use measures; it is also an imperfect yet reasonable indicator of other measures of biodiversity. We also collect alternative outcomes—discrete land use classes and species counts—for use in robustness analysis.

The data reveal four key facts about Europe’s land use and protection policies. First, large parts of Europe have undergone land-use change over the last 40 years. We find that some parts of Europe have been greening and reforesting, while other parts of Europe have seen vegetation

degradation (Winkler, Fuchs, Rounsevell, et al. 2021).¹ Second, forest loss and degradation take place on more populated and agriculturally suitable land. Agriculture and urbanization continue to cause deforestation, and grasslands across Europe have been converted into monocultures of high-yield rye. Indeed, the EU reports 81% of listed protected species and habitats have poor conservation status (European Environment Agency 2020). Third, EU land-protection targets greener areas and extends across both sparsely and densely populated regions. Fourth, protection is strictest on sparsely populated, initially green land, whereas most protected areas are subject to less strict protection when land is more populated and initially less green. Overall, Europe exhibits different land-use dynamics compared to heavily studied illegal deforestation in the tropics (Burgess et al. 2012; Balboni, Burgess, and Olken 2023; Burgess, Costa, and Olken 2023). In this context, it is an open question whether protection efforts have been effective in safeguarding nature under pressure or have contributed to expanding natural areas where land-use pressures have lessened.

Based on these observations, we model a decision maker who can protect new plots each year by restricting economic activity to increase ecological value. Contributing to existing models of optimal siting of protection (Weitzman 1998; Assunção et al. 2022), we first focus on the gradual increase in the protected surface and examine the treatment effects we expect if a social planner were to implement an optimal protection policy over time. The solution to the dynamic optimization problem yields four predictions. First, a social planner should protect only when biodiversity gains are present, thereby yielding positive treatment effects. Second, treatment effects could increase or decrease across successive cohorts of protection. Third, we contrast the social planner benchmark with a setting in which the EU delegates siting decisions to local policymakers. In our framework, the EU sets an area-based protection target, while local policymakers choose the location and the stringency of protection.² When the local policymaker places little weight on ecological outcomes but either derives green-glow benefits from designation itself or faces a binding quantity target for protected land, the model can generate protection paths with small or near-zero ecological treatment effects. Fourth, when protection stringency is a choice variable, a local policymaker facing these same incentives may optimally implement weak protection on many plots, resulting in near-zero ecological treatment effects.

To test these predictions, we employ a staggered difference-in-differences design. Because of selection into treatment, treated and untreated areas differ on several dimensions. Moreover, time-varying selection can cause treatment effects to vary by cohort (the calendar year in which land first gets protected) and event time (years relative to treatment). Typical two-way fixed effects estimators are biased in such settings (Chaisemartin and D'Haultfœuille 2020; Goodman-Bacon 2021; Sun and Abraham 2021). To overcome this challenge, we apply the doubly-robust estimator of Callaway

1. Forests have expanded by more than 30% since 1900, an area the size of Portugal (Eurostat 2021); an example of the “forest transition” (Barbier, Burgess, and Grainger 2010). Yet, this greening trend has slowed down substantially in more recent decades.

2. In practice, member states often further delegate implementation to provinces or municipalities. Our model abstracts from these layers and treats the local decision maker as a single entity representing all national and subnational authorities.

and Sant'Anna (2021). This estimator combines cohort-specific inverse probability matching with an outcome regression adjustment to compare protected areas with observably similar unprotected land and to control for time-varying differences in observable plot attributes.

In addition to time variation in treatment effects, previous literature—described below—has revealed that the effect of land protection is highly context specific. This underscores the importance of testing for treatment-effect heterogeneity along many dimensions of our covariate space, such as land, soil, and climate attributes; the degree of local economic pressures from agriculture and forestry; and economic, political, and institutional factors. Using our expansive data, we use the non-parametric causal random forest method of Wager and Athey (2018) to estimate highly-granular conditional average treatment effects (CATEs).

The results are sobering. First, the Europe-wide average treatment effect (ATE)—aggregated across countries, cohorts, and event time—is statistically and economically close to zero. We do not find evidence for a statistically precise, meaningful causal contribution of protection to our first measure—vegetation cover—in any of the five largest EU member states; a few, mostly small, countries show moderate treatment effects on either side of zero. Second, up to three decades after treatment, event-time treatment effects indicate a zero effect of protection on vegetation cover. Third, we find no trend in cohort-level ATEs. Land protected later in time does not contribute more to vegetation cover than land protected early in our sample. Fourth, these three findings are identical for our second measure—nightlights. Fifth, the treatment effects are not meaningfully heterogeneous across the covariate space such as initial greenness, population density, or measures of agricultural land productivity; the zero effect is stable and pervasive across a wide range of CATEs. Additionally, modest differences in treatment effects between countries do not correlate with political or land endowment variables and treatment effects are small for all levels of strictness. Sixth, we show that our conclusions hold up when using discrete land-use data. The limited species count data do not suggest a relationship between protection and biodiversity either.

These findings provide new insights into the limitations of area-based protection targets. Our model explains how the observed siting and stringency choices arise from the political economy of EU land protection. The EU delegates the implementation of its 30% area target to local authorities who might place less weight on ecological benefits than a social planner. Participatory procedures, whereby local constituents can object to costly sitings, together with green-glow benefits from designation, create incentives to protect the least contentious land first. Due to the delegation issue, current policy does not seem to weigh the full societal economic benefits and losses. Addressing that would require a substantial change in how the regulation is implemented, for example using cost-benefit analysis where costs and benefits beyond the local constituents are accounted for, and decision making at higher levels of government.

The absence of detectable ecological effects can be explained by two siting patterns. First, strict protection is predominantly applied to infra-marginal land where development pressure is low and protection has little bite. We find that 4% of protected areas are under strict protection and on land

facing low development pressure. Although such protection may safeguard against very long-run development threats, we find no evidence of such risks in areas that have been protected for several decades. Second, land facing substantial pressure can be designated a weak protection regime that imposes only limited constraints on economic activity. We find that 54% of protected areas are weakly protected and on higher-pressure land.

Finally, we note an important caveat regarding our outcome measure. Our primary indicator, NDVI, captures vegetation cover—an important dimension of biodiversity at scale—and is correlated with key animal species such as birds, but it is unlikely to reflect all aspects of biodiversity that land-protection policies intend to preserve.

The field of economics is increasingly interested in the decline of forests and biodiversity because of species extinction risk (Taylor and Weder 2024), ecoservice losses (Druckenmiller and Taylor 2022; Frank 2024), and financial risks (Kedward, Ryan-Collins, and Chenet 2023). Protected area policies are the primary policy instrument today to limit biodiversity declines and the EU’s flagship nature protection strategy. They are expected to affect 30% of global landmass by 2030. It is essential to understand the contribution of protected areas to mitigating vegetation and biodiversity loss, as well as their land-use impacts, which could impact economic performance and growth (Herrendorf, Rogerson, and Valentinyi 2014). While the EU’s policy has impressive scale, our results indicate the impact on vegetation cover and economic performance is likely to be minimal.

There are a handful of global studies on the effect of land protection that include Europe (Joppa and Pfaff 2009; Joppa and Pfaff 2010; Abman 2018; Maxwell et al. 2020; Wolf et al. 2021) and a large literature on protected-area policies in tropical forests (e.g., Andam et al. 2008; Sims 2010; Pfaff et al. 2015; Sims and Alix-Garcia 2017; Souza-Rodrigues 2019; Assunção et al. 2022; Keles, Pfaff, and Mascia 2022; Cheng, Sims, and Yi 2023; Rico-Straffon et al. 2023). This literature has generally found a modest impact on forest cover, but stronger positive effects are observed in well-enforced areas experiencing economic development pressure (Börner et al. 2020; Assunção, Gandour, and Rocha 2023; Reynaert, Souza-Rodrigues, and van Benthem 2023). Our study is unique in that our data span four decades at high frequency and has fine geographic resolution, allowing us to study the staggered introduction of land-protection policies using recent econometric advances that yield unbiased estimates in such settings and allow for unprecedented opportunities to study how treatment effects evolve over event time and cohorts. Prior studies of land protection—both global and national—have lacked one or more of these elements, typically covering relatively short time periods (about a decade or less), leveraging cross-sectional or limited panel variation through standard matching estimators, and/or lacking a causal research design.

A rapidly-growing literature studies the economics of conservation related to the political economy (Harstad and Mideksa 2017; Harstad 2023), trade policy (Hsiao 2021), and mechanism design of US conservation markets (Aronoff and Rafey 2023; Aspelund and Russo 2024). Our paper contributes to this literature by focusing on the environmental outcomes of a large-scale, regulation-based conservation program in Europe, highlighting that weak regulation and site selection in

areas with low opportunity costs are also issues in advanced economies, despite many EU countries ranking among the top for effective governance. Such evidence from advanced economies has been lacking despite their importance for meeting the global 30x30 target and their land-use and enforcement dynamics being different from tropical forest countries.³ The implementation of the EU's policy and its follow-up framework, the Nature Restoration Law, was heavily contested by member states and is expected to be at the center of political tensions over land use in the EU in the coming decade (Hancock and Jones 2024). Our theoretical and empirical results suggest that local political economy factors undermine the ambitious international targets ratified by the EU. Decentralized decision making by local jurisdictions whose conservation preferences do not align with the EU results in poor siting, weak protection, and negligible biodiversity improvements. Even so, the EU often conditions trade tariff reductions on nature protection efforts in other parts of the world without a critical assessment of its own contributions.

The rest of the paper proceeds as follows. Section 2 provides details about the EU's land-protection policies. In Section 3, we offer a conceptual framework and develop testable predictions. Section 4 describes the data sources and Section 5 describes four facts from the data. Section 6 outlines the empirical methodology and Section 7 presents results. Section 8 concludes.

2 Protected-area policy in the European Union

The EU specifies biodiversity strategies for each decade that translate the ratification of international agreements into specific EU goals.⁴ The 2010 and 2020 strategies progressively increased the scope of the EU's protected area policy, which will continue into the next decade. Europe's Green Deal and Nature Restoration Law aim to protect 30%, and rehabilitate at least 20%, of its land by 2030 (European Parliament 2023) in line with the ratification of COP 15.⁵ The member states adhere to these targets by assigning areas under the Natura 2000 policy (European Union 2009). This policy combines two earlier EU directives: the habitat directive (The Council of the European Communities 1992) and the birds directive (The Council of the European Communities 1997). The directives describe a list of species and habitats requiring conservation measures. The list is split into annexes based on the extent to which species and habitats are threatened. The directives list a set of restrictions for each annex, such as a restriction on land use to preserve the habitat of endangered species. The directive requests member states to take measures to maintain the animal population's size and habitat's territorial presence while considering economic requirements. Every six years, all EU member states report on the state of listed species and habitats.

The directive requires countries to submit a standardized report on their protected areas to the

3. The Auffhammer et al. (2021) and Bahrami, Gustafson, and Steiner (2024) studies of the effect of protection on land-market impacts in the US are rare exceptions.

4. See, for example, the communication on the 2020 strategy in European Commission (2011).

5. The Nature Restoration Law officially entered into force on August 18, 2024, marking a significant milestone in the EU's environmental policy. Many member states oppose the 2030 measure because of the agricultural restrictions embedded in the legislation.

European Commission, following International Union for Conservation of Nature (IUCN) guidelines. The Commission evaluates member state proposals and may amend them. This could lead to a back-and-forth between the Commission and member states over revisions to protection proposals. Member states then ratify their plans into national laws that specify the legal status of protected areas.

While the EU directives describe species and habitats needing protection, member states are responsible for translating these guidelines into actual policy. First and foremost, member states decide on the siting of protected areas. In line with the EU list of species and habitats, they identify territorial regions contributing to improving listed species and habitats. The member states then draw the boundaries of the areas in cooperation with ecologists, local communities, and farmers, and specify restrictions on different economic activities. Finally, they implement national laws and regulations and enforce the policy. The Commission oversees the member state plans and gives feedback but has no direct regulatory or enforcement power. This legal setting leads to potential variation in the policy's siting, restrictions, and enforcement across member states.

When studying the siting policies of member states more closely, it is striking that the procedure allows for a great deal of input from local communities and farmers. For example, the French procedure is detailed in Hassan Souheil and Douillet (2011) and is titled "Dialogue for Natura 2000." Under this plan, the French state contacts the prefect of a region if their experts in the Ministry of the Environment identify a region where protection should be implemented. The local prefect then establishes a steering committee ("comité de pilotage") to bring together all possible stakeholders to formulate the objectives of the protected areas and the siting. The report lists a dozen involved official structures (such as government bodies, hydrological agencies, and sports organizations) and two dozen stakeholders (ranging from hunters and farmers to tourists and scientists) as possible steering committee members. For each validation of a Natura 2000 area the collective procedure allows for participation and negotiation with local stakeholders. Participation is key; the document advises that "each stakeholder, each inhabitant, is legitimate to be involved closely or remotely in the Natura 2000 process simply because of their link to the area concerned. Leaving no one behind is a good way to avoid local stakeholders feeling a lack of consideration and thus to limit opposition" (Hassan Souheil and Douillet 2011). The regional steering committee validates a formal document (DOCOB) that is passed up to the French government, which in turn includes it in the national protected area plans submitted to the EU Commission.

Most member states' procedures incorporate local participation in the siting decisions. The Netherlands delegates the governance of protected areas to provinces in collaboration with local stakeholders. Each siting decision is publicly available and revised after a six-week period of public feedback.⁶ Similarly, the UK Department for the Environment, Food and Rural Affairs (Defra) oversees the UK's siting decisions after Natural England (Defra's executive body) identifies a possible area for siting. Natural England carries out a public consultation, to give everyone who

6. See <https://www.natura2000.nl/werkwijze/aanwijzing-natura-2000-gebieden-0>.

might be affected by the designation or who has relevant scientific information an opportunity to comment. This includes landowners and occupiers, local planning authorities, and other interested organizations. The results of the consultation are reported back to Defra, which may ask Natural England to try and resolve any remaining objections to the designation or provide more scientific information to support the proposal.⁷ Participation of local stakeholders seems to be the key unifying element of the procedure across EU member states.

Because the actual implementation is delegated to lower government bodies, the member states report many types of protected areas to the EU Commission. The policy thereby covers municipal, regional, and national protected areas. Some areas restrict all or most human activity (e.g., strict nature reserves and national parks; 7.6% of the EU's protected landmass), while others allow some industrial and agricultural activities (e.g., habitat or species management areas; 47% of the EU's protected landmass). We discuss the breakdown of these categories in Appendix A.3. Member states have different policies regarding protection and land ownership.⁸ Furthermore, the policy encompasses protected areas beyond merely aiming to improve biodiversity. Some are protected cultural heritage, and others serve recreational goals such as ecotourism.

Our analysis studies protected areas under the early EU directives or member state policies before 2009 and the complete EU Natura 2000 program after 2009. In 2020, the EU released an evaluation of the Natura 2000 network (European Environment Agency 2020). The report describes the difficulty of such an evaluation: "Measuring the ecological effectiveness of a network of protected areas is difficult, as baseline data are scarce and the data have many limitations, such as the lack of data enabling comparison of the conservation status of and trends in species and habitats inside and outside of the Natura 2000 network." The evaluation is based on expert opinion and member state surveys from the recurring six-year reporting requirement. However, none of the member states collect high-quality quantitative data that allow causal evaluation. Our study aims to provide large-scale, long-term causal evidence for the effectiveness of the EU's land-protection policy.

The EU's evaluation states that, between 2013 and 2018, the proportion of species listed in the birds directive with poor and bad status increased by 7%, corresponding to 81% of habitat assessments having a poor or bad conservation status, and only 15% of habitats having a good status. The main pressures that deteriorate the status of habitats are agricultural activities and urbanization. Agriculture is reported to affect habitats through changes in grassland management, landscape fragmentation, land-use conversion, and drainage. Member states reported more than 20,000 areas under severe pressure, confirming that the protection policy encompasses land where nature could potentially expand.

The report ends with actions that could improve the performance of the network. The first recommendation is to improve site selection: "Inefficient site selection has been linked to politically-motivated selection and giving low priority to conservation objectives compared with economic

7. See https://consult.defra.gov.uk/natural-england/crouch-roach-estuaries/supporting_documents/European%20leaflet%20Natura%202000.pdf.

8. Unfortunately, comprehensive ownership data across European countries are not available to us.

objectives.” Furthermore, the report states that management and monitoring could be more effective and that authorities should prioritize ecological performance. The poor evaluation report of the Natura 2000 network led the EU to formulate a more ambitious policy in the Nature Restoration Law of the EU Green Deal in 2024.

3 The policymaker’s objective function

We formalize a policymaker’s decision to protect land to state testable predictions. Consider a choice over protecting grids $i \in \mathbb{I}$, where $\mathbb{I} = \{1, 2, \dots, I\}$. Treated areas are protected at $t = g(i)$ so that g indicates the treatment cohort. Control areas are never protected. The decision maker selects sites i for protection, and we denote the treatment status of each plot with $D_{it} \in \{0, 1\}$. Protection is irreversible: $D_{it} = 1 \forall t \geq g$. Below, we generalize to a discrete-continuous decision: a decision maker selects the protection location as well as its stringency.

Land has economic use value v_{it} and ecological value e_{it} , including biodiversity, carbon capture, and other ecosystem services. Generating economic use value v_{it} requires human activity, denoted by land-use intensity a_{it} . Land use before protection, $a_{it}(0)$, is chosen by private actors to maximize their private payoffs. We assume that $v_{it}(a_{it})$ is continuous in land-use intensity and that economic use value is zero in the absence of human activity, $v_{it}(a_{it} = 0) = 0$. For each plot i , economic use value has a fixed sign for all relevant levels of land-use intensity: land use on a given plot either generates positive economic value or negative economic value, but not both.⁹ The absolute economic value increases with land-use intensity: $|v_{it}(a_{it})|$ rises with a_{it} . Ecological value equals $e_{it} = e_{it}(a_{it})$, where $e'_{it}(a) < 0$: more intensive land use reduces ecological value.

Our framework assumes that v_{it} , e_{it} , and a_{it} are constant over time except when treatment occurs, implying that their values before and after treatment only vary cross-sectionally: $a_{it < g} = a_i(0)$, $v_{it < g} = v_i(a_i(0))$, $e_{it < g} = e_i(a_i(0))$ and $a_{it \geq g} = a_i(1)$, $v_{it \geq g} = v_i(a_i(1))$, $e_{it \geq g} = e_i(a_i(1))$. The source of cross-sectional variation in v_{it} and e_{it} are differences in land’s economic and ecological productivity. Our empirical predictions below generalize to time-varying v_{it} and e_{it} , and the empirical analysis explicitly accounts for this by modeling rich time-varying treatment heterogeneity.

As long as a plot is not protected ($D_{it} = 0$), the resulting policy-maker-relevant economic use value $v_i(a_i(0))$ can be positive or negative depending on the plot. The unrestricted land use reduces ecological value below its maximum: $e_i(a_i(0)) \leq e_i(a_i = 0)$.

After protection ($D_{it} = 1$), the policy prohibits all land use for all $t \geq g(i)$, setting $a_{it \geq g}(1) = 0$ and $v_i(a_i(1)) = 0$. The ecological outcome of plot i attains its maximum value, $e_i(a_i(1)) = e_i(a_i = 0)$, after protection.¹⁰ This same ecological value would arise if the land had been idle even in the absence of protection (i.e., if $a_i(0) = 0$).

9. An example of inefficient land use is activity that generates private profits but yields negative economic use value from the policymaker’s perspective once the costs of road construction, maintenance, and other public provisions are accounted for.

10. We consider e to include any gains in tourism value from protection.

3.1 Planner solution

In each period t , a unit of land generates payoffs to a planner. The planner's present value (at time $t = s$) of a non-protected unit of land, $V_{it}(0)$, and of a protected unit, $V_{it}(1)$, equal:

$$V_{it}(0) = \sum_{t=s}^{\infty} \delta^{t-s} [v_i(a_i(0)) + e_i(a_i(0))] \quad \text{and} \quad V_{it}(1) = \sum_{t=s}^{\infty} \delta^{t-s} [v_i(a_i(1)) + e_i(a_i(1)) - c(1)] \quad (1)$$

where δ equals the discount factor and $v_i(a_i(1)) = 0$. Protection comes with an administrative, monitoring, and enforcement cost of $c(1)$.¹¹ The difference in the present value (at time $t = g$) of a plot i that is protected at $t = g$ relative to the no-protection counterfactual equals:

$$V_{ig}(1) - V_{ig}(0) = \sum_{t=g}^{\infty} \delta^{t-g} [v_i(a_i(1)) - v_i(a_i(0)) + e_i(a_i(1)) - e_i(a_i(0)) - c(1)], \quad (2)$$

which shows how land protection is a trade-off between changes from restricted economic activity, ecological gains, and implementation costs. We define the ecological treatment effect of protection as $\Delta e_i = e_i(a_i(1)) - e_i(a_i(0))$, which will function as our primary empirical object of interest.¹²

Objective function: In the data, we witness gradual protection over several decades. Therefore, we assume the planner has a resource constraint that limits the maximum amount of plots that can be treated in each period to equal \bar{I} . The planner faces the dynamic discrete choice problem of where and when to apply protection:

$$\max_{\{D_{it}\}_{i=1}^I} \sum_0^{\infty} \delta^t \sum_{i=1}^I [v_i(a_i(D_{it})) + e_i(a_i(D_{it})) - c(D_{it})] \quad (3)$$

$$s.t. \# \{ i \mid D_{it} = 1, D_{it-1} = 0 \} \leq \bar{I} \quad \text{for all } t \geq 0. \quad (4)$$

Here, (4) is the constraint capturing that the decision maker cannot protect all the land at the same time—no more than \bar{I} of the D_{it} can switch from zero to one for each treatment cohort g .

Characterizing the solution: The planner finds the optimal allocation according to the following algorithm:

Step 0: Each period t starts with the remaining unprotected land $S_t \subseteq \mathbb{I}$.

Step 1: The decision maker computes each plot's net present value change from protection defined as:

$$\phi_i = \frac{\Delta e_i - v_i(a_i(0)) - c}{1 - \delta}$$

11. We present a model with homogeneous implementation costs, but extensions to heterogeneous implementation costs could be interesting in settings where researchers can access data on such costs.

12. Note that the empirical analysis allows for time-varying Δe_{it} as we account for temporal changes in e_{it} and v_{it} before and after treatment.

Step 2: Define the set S^* with the \bar{I} most positive ϕ_i values in the set S_t , breaking ties arbitrarily:

$$S^* = \{i_1, i_2, \dots, i_{\bar{I}}\} \text{ such that } \phi_{i_1} \geq \phi_{i_2} \geq \dots \geq \phi_{i_{\bar{I}}} \geq \phi_j \text{ for all } j \notin S^*.$$

Step 3: For each $i \in S^*$, set $D_{it} = 1$. For all other units $j \notin S^*$, keep $D_{jt} = 0$.

Step 4: The stock of untreated land updates: $S_{t+1} = S_t \setminus S^*$.

Repeat Step 0-4 until no land is left with $\phi_i > 0$.

We now use this characterization of an optimal protection policy to formulate testable predictions to connect the theoretical framework of optimal siting to the empirical framework that estimates the treatment effects of the EU's protected area policy.

Prediction 1: positive ecological effects for optimally protected land

A planner protects plots with $\phi_i > 0$, which implies that $\Delta e_i - v_i(a_i(0)) > c$. Under strict protection, any plot with positive baseline activity $a_i(0) > 0$ necessarily exhibits a strictly positive treatment effect, $\Delta e_i > 0$.

For plots with negative economic use value, $v_i(a_i(0)) < 0$, protection reduces a loss-making use and raises ecological value, implying positive ecological treatment effects; but it is beneficial only if these gains are sufficient to offset the implementation and enforcement cost c . For plots with positive economic use value, $v_i(a_i(0)) > 0$, protection involves a trade-off between foregone economic use and improved ecological outcomes, and protection occurs only when the ecological gain Δe_i is sufficiently large to compensate for both lost economic value and implementation costs.

Furthermore, the planner would not engage in the protection of land that is not under any economic pressure—i.e., plots with no baseline activity ($a_i(0) = 0$) for which $v_i(a_i(0)) = v_i(a_i(1)) = 0$ and hence $\Delta e_i = 0$ —as long as $c(1) > 0$; such efforts would be considered a wasteful allocation of resources.

Our framework assumes v_i constant over time except when treatment occurs. However, our prediction generalizes to the settings where v_{it} is time-varying beyond treatment. For example, land at risk of future economic development ($a_{it}(0)$ increases over time) may warrant protection as soon as the net present value of ecological savings outweighs the costs of preventing development. Even in areas where economic activity has previously declined and nature has regrown, a protection policy that further restricts a_{it} could have meaningful effects by choosing a path of $v_{it}(a_{it}(1))$ that declines faster than $v_{it}(a_{it}(0))$, accelerating natural regrowth beyond what we would see without protection. This is particularly relevant in Europe, where we observe reforestation in some areas but increased pressure in others. A protection policy can accelerate greening and reforestation in certain areas and prevent the loss of vegetation in other areas that come under economic pressure.

We might also expect that e_i adjusts gradually after treatment, resulting in heterogeneity in the treatment effect over time. Nature may require time to regenerate after protection, or, conversely,

untreated (control) plots might experience gradual degradation. These dynamics will manifest as event-time heterogeneity in our treatment effects.

Prediction 2: heterogeneity in treatment effects across protection cohorts

The data allow us to examine heterogeneity in treatment effects across cohorts. Under the planner's solution, plots are protected in decreasing order of the net gain from protection, ϕ_i . Plots with the highest values of ϕ_i are protected first, while plots with lower net gains are protected in later cohorts. The evolution of average ecological effects across cohorts depends on how ecological gains Δe_i covary with baseline economic use value $v_i(a_i(0))$.

To illustrate how cohort-specific treatment effects may rise or fall under the planner's solution, suppose that baseline economic use value is positive on all plots, $v_i(a_i(0)) > 0$. Assume further that baseline land-use intensity $a_i(0)$ is monotonically increasing in baseline economic use value $v_i(a_i(0))$, and that ecological value $e_i(a)$ is monotonically decreasing in a both within and across plots. Under these assumptions, ecological gains from protection, Δe_i , are increasing when plots are ranked by their baseline economic use value $v_i(a_i(0))$. We can then characterize the planner's protection path by evaluating:

$$\frac{\partial \phi_i}{\partial v_i(a_i(0))} = \frac{1}{1-\delta} \left(\frac{\partial \Delta e_i}{\partial v_i(a_i(0))} - 1 \right).$$

The derivative captures how the net present value of protection changes when we vary lost economic production, allowing us to see if ecological gains or economic losses dominate the order of protection and the associated treatment effects across cohorts.

If $\partial \Delta e_i / \partial v_i(a_i(0)) > 1$, ecological gains rise more than one-for-one with baseline economic value and the planner protects high- v plots earlier, so average ecological treatment effects tend to be larger in early cohorts. If instead $\partial \Delta e_i / \partial v_i(a_i(0)) < 1$, foregone economic use dominates and the planner protects low- v plots earlier; in that case, average ecological treatment effects may increase across cohorts as protection expands to more valuable (and more ecologically responsive) land.

The cross-sectional monotonicity assumption is strong, and in general, we cannot assign a strict order for the treatment effect size across cohorts. The optimal path of protection depends on the correlation between v_{it} and e_{it} across i .¹³

3.2 Siting under delegation

As described in Section 2, the EU establishes quantity targets based on international agreements; however, the actual siting decisions are largely delegated to national or local authorities. We discuss the potential impact of such delegation and quantity targets on treatment effects.

13. See Weitzman (1998) and Metrick and Weitzman (1998) and the subsequent literature for theoretical models of optimal siting.

We define the problem of the local decision maker as follows:

$$\max_{\{D_{it}\}_{i=1}^I} \sum_0^{\infty} \delta^t \sum_{i=1}^I \left[v_i(a_i(D_{it})) + (1 - \alpha_1) e_i(a_i(D_{it})) - c(D_{it}) + \alpha_2(D_{it}) \right] \quad (5)$$

$$s.t. \# \{ i \mid D_{it} = 1, D_{it-1} = 0 \} \leq \bar{I} \quad \text{for all } t \geq 0 \quad (6)$$

$$s.t. \# \{ i \mid D_{it} = 1 \} \geq \tilde{I} \quad \text{before } t > t^{target} \quad (7)$$

with $0 \leq \alpha_1 \leq 1$, $\alpha_2 \geq 0$ when $D_{it} = 1$, and t^{target} the date by which the EU requires \tilde{I} to be protected. There are two differences between the local authority's problem and the planner's problem. First, there might be preference misalignment. When $\alpha_1 > 0$, local authorities place less weight on ecological benefits, while $\alpha_2 > 0$ implies a green-glow effect capturing political benefits from implementing an environmentally friendly policy (even if its impact might be limited). Second, the local decision maker might face quantity targets. Under international agreements such as the 30-by-30 target, the optimization problem changes because the EU delegates the protection of vast amounts of land to local authorities. The quantity agreement introduces an additional constraint (7): the number of protected areas must reach the level \tilde{I} by $t = t^{target}$.

Absent a binding quantity target, a higher degree of preference misalignment (i.e., as α_1 approaches one) leads the local policymaker to protect less land than the planner. As α_1 increases, the net present value of protection for the local authority places less weight on ecological gains, so that the local valuation of protection is strictly lower than the planner's ϕ_i for any given plot. The local decision maker still protects land that has negative $v_i(a_i(0))$ or high enough ecological value, and we expect positive treatment effects whenever land is protected, but the ranking of plots ϕ_i will be increasingly driven by the economic loss $v_i(0)$ rather than the ecological gain. When $\alpha_1 = 1$, the local policymaker will only protect land with negative $v_i(a_i(0))$ (as long as it warrants spending c).

There could be two reasons why $\alpha_1 > 0$. First, local politicians may not internalize biodiversity externalities as much as EU-level politicians because the benefits of biodiversity are not necessarily local, whereas the economic losses of protection are—as a result, they may want to “un-do” the EU policy to some extent. Second, local stakeholder participation might favor the representation of those affected by the losses from protection. As described in Section 2, stakeholder participation seems to be the norm in siting policies. While NGOs and environmental organizations usually participate in siting discussions, local economic stakeholders often succeed in exerting pivotal influence on siting decisions.

In the case of green-glow benefits, local decision makers may protect infra-marginal plots that are not at risk of economic development and, therefore, without ecological gains ($v_i(a_i(0)) = 0, \Delta e_i = 0$), as they now gain α_2 with each protection that otherwise leaves the land unchanged. However, whenever land is available with $v_i(a_i(0)) \leq (1 - \alpha_1) \Delta e_i$, a green-glow decision maker will protect this land before land without any land use, regardless of how large α_2 is (this is also trivially satisfied for land with negative $v_i(a_i(0))$). The net gain from protection is larger than for infra-marginal

land and such protection has positive treatment effects. Hence, we would witness zero effects after the protection of land with higher ϕ_i . Once the local policymaker runs out of infra-marginal land to protect, we could even see protection where economic costs outweigh ecological benefits, which would be associated with positive treatment effects again.

Under the additional constraint of a quantity target, a decision maker may protect a number of plots even if their marginal net benefit of protection is negative. In this case, once the optimal set of plots (with $\phi_i > 0$) is exhausted, additional protection is applied merely to meet the quantity target, and the local policymaker selects the least costly plots to satisfy the quantity constraint. Initially, plots with $0 \leq \Delta e_i - v_i(a_i(0)) \leq c$ will be protected and treatment effects will be positive. Next, protection happens on land for which $\Delta e_i = v_i(a_i(0))$. This includes infra-marginal plots where $a_i(0) = 0$; their net cost of protection is c . Whenever the decision maker runs out of cheap, infra-marginal land to protect, as could be the case with ambitious quantity targets, treatment effects may gradually become larger again as the decision maker must start protecting plots with $v_i(a_i(0)) > \Delta e_i > 0$.

So far, all the scenarios we discussed—preference misalignment, green glow, and quantity targets—lead to a protection policy where we initially expect the treatment effects to be positive as the local policymaker selects the land with the most positive ϕ_i first, but eventually could lead the local authority to select infra-marginal land for protection.

Prediction 3: null effects from strict protection on land with low land-use intensity

When green-glow benefits ($\alpha_2 > 0$) or area-based targets are combined with limited internalization of ecological gains (α_1 close to one), protection decisions on land with zero or very low baseline activity can become optimal from the local policymaker’s perspective, even in early cohorts.¹⁴ In this case, protection decisions are driven primarily by the green-glow benefit or by the need to satisfy the quantity constraint, rather than by realized ecological gains—the local authority places low value on these yet faces foregone economic use costs $v(a)$ that increase with the reduction in human activity. Optimal implementation then targets plots with very low baseline land-use intensity $a_i(0) \approx 0$, for which strict protection implies little change in activity and foregone economic use value is low. Such protection induces only negligible changes in ecological value, $\Delta e_i \approx 0$, despite formally binding restrictions on land use. Moreover, if the administrative and enforcement costs of protection are increasing in the baseline land-use intensity—so that it is particularly costly to limit activity on highly used plots—this selection toward low-intensity land is further reinforced, leading to an even larger share of protected plots with null ecological effects.

In contrast to the planner benchmark—under which strict protection is targeted toward plots with substantial baseline activity and therefore yields strictly positive ecological gains—a local policymaker who places little weight on ecological outcomes, and who is motivated by green-glow benefits or quantity targets, may instead direct strict protection toward low-intensity land, generating widespread null ecological effects.

14. Note that the policymaker would protect areas with $(v < 0, a > 0, \Delta e_i > 0)$ before protecting land with $a = 0$.

3.3 Siting and strictness under delegation

Until now, we have modeled protection as a discrete decision, but in practice the strictness of protection is also a policy choice. For each protected plot, the local policymaker can restrict only certain economic activities rather than prohibit all use. We therefore introduce a continuous protection stringency parameter $d_i \in [0, 1]$, which is constant over time after treatment. Protection affects land-use intensity according to $a_i(1) = (1 - d_i) a_i(0)$, where $d_i = 1$ corresponds to strict protection with no remaining human activity, and values of d_i close to zero correspond to weak protection imposing few economic restrictions. Such low-stringency protection may be particularly relevant for certain categories of EU conservation areas, such as species management areas (see Appendix A).

This discrete-continuous optimization problem provides further flexibility to local policymakers in implementing a protection policy, as the protection level d_i can be chosen to optimize the protection payoff ϕ_i for each plot and thereby increases the number of plots with a positive net present protection payoff ϕ_i .

Prediction 4: null effects from weak protection on land with high land-use intensity

When protection stringency is a choice variable, a local policymaker who places little weight on ecological outcomes (α_1 close to one), but derives benefits from designation itself ($\alpha_2 > 0$) or faces binding quantity targets, may optimally implement weak protection on many plots. In addition to protecting infra-marginal land with low baseline activity, the policy can now apply $d_{it} \approx 0$ to prevent steep protection cost. Hence, the set of land to which inconsequential cheap protection can be applied expands.

Under variable stringency, post-treatment land-use intensity equals $a_i(1) = (1 - d_i)a_i(0)$. Ecological gains therefore depend on the reduction in activity induced by protection:

$$\Delta e_i(d_i) = e_i((1 - d_i)a_i(0)) - e_i(a_i(0)).$$

For low stringency levels ($d_i \approx 0$), the reduction in land-use intensity is small even on plots with high baseline activity, implying $\Delta e_i(d_i) \approx 0$. As a result, variable protection stringency expands the set of circumstances under which protection generates null or near-null ecological effects, including on land with substantial baseline use. Compared to strict protection, weak protection thus provides an additional channel through which delegated implementation can yield widespread ecological inframarginality.

Overall, our framework highlights the conditions under which land-protection policies may generate null ecological treatment effects. Such zero or near-zero effects arise when protection decisions are driven by preference misalignment—specifically, when local policymakers place little weight on ecological outcomes (high α_1)—together with incentives to designate land for political or administrative reasons, such as green-glow benefits or binding area-based targets. In this case, they optimally select land with low baseline use or apply weak protection stringency in order to

limit economic and enforcement costs, resulting in little or no change in ecological outcomes.

Although our empirical focus is on Europe, the mechanisms highlighted by the model are likely to be even more salient in low-income, forest-rich countries, where local authorities often face limited enforcement capacity and severe budget constraints. In such settings, local decision makers may strongly under-value broader societal ecological benefits (high α_1) and, when faced with area-based protection targets, may rationally implement land-protection policies that have little or no ecological impact.

4 Data

We collect six types of data to assemble two remote-sensing datasets spanning the entirety of the European Union between 1985-2019.¹⁵ The most granular dataset (to analyze vegetation cover) divides Europe into 117 million equal-sized grids of 300 by 300 meters. The second dataset (to analyze nightlights) has grids of one square kilometer. Our data deliver comparable and consistent measures across space and time (see Appendix A for details):

Policy rollout. For every protected area in the EU, we assign the date of its initial protection based on data from the Common Database on Designated Areas (CDDA), consolidating land-protection policies across 39 European countries. Focusing solely on terrestrial protection and excluding marine reserves, our dataset includes details on 118,511 distinct areas that were protected between 1800 and 2019.^{16,17} We establish a grid cell as protected if any non-zero fraction of its land area falls under a conservation agreement.

Vegetation cover. We aggregate satellite images from the Landsat 5, 7, and 8 data to construct a continuous normalized difference vegetation index (NDVI) at a bi-annual frequency, with higher values on a 0-100 scale indicating denser and richer vegetation. We rescale NDVI indices to be between 0-100 instead of the [-1,1] range standard in the remote-sensing literature. We focus on the [0-1] range and scale the index to 0-100. We drop observations with NDVI less than 0, as this range corresponds to snow, water, and clouds.¹⁸ The remote-sensing measures start in 1985. We use biennial aggregation to reduce missing data problems caused by cloud coverage and focus on the summer months when perennial vegetation is most visible. We use NDVI because we can construct the measure with early Landsat data, allowing us to obtain a panel that is four decades

15. We use the term EU, but the data include the 27 member states as well as Albania, Bosnia, Montenegro, Macedonia, Norway, Switzerland, and Serbia. We exclude Iceland, Malta, and Liechtenstein due to missing data issues.

16. France has missing foundation dates for 1,447 areas, representing only 0.3% of the protected area surface in France. We add the foundation years for the 127 largest French areas with missing foundation dates from internet sources.

17. See Villasenor-Derbez, Costello, and Plantinga (2024) and McDonald et al. (2024) for an analysis of the impacts of marine protected areas under the 30x30 target.

18. We also present robustness to samples with NDVI above 40 to focus on land with vegetation and exclude urban grids, bare soil, and rocky landscapes. See Appendix D and the discussion in Section 7 for details and evidence that our results are robust to the chosen NDVI threshold.

long. Alternative measures such as Vegetation Cover Fields (VCF) are only available since 2000.¹⁹

Previous research has shown that NDVI correlates equally well as other vegetation measures with the frequently-used biodiversity marker bird species richness (Nieto, Flombaum, and Garbulsky 2015; Hobi et al. 2017). The finding that NDVI is an excellent predictor for bird diversity is replicated in our study area for French (Bonthoux et al. 2018) and Mediterranean landscapes (Ribeiro et al. 2019). NDVI has shown to be effective in measuring the treatment effect of policies targeting land use across different sensors (Lassiter 2022). Nevertheless, it is important to mention that the EU's protection policies have multiple biodiversity objectives which are not perfectly captured by NDVI. NDVI will not be able to detect certain types of biodiversity improvements, and will be at best positively correlated with others. For example, species management areas may lead to improved animals counts, including insects, but we lack evidence on how vegetation cover correlates with many animal species other than birds. However, as we discuss below, direct species count data availability is limited.

Furthermore, we think it is helpful to focus on continuous measures. First, it allows us to capture gradual changes in vegetation that do not necessarily lead to land use re-categorization. For example, new construction such as housing or large-scale solar panel installation, or agricultural expansion, will result in decreases in a plot's NDVI score even when the discrete land-use category might be unaffected. Using discrete land-use data from the Historical Land Dynamics Assessment Program (HILDA), we verify that NDVI, on average, is highest for forested land (62, and more than 25% of forested grids have NDVI above 73), 53 for grassland, and 49 for cropland. Additionally, the continuous measure avoids classification errors that plague categorical land-use classifiers (see Alix-Garcia and Millimet 2022; Torchiana et al. 2023). Continuous measures are critical because our data encompass various protected areas. In 93% of the protected areas some economic exploitation of the land is allowed for and we might expect some of these areas to have gradual greening rather than discrete land-use changes.

Nightlights. We rely on Li et al. (2020) for a 1992-2018 one square kilometer panel of remotely-sensed nightlights. Here, our goal is to measure human presence on a granular scale. If protected areas limit economic activity, we expect outward migration from the area and reduced traffic, which could reduce nightlights. Nightlights have been used as a proxy for economic development and GDP in remote areas (Donaldson and Storeygard 2016) and urban/settled areas alike (Gibson et al. 2021), but this approach has also received criticism about unstable relationships and data inconsistencies (Chen and Nordhaus 2011; Bickenbach et al. 2016). We do not aim to interpret nightlights as a GDP/economic development measure. Still, we think it is useful as a measure complementary to NDVI, capturing the degree of human presence in the area.

HILDA. We obtain discrete land-use data via HILDA dating back to 1900 at a decadal frequency and a resolution of one square kilometer (Fuchs et al. 2015). HILDA classifies each grid as settlement, cropland, forest, grassland, other land, or water. The data are constructed by harmo-

19. The Enhanced Vegetation Index (EVI) targets the measurement of tropical forests not present in the EU.

nizing historical land cover information such as national inventories, maps, and aerial photographs with remote-sensing data. HILDA allows us to investigate long-term trends in EU land use. We are mainly interested in forests, grassland (which includes pastures), and cropland. Appendix Figure A.4 shows land-use shares between 1900 and 2010, and Appendix Table A.5 reports the EU's land-use transition matrix for that period.

Species counts. We use the BioTIME dataset, which is the largest available aggregation of species count studies across space and time, see also Liang, Rudik, and Zou (2023). Species count data allow for more direct measures of biodiversity. However, many count data suffer from one or all of the following issues: nonrandom location of counts, short panels with recent coverage only, a low number of species, model-based projections across space instead of raw data, and limited regional coverage. Each of these issues is problematic for the goal of our study: a comprehensive, long-term, EU-wide evaluation of protection. None of the species count data sources, including BioTIME, allow us to achieve a similar causal research design as with remote-sensing data.²⁰

Control variables. Finally, we add data on bio-geographical regions from the European Environment Agency; climate zones, soil properties, and topography from the European Soil Data Centre; precipitation from the European Centre for Medium-Range Weather Forecasts; and solar radiation from WorldClim. We merge these controls with the NDVI and nightlight data to control for each grid's natural vegetation growth propensity.

5 Descriptive evidence from 100 years of land-use data

In this section, we provide four empirical facts regarding land use in the European Union.

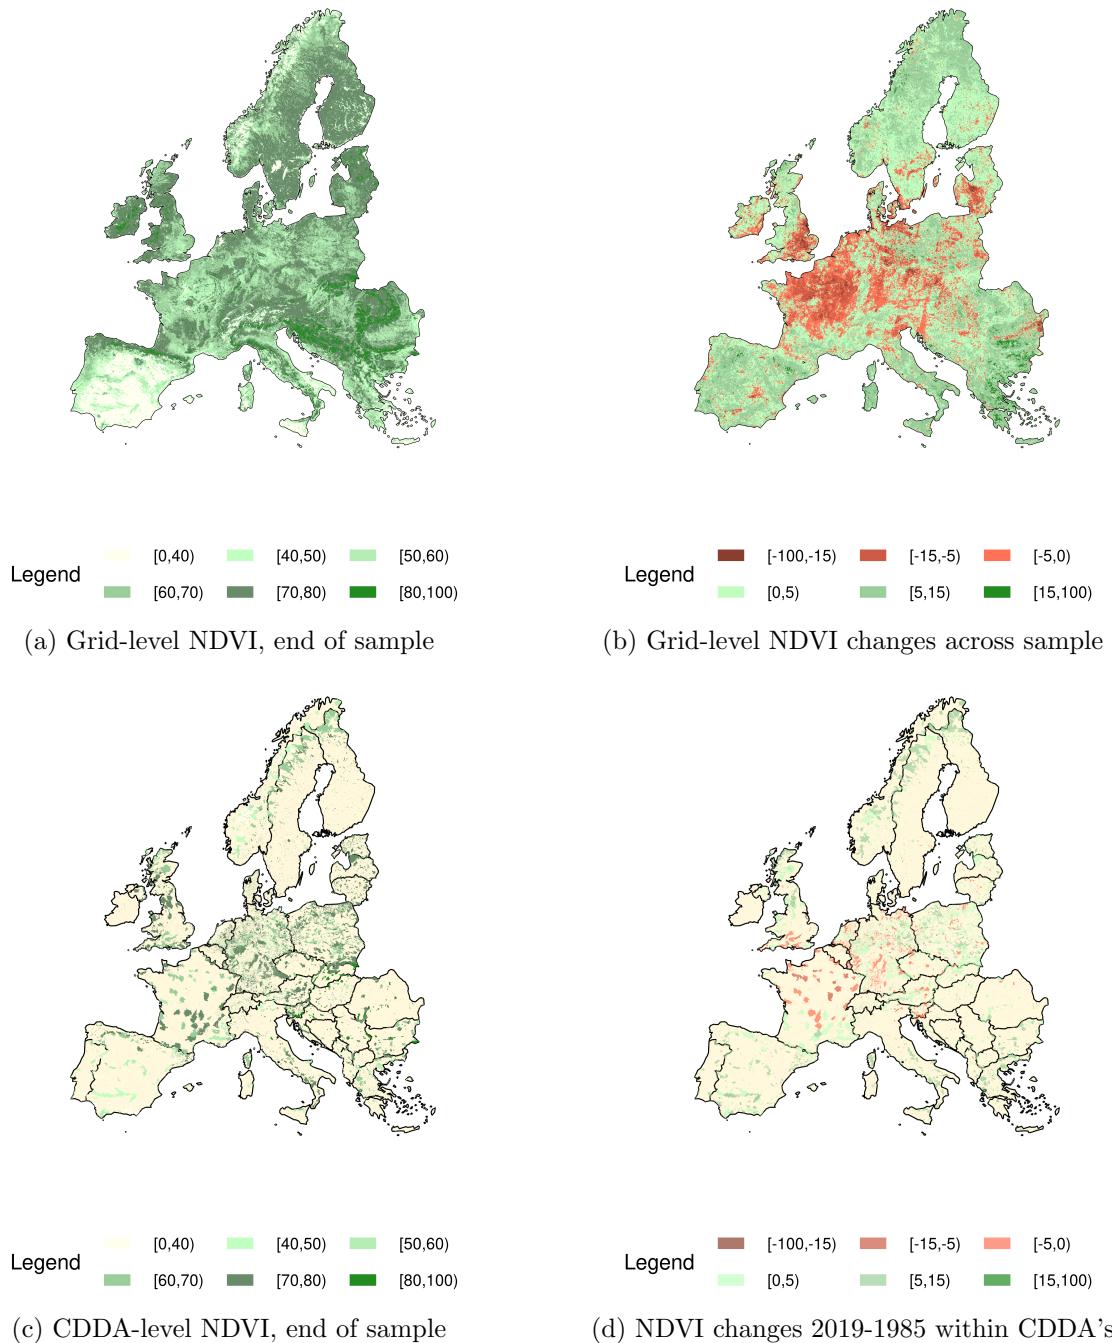
Fact 1. Large parts of Europe have experienced land-use transition in the last 40 years. Figure 1 maps the levels and changes in NDVI in Europe. Panel (a) shows that NDVI indices range between 50 and 100 in many areas. Panel (b) plots changes in the NDVI index between 1985 and 2019. We document decreasing NDVI in red and increasing NDVI in green. The map reveals a substantial greening in certain parts of Europe, but vegetation degradation in others. The greening is most pronounced in southeastern Europe and Scandinavia. NDVI decreases in large parts of France, Germany, and the Baltic states. This land seems to be under pressure from economic development, and stringent protection may prevent vegetation loss. Many parts of the EU have seen small changes in NDVI, not larger than 5 points in absolute value.

Panels (c) and (d) replicate panels (a) and (b) but only for areas that had received protection by 2019.²¹ Many protected areas have high greenness levels in 2019. Most of them experienced

20. Promising improvements in more direct measures of biodiversity include several databases of animal tracking data, such as the Global Biodiversity Information Facility, Movebank, the PanEuropean Common Bird Monitoring Scheme, eBird, eButterfly, and the European Bird Census Council. See <https://www.gbif.org/>, <https://www.movebank.org/cms/movebank-main>, <https://pecbms.info/>, <https://ebird.org/home>, <https://www.e-butterfly.org/> and <https://www.ebcc.info/>.

21. As visible in Panel (c), Ireland appears to have only a few, small protected areas in our data. This is an artifact of the shapefile of the European CDDA database: Ireland has several large Special Areas of Conservation according

Figure 1: NDVI changes inside and outside protected areas



NOTES: Panel (a) plots the NDVI across 2015-2019 for each grid. Panel (b) plots the change in NDVI between 2015-2019 and 1985-1989. Panel (c) plots NDVI within each CDDA for 2015-2019. Panel (d) plots the change in NDVI within each CDDA between 2015-2019 and 1985-1989.

increases in NDVI between 1985-2019, although there are also many protected areas with small positive or negative NDVI changes. However, comparing (b) and (d), many are also situated in areas that are greening regardless of land protection. The changes in greenness over time in protected areas vs. non-protected areas that are similar on observables, together with the staggered implementation of protection, form the primary identifying variation we exploit in this paper.

Fact 2. Economically valuable land continues to be deforested. Table 1 demonstrates that economically valuable land has experienced deforestation in the EU. Compared to land that remained forested, deforested land is more populated, more urbanized (higher nightlights), and more agriculturally suitable (longer growing seasons and potassium-rich soil).²² The table also highlights that the EU has protected land at risk of deforestation at a very similar rate as intact forests. The balance table thus suggests that land protection does occur on land at risk of development.

Table 1: Average of key variables among land classified as forest in 1900 by whether that land was deforested, 1900-2010

	Not deforested, 1900-2010		Deforested, 1900-2010	
	Mean	Std. dev.	Mean	Std. dev.
Nightlights in 2010	6.3	9.3	10.8	12.1
Percent of grid protected	24.2	45.9	28.7	46.0
Population density in 2000	46.7	143.4	103.8	273.8
Crop suitability	6.2	1.3	6.3	1.3
Forest suitability	3.5	1.5	4.0	1.4
Grassland suitability	5.6	1.4	5.6	1.6
Slope steepness	2.3	3.1	2.0	3.0
Solar radiance	10.7	2.2	11.4	2.0
Precipitation	767.3	266.9	791.2	261.4
Potassium	147.1	55.9	175.9	71.5
Nitrogen	2.2	1.0	2.3	1.0
Growing season length	222.5	67.0	256.0	57.1

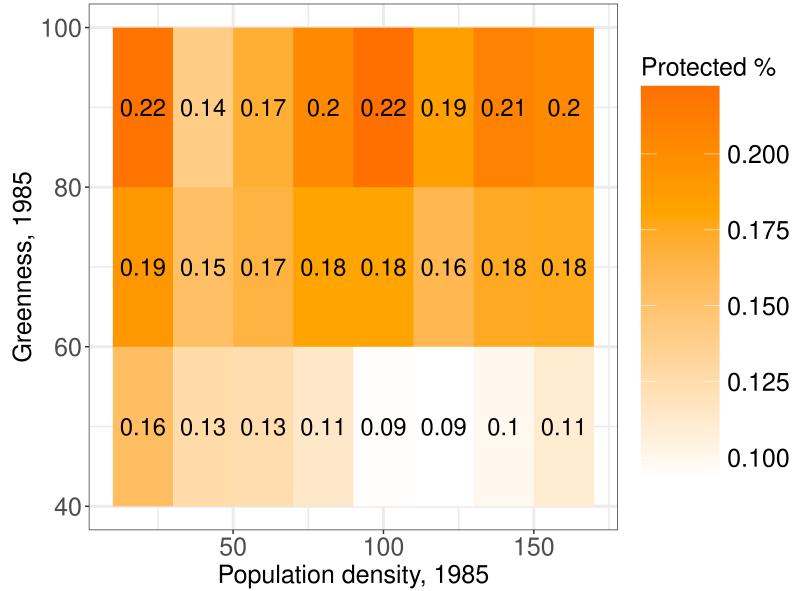
NOTES: Table presents balance of several key variables against an indicator based on the discrete land-use classifications provided by the HILDA data. Not deforested indicates 1 km × 1 km grid cells which were coded as forest in 1900 and were still coded as forest in 2010. Deforested indicates areas which were forest in 1900 but were coded as any other category in 2010. Percent of grid protected indicates the percent of the 1 square kilometer grid which contains protected areas, regardless of their designation year and CDDA designation. Means of time-varying variables are calculated in a specific cross-section, as indicated in the table. Units for all variables are indicated in Appendix Table A.4. Appendix Table A.6 replicates the table for the years 1990-2010 and 2000-2010.

Fact 3. Land-protection targets greener areas and extends across both sparsely settled and densely populated regions. We explore protected area siting visually in Figure to the National Park and Wildlife Service (see <https://www.npws.ie/maps-and-data>) but these do not appear in the CDDA.

22. We re-create the table for recent decades in Appendix Table A.6. Land deforested more recently tends to be located in more populated areas, likely reflecting land pressure from urbanization, especially in the 1990s.

2. The figure illustrates the probability of land protection since 1985, categorized by the initial greenness and population density of each land area, as of 1985. Then, for example, among all land that had a population density of less than 20 people per square kilometer and greenness above 80 in 1985, 22% by area received protected status after 1985. The figure shows that protection is more targeted towards land with higher greenness, but is targeted equally across population bins. Densely and sparsely populated land receive similar rates of protection.²³

Figure 2: Share of protected areas in population density and starting greenness bins



NOTES: Population density and greenness are divided into increments of 20, from [0, 160] and [40, 100], respectively. Resolution of a grid cell is 300 × 300 meters. Grid cells are considered protected if any non-zero share of the grid cell protected. Greenness-by-population density cells with fewer than 100 grid cells are omitted. Sample includes all land which had a greenness of at least 40 at some point in the sample.

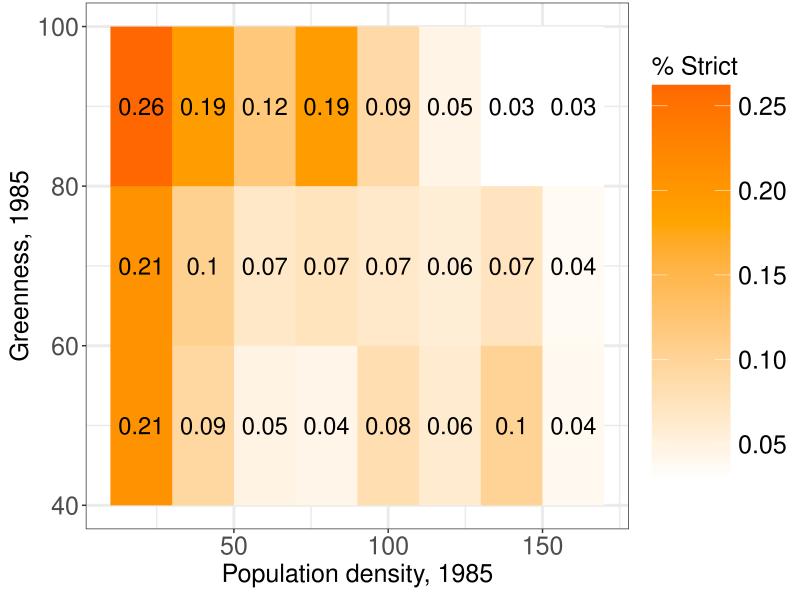
Fact 4. Protection is strictest on sparsely populated and initially green land, while the majority of protected areas are subject to less strict protection. We assess if there are spatial patterns in the stringency of protection. Figure 3 describes the likelihood that land protection is in IUCN Categories Ia, Ib, or II by greenness and population density. These three categories impose strict restrictions on human interaction: only scientific, ecological, and recreational use are permitted (though in some cases, an exception is made for use by indigenous peoples), while other categories allow for some economic use of the area.²⁴ Then, for example, among protected areas founded after 1985 which had 1985 population density less than 20 and

23. Appendix Table A.14 presents the results from a linear probability model of land protection. We find that more economically valuable land, as measured by nightlight luminosity, the presence of cropland, or settled land use before protection, has a lower probability of protection.

24. The criteria for these three categories explicitly reference “pristinity” of nature; others do not. For more discussion of IUCN categories, see Appendix Table A.3. Adding areas of Type III (natural monuments) to the “strict” designation does not alter the overall pattern of results.

greenness above 80, 26% are strict. The figure reveals a stark pattern: post-1985 strict land protection occurs on low-density or already-green land as of 1985.

Figure 3: Share of protected areas by starting greenness and population density which are strict



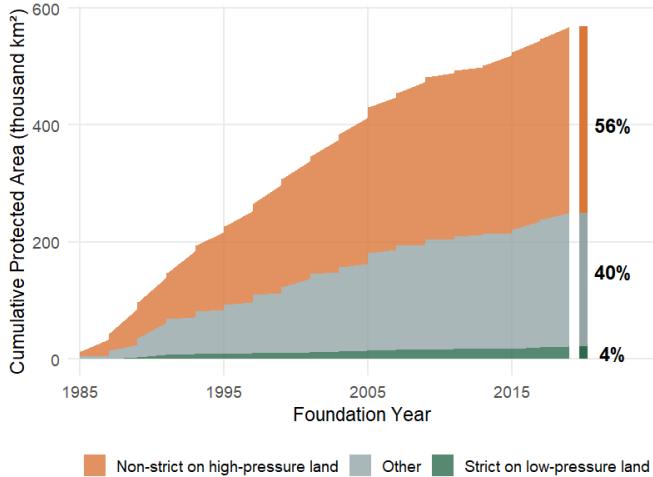
NOTES: Figure is constructed exactly following Figure 2. Entries are populated with the area-weighted average strictness of protected areas in each cell. Strictness is calculated using the International Union for the Conservation of Nature (IUCN) classification, with Ia, Ib, and II areas mapping to 1, and III or below mapping to 0. Ia, Ib, and II areas place the strictest limitations on human activity (only scientific, ecological, and recreational use, respectively). Thus, stricter areas have high (dark orange) values of the strictness index.

These stark spatial patterns reveal that strict protection is less likely to be applied when nature appears under pressure. As the European Environment Agency (2020) points out, urbanization and agriculture are the main causes of nature degradation in the EU. Urbanization correlates by definition with population density, and almost all agricultural areas have NDVI values well below 80. Therefore, Figure 3 reveals that protection is likely to be less strict where we expect pressure on nature. Such areas are far more likely to be species management areas or other types of protection that do not substantially restrict human activity. Yet, we acknowledge that population density and initial greenness are not perfect proxies for pressure, and biodiversity could also be under pressure in ways we do not measure in our analysis.

Figure 4 illustrates how protection stringency varies with development pressure, contrasting strict protection on low-pressure land with weaker protection on land facing higher pressure. By 2019, 56% of the EU's protected areas were located on land with relatively high population density and middling or low initial greenness, but these areas are subject to weaker protection mandates (IUCN Categories Ia, Ib, and IUCN Category II, national parks). In contrast, strict protection on low-pressure land accounts for only 4% of total protected area by 2019.²⁵

25. Strict protection is a small share of overall protection. In Appendix Table A.3, strictly protected areas account

Figure 4: Cumulative share of land protection by strictness and land pressure



NOTES: Low pressure land has initial greenness above 70 and population density below 100 people per sq. km; high land pressure land is all other land. Strict protection is defined as IUCN category Ia, Ib, or II. Limited to foundation years between 1985-2019.

In summary, the European Union has been experiencing a major land transition over the last 40 years (Fact 1). Significant parts of the EU are still under threat of deforestation or vegetation degradation (Fact 2). Despite such land pressure, protected areas target greener land across the whole population density distribution (Fact 3). Land protection is most stringent on undeveloped land with a low population density but the majority of protected areas are not subject to strict protection (Fact 4). The descriptive results support the hypothesis that land protection could be useful to preserve or strengthen greenness in the EU; they also suggest that siting often occurs in already-green, low-populated areas and that the EU applies weaker restrictions in at-risk areas. These may inhibit positive treatment effects on vegetation growth.

6 Methodology

As in Section 3, the unit of observation is a grid, $i \in \mathbb{I}$. We observe every grid i at a biennial frequency. Define periods $t \in \mathbb{T}$, where $\mathbb{T} = \{1985, 1987, \dots, 2019\}$ is the sample of years we have data for. Treated areas correspond to areas that are protected at some time $g(i) \in \mathbb{T}$, and control areas are assigned $g(i) = \infty$ to indicate that they are not treated before or during \mathbb{T} .²⁶ Define G as the group of all units treated at time g . We use subscript t for calendar time 1985 to 2019 and we define event time $e \equiv t - g$ as the number of years before or after treatment. We define the treatment indicator $D_{it} \in \{0, 1\}$ as equal to one after plot i is protected: $D_{it} = \mathbb{1}[t \geq g(i)]$.

for 8.2% of total protected areas by count and 17.9% by land area.

26. Because less than a quarter of land is treated at the end of our sample, we select control units from ‘never-treated units’ in our sample, and report robustness with control units from ‘not-yet-treated units’ in Appendix D.

We aim to estimate the treatment effect of protection on the plot-level outcomes of vegetation cover and nightlights. We define the treatment effect of protection at time g on an outcome variable Y_{it} as:

$$\tau_{igt} = \mathbb{E}[Y_{i,t} - Y_{i,g-1}|G, D_{it} = 1] - \mathbb{E}[Y_{i,t} - Y_{i,g-1}|G, D_{it} = 0], \quad (8)$$

comparing the difference in outcomes of grid i between periods t and $g - 1$ in treatment cohort G with the unobserved difference in counterfactual outcomes should grid i not have been treated.

6.1 Staggered difference-in-differences design

We apply the doubly-robust difference-in-differences estimator of Callaway and Sant'Anna (2021) to obtain country, event-time, and cohort-specific average treatment effect estimates θ_{cgt} of (8).²⁷ We estimate average treatment effects at the country level because EU member states have jurisdiction over the policy and differ in their site selection and enforcement. In addition, we might expect treatment effect heterogeneity across event time e . For example, when treated plots regenerate after protection, we expect vegetation cover to increase gradually. Finally, there is reason to expect cohort-specific treatment effects whenever site selection differs across cohorts g when more and more land gets protected (see Section 3 for details).

The doubly-robust estimator is a crucial improvement in our setting relative to standard two-way fixed effects methods for two reasons. First, we observe a staggered introduction of the European policy over a 35-year window from 1985 to 2019, where time-varying site selection could cause cohort-specific treatment effects. In such settings, standard two-way fixed effects estimation is biased. To obtain unbiased estimates, the doubly-robust estimator allows for cohort-specific matching probabilities, outcome correction models, and treatment effects.²⁸ Cohort-specific matching addresses time-varying selection bias stemming from the correlation between unobserved plot-level economic activity and treatment status. The outcome correction controls for time-varying differences in observable weather-related plot attributes. This is essential because of the weather-induced variability in NDVI. Furthermore, we can estimate specific dynamics for each cohort to explore the policy's potential heterogeneity across years-since-treatment and across cohorts.

We specify a model for the propensity weighting and the outcome correction procedures based on variables that drive vegetation growth. We rely on previous literature assessing land conditions and yields (Schlenker and Roberts 2009) to select yield-relevant variables for the inverse probability weighting. Fixed land factors include a measure of soil suitability for agriculture, elevation, climate zones, and biogeographical regions. Time-varying variables are rainfall, heating-degree days, and the length of the growing season. Additionally, we include the average greenness, population density, and growing season length, all measured in 1985, and the greenness trend between 1985-1989. To

27. We use a difference-in-differences approach instead of a regression discontinuity design (RDD) as in Turner, Haaghout, and van der Klaauw (2014). Protected area boundaries often follow land-use breaks (field edges, roads, villages), so land cover and protection change simultaneously at the border, invalidating the RDD identification assumptions as the different land covers cause discrete jumps in NDVI at the border of protected areas.

28. For completeness, we also report results from standard two-way fixed effects estimation in Appendix C.2.

test for parallel trends, we select plots treated at least three two-year periods after the first year of our outcome measures so that the first-treated cohort is in 1991.

The outcome regression adjustment linearly projects control units' change in outcomes between $g - 1$ and t on plausibly exogenous observables, separately for every cohort group G :

$$m_{gt}(X) = \mathbb{E}[Y_{i,t} - Y_{i,g-1}|G, X, D_{it} = 0].$$

By subtracting the fitted values of $m_{gt}(X)$ from the outcome difference of treated units, the procedure corrects for the confounding effects of time-varying observable differences in control variables.

The procedure involves estimating the propensity scores, outcome correction, and dynamic average treatment effect by country-cohort. We estimate standard errors using the multiplier bootstrap procedure. Because multiple plots are assigned to treatment simultaneously, we cluster at the assignment level of the CDDA identifier (Bertrand, Duflo, and Mullainathan 2004) in addition to the default plot-level clustering used by the Callaway and Sant'Anna procedure. We provide details on the estimation in Appendix B.

The estimated average treatment effects θ_{cgt} vary across countries, over event time $e = t - g$, and across treatment-assignment cohorts g . Define the set of countries as \mathcal{C} and the set of cohorts as \mathcal{G} . Recall we defined \mathbb{T} as the set of time periods in our study sample; we refer to the last year of the sample as T . We compute country-specific treatment effects by averaging over the treated time periods and cohorts:

$$\theta_c = \sum_{g \in \mathcal{G}} w_{cg} \sum_{t > g}^T w_e \theta_{cgt} = \sum_{g \in \mathcal{G}} \frac{N_{cg}}{\sum_{g \in \mathcal{G}} N_{cg}} \sum_{t > g}^T \frac{1}{T - g + 1} \theta_{cgt}, \quad (9)$$

where N_{cg} is the number of observations in country c that were treated in foundation year g , and w are weights. We average treatment effects across all cohorts; within each cohort, we average over all treated periods for that cohort in our sample. Because our panel is balanced, every event time has an equal weight $w_e = \frac{1}{T - g + 1}$ within its cohort.

We compute the overall treatment effect estimate across the EU as a weighted average of the country-specific θ_{cgt} parameters where the weights depend on the number of treated observations for each tuple $\{c, g, t\}$. Define the number of plots i treated in cohort g and observed in period t within country c as N_{cgt} . Then, define the EU-wide θ_{gt}^{EU} as:

$$\theta_{gt}^{EU} = \sum_{c \in \mathcal{C}} w_{cgt} \theta_{cgt} = \sum_{c \in \mathcal{C}} \frac{N_{cgt}}{\sum_{c \in \mathcal{C}} N_{cgt}} \theta_{cgt}. \quad (10)$$

Generally, such a weighted sum requires estimating the weights w_{cgt} . In our setting, we observe the true weights on observations because our data cover all land in the EU. As a result, we directly adjust standard errors without calculating a separate standard error for the weights. We obtain the overall θ^{EU} as the arithmetic mean over cohorts and their treated periods. The EU-wide

aggregation is asymptotically valid when treatment effects are i.i.d. across countries.²⁹

We aggregate θ_{gt}^{EU} over event time and cohorts. We first recast θ_{gt}^{EU} as θ_{ge}^{EU} by setting $e = t - g$. Aggregating over cohorts gives an event-study type estimator θ_e^{EU} indexed by event times:

$$\theta_e^{EU} = \sum_{g \in \mathcal{G}} w_g \theta_{ge}^{EU}, \quad (11)$$

where w_g equals the fraction of cohorts observed at event time e . Finally, aggregating over event times gives a measure of the average effect of being treated at time g :

$$\theta_g^{EU} = \frac{1}{T - g + 1} \sum_{e=0}^{T-g} \theta_{ge}^{EU}, \quad (12)$$

where we weight each event time within cohort g the same, implying a balanced panel.

Event-study style aggregations represent multiple underlying mechanisms if the composition of cohorts represented in each θ_e is different.³⁰ To ensure that results are not driven by sample composition alone, we construct a series of panels requiring units to be in-sample for at least three periods before and after treatment. This functionally excludes the earliest treated and latest treated plots from the sample. We also use an equivalent aggregation at the country level to examine heterogeneity in the results across countries. We aggregate within-country and across event time to obtain θ_{cg} .

6.2 Conditional average treatment estimator

In addition to differences in treatment effects across countries, event times, and protection years, we anticipate treatment effect heterogeneity due to factors such as pre-protection greenness, population density, varying soil characteristics, or weather conditions impacting vegetation regeneration ability. Absent knowledge of the universe of treatment effects, we use heterogeneity to understand how to target land protection in the context of our framework in Section 3. For example, previous agricultural land may green differently than land that was already forested. We use the conditional average treatment effect (CATE) framework put forward by Chernozhukov et al. (2024) and test directly whether there is a statistically significant deviation from the average treatment effect across the covariate space. The CATE framework estimates a high-dimensional, nonparametric function:

$$\tau_{igt}(x) = \mathbb{E}[Y_{i,t} - Y_{i,g-1}|G, D_{it} = 1, X = x] - \mathbb{E}[Y_{i,t} - Y_{i,g-1}|G, D_{it} = 0, X = x] \quad (13)$$

29. Our standard errors for the EU-level effects are likely smaller relative to a confidence band that would incorporate covariance in country-level treatment effects. However, EU member states independently decide which land to protect, and we do not observe coordination between member states in practice. The choice of aggregation technique is also motivated by the computational cost of the doubly-robust estimator.

30. Because we aggregate across cohorts, the sample is densest close to event time zero, with fewer units used to calculate effects in the earliest pre-periods and latest post-periods. For example, if a unit is treated in 2019, it has only one post-period (2019) but has 17 biennial pre-periods dating back to 1985. No other cohort is represented in the dynamic effect for the bin $t - g = -34$ as 2019 is the last in-sample cohort.

While we could introduce additional heterogeneous treatment effects in the difference-in-differences framework, we have a large set of covariates to consider. We therefore apply the honest random forest estimator of Wager and Athey (2018) to select statistically-relevant dimensions of heterogeneity, avoid multiple testing problems, and return a causally-valid conditional average treatment effect. Both Callaway and Sant’Anna (2021) and Wager and Athey (2018) are doubly-robust difference-in-differences methods as we apply them: the first is semiparametric and the second nonparametric. In contrast to current literature, where third differences are chosen by the researcher, a nonparametric method has the advantage of choosing salient dimensions of heterogeneity directly from the data. Thus, the data inform how to construct $\tau_x = \mathbb{E}[\tau|X = x]$ for covariate X .

We adapt the approach of Wager and Athey (2018) and the associated ‘grf’ software package, which is designed for a cross-sectional comparison of treated and control means, to account for the variability in treatment selection across cohorts. We provide a brief overview of our approach here, with more details in Appendix B.

For each cohort, we restrict attention to a single pre-period $g - 1$ and a single post-period, which we fix to the last year of the data, $t = 2019$. Estimation requires four steps, each involving the calibration of a separate random forest. First, we calculate a cohort-specific propensity score $p_g(X)$ in period $g - 1$ which predicts selection into treatment. Second, we calculate a regression-adjustment, $m_{gt}(X)$, using a regression tree in the period $g - 1$. These first two steps are identical to the staggered difference-in-differences design; however, here we estimate both steps using a random forest instead of a parametric estimator. Lastly, we use both calibrated trees to calculate a third causal random forest in period $g - 1$, which provides the relative pre-period expectation and a fourth causal random forest in period $t = 2019$, which provides the relevant post-period expectation:

$$\tau_{igt}(x) = \overbrace{(\mathbb{E}[Y_{i,t}|G, D_{it} = 1, X = x] - \mathbb{E}[Y_{i,t}|G, D_{it} = 0, X = x])}^{\text{post-period}} - \underbrace{(\mathbb{E}[Y_{i,g-1}|G, D_{it} = 1, X = x] - \mathbb{E}[Y_{i,g-1}|G, D_{it} = 0, X = x])}_{\text{pre-period}}$$

In these final steps, the random forest imputes counterfactual outcomes for each treated unit by splitting the data into groups using controls X (our approach is related to Knittel and Stolper 2021). Each split is chosen to maximize the between-split heterogeneity in treatment-control differences. Splits with excessive imbalance in (propensity-weighted) treated and control units are penalized. As a result, the finest split of each “tree”—a leaf—grown by the random forest contains a set of treated units and a counterfactual group comprised of their nearest neighbors. In each period, we thus obtain a vector of treatment-control differences associated with each grid cell. We then evaluate grid-level treatment effects by subtracting the pre-period effect from the post-period effect to arrive at a CATE difference-in-differences estimator.

7 Results

In this section, we present average treatment effects on NDVI and other outcomes. Next, we investigate (the lack of) treatment effect heterogeneity and its possible explanations. Finally, we link our empirical results to the four theoretical predictions we formulated in Section 3.

7.1 Average treatment effects

Average treatment effects on NDVI by country

Figure 5 presents the EU-wide average treatment effect θ_{EU} on vegetation greenness aggregated across all countries, event times, and cohorts, along with the country-specific ATEs θ_c ordered from small to large (red and blue bars). Detailed results for Figure 5 and all subsequent figures are presented in Appendix C. The top panel summarizes θ_{EU} across different econometric specifications. The preferred EU-wide ATE via the doubly-robust Callaway and Sant'Anna estimator is 0.02 with a standard error of 0.07 for a change in the NDVI index (which ranges between 0 and 100).³¹

We cannot reject a zero effect of protection on vegetation greenness, and the confidence interval is tight around zero: the EU-wide average treatment effect estimate is a precise zero. The figure also displays the estimates from two naive, and biased, two-way fixed effects estimators (with and without matching).³² Their EU-wide ATE of 0.5 is similarly small. In Appendix D, we conduct several robustness checks. By looking at the first differences in greenness as an outcome, we confirm that the null result and pre-treatment parallel trend are robust to changes in functional-form assumption (Roth and Sant'Anna 2023). We limit the sample to NDVI between 40 and 100 to focus on greener plots and corroborate the zero result. We use not-yet-treated plots instead of never-treated plots as controls and confirm our results. We focus on a limited high-population sample, for which we do not find parallel trends before treatment, but find no evidence for positive treatment effects. Finally, we estimate a spatial first difference model and also find a zero effect.³³

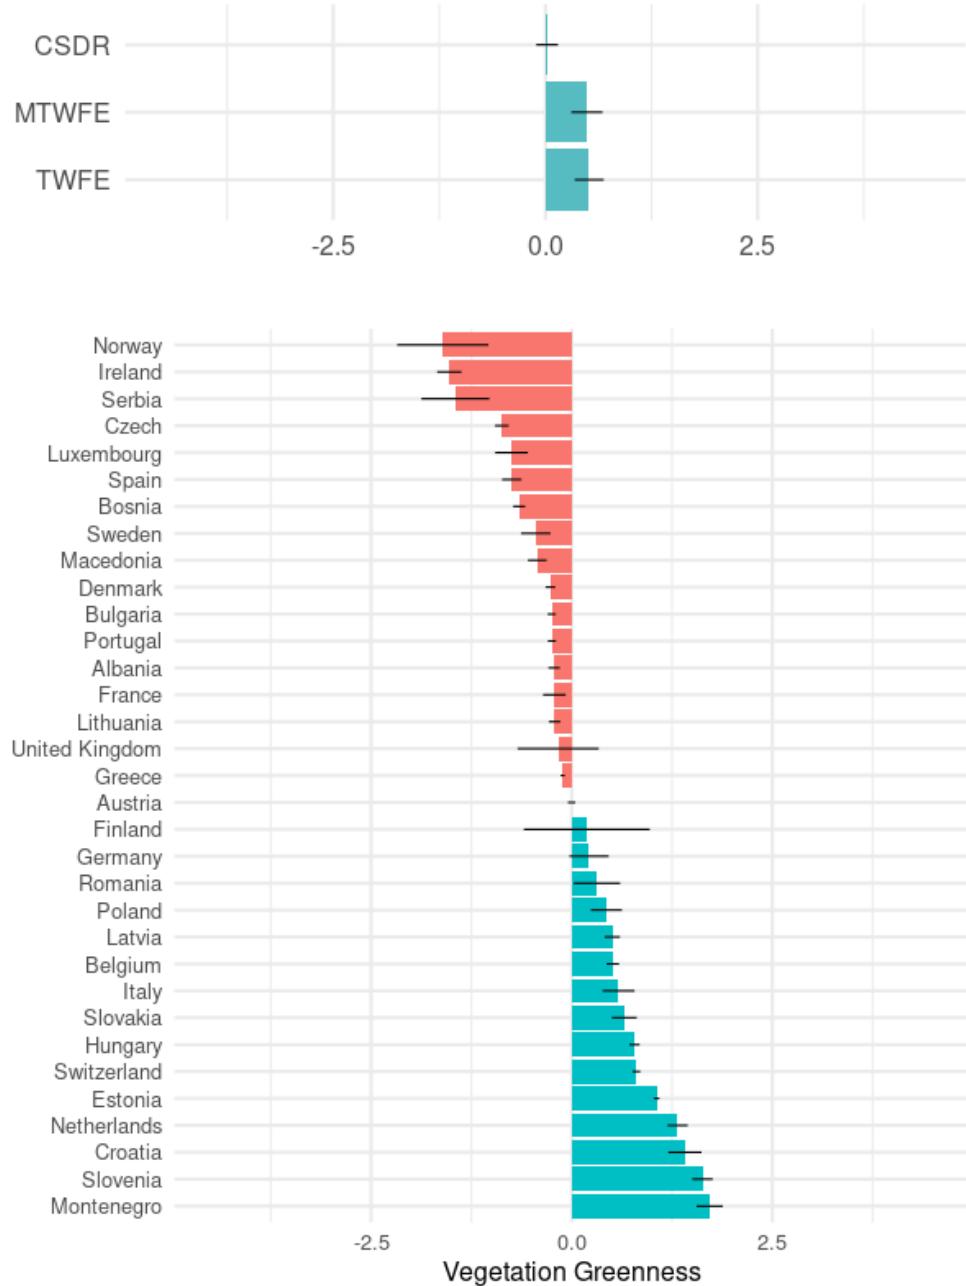
There are some differences across countries (bottom panel). However, the size of all estimated treatment effects in absolute terms is typically small— θ_c varies between -1.5 and 2.5. In 18 out of 32 of the countries, the treatment effect is less than a 0.5 point change in the NDVI index, which is small relative to NDVI index difference of 14 points we see on average between forest (63) and cropland (49) or the 10 point difference between forest and grassland (53). The treatment effects are also small relative to the standard deviation of within-country mean-adjusted greenness

31. Appendix C.1 shows balance tables with and without cohort-specific matching.

32. For the two-way fixed effects regression, we estimate a difference-in-differences model with grid- and time- fixed effects, λ_i and λ_t : $Y_{it} = \beta D_{it} + \lambda_i + \lambda_t + \epsilon_{it}$. We report the dynamic estimates from the two-way fixed effects model in Appendix C.2.

33. Because our results indicate a zero treatment effect of protection on greenness, we are not concerned about spillovers. Robalino, Pfaff, and Villalobos (2017) show empirically that protected areas in Costa Rica facing greater threats of deforestation present larger spillovers on nearby land. Leite Mariante and Salazar Restrepo (2024) develop a model to identify the general equilibrium spillovers of conservation in the Brazilian context where conservation has positive treatment effects on forests.

Figure 5: Treatment effects on NDVI, EU-wide and by country



NOTES: Top panel reports treatment effects estimated through three distinct estimators: the Callaway and Sant'Anna doubly-robust procedure (CSDR), two-way fixed effects (TWFE), and matched two-way fixed effects (MTWFE). Bottom panel reports CSDR treatment effects θ_c aggregated by country with bootstrapped confidence intervals. Blue bars indicate positive ATEs; red bars indicate negative ATEs. Horizontal black lines indicate 95% confidence intervals. Greenness varies from 0 to 100. Appendix Table A.10 presents the regression results.

(14). Save for Croatia, the Netherlands, Montenegro, Slovenia, Serbia, and Ireland, every country has a point estimate less than 1 in absolute value. In the distribution, more than half of the countries have small negative treatment effects, and about a third of the treatment effects are slightly positive. Of the five largest countries by surface area (France, Spain, Sweden, Germany, and Finland), only Germany has a positive treatment effect (0.2 NDVI points). Altogether, Figure 5 shows some meaningful treatment effects in a handful of mostly smaller EU countries, but generally large positive treatment effects are rare, and the Europe-wide average treatment effect is small by any standard.

Average treatment effects on NDVI across event time and cohorts

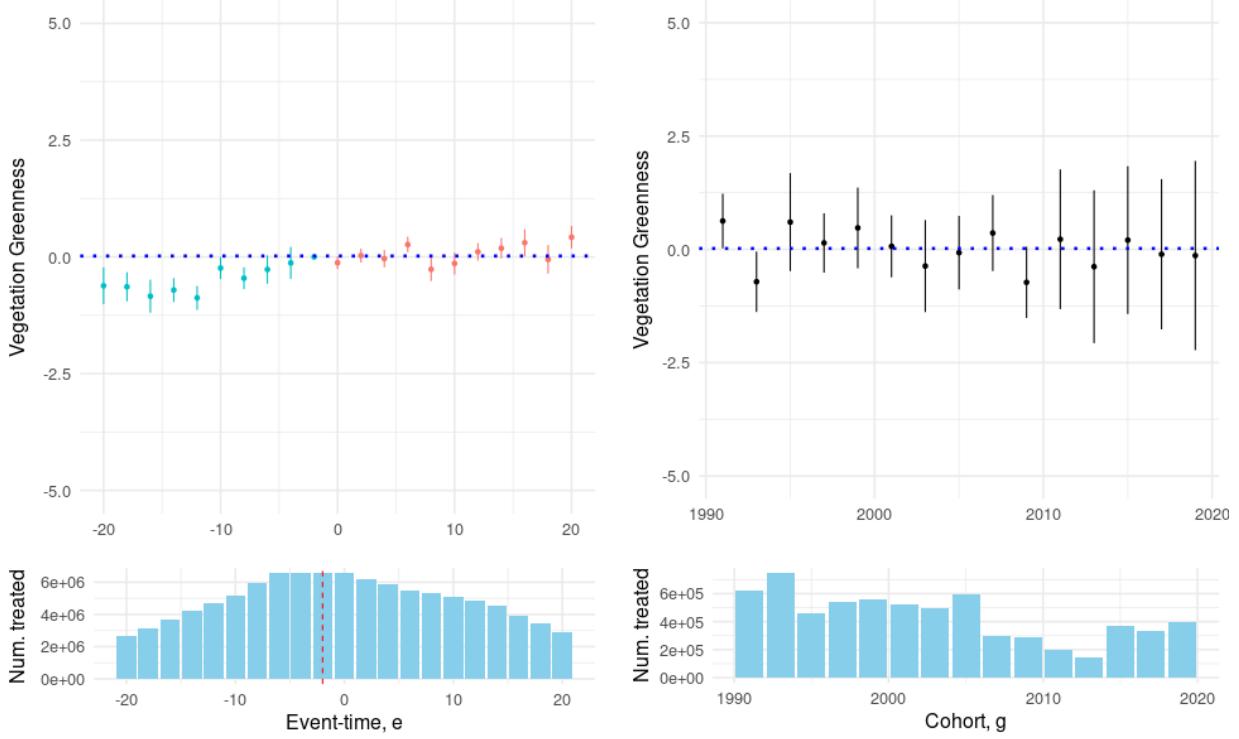
Figure 6 (left panel) shows the ATE estimates θ_e^{EU} and associated 95% confidence intervals—aggregated across the entirety of the European Union—over event time e , with $e = 0$ indicating when protected areas were established. These dynamics reveal no post-protection upward trend in treatment effects over time. The effect of conservation on greenness is consistently close to zero—almost all confidence intervals are contained within $[-1, 1]$ up to 20 years post-protection. Even 30 years after treatment, we find no evidence of a positive protection effect. The lack of a pre-trend is reassuring and suggests the matched controls are appropriate comparisons for the treated areas.³⁴ This does not imply that vegetation greenness has been constant over time. Certain parts of Europe have been greening, but treatment and control plots in those areas have been greening in parallel, both before and after the establishment of land-protection policies.

Even if the average treatment effect is zero, we investigate if there is a time trend in the treatment effects θ_g^{EU} by cohort of protection g . As discussed in Prediction 2 of Section 2, one might expect that land with low opportunity costs is protected first (small treatment effect). As time progresses and the land with the lowest opportunity cost has been protected, local governments might focus on areas with higher opportunity costs (larger treatment effect). Figure 6 (right panel) shows the EU-wide cohort-level treatment-effect estimates. There is no time trend in the treatment effects across cohorts—the effect of land protection is close to zero regardless of when the land got protected.

To formally establish this result, we compute two test statistics applied at the country level in Appendix Table A.13. First, we test if there are statistically-significant linear trends in the country-cohort level estimates θ_{cg} . Second, for each country, we split the estimate θ_{cg} in an early-treated and late-treated group and test for a difference in the treatment effect size between the two groups. With these two tests, we find no evidence of any meaningful positive trends across cohorts of protection at the individual country level.

³⁴ There are some pre-treatment coefficients that are significantly different from zero, but these are close to zero and would likely not imply large changes to the post-treatment confidence intervals (Rambachan and Roth 2023).

Figure 6: Treatment effects on NDVI over event time and by cohort



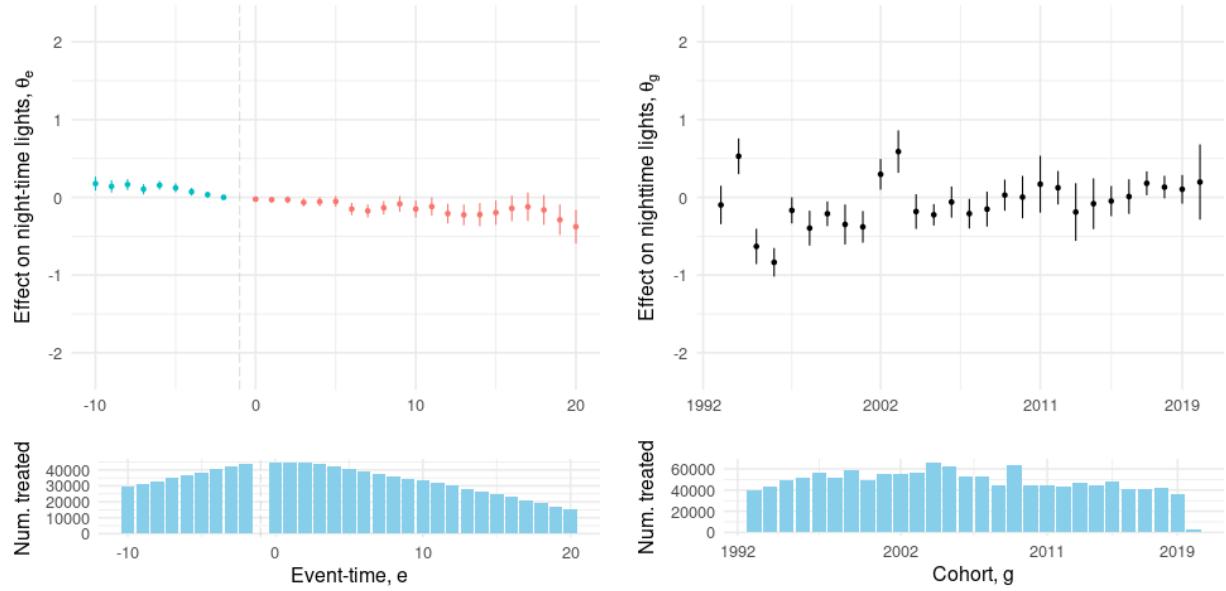
NOTES: Treatment effects on greenness aggregated by event-study period θ_e^{EU} in the left panel, and by cohort θ_g^{EU} in the right panel. Bottom panels indicate the number of treated grid cells for each event-time or cohort. The event-time design duplicates observations across event periods, while the cohort design splits the sample across cohort dates. Event-times in the left panel trimmed to $[-20, 20]$. Dashed blue line shows the overall EU-wide effect from Appendix Table A.10. Greenness varies from 0 to 100. Both panels use the Callaway and Sant'Anna doubly-robust estimator with 95% bootstrapped confidence bands. Appendix Tables A.11 and A.12 present the regression results.

Average treatment effects on other outcomes

Having established a mostly-precise zero effect of the effect of land protection on NDVI across countries, event times, and cohorts, we now present evidence on the effect of protection on nightlights. One might hypothesize that protection reduces human activity relative to matched controls, which would manifest itself through a decrease in the brightness of nighttime light.

Figure 7 shows that—whether broken down by event time or by cohort—there is no evidence that land protection reduced nightlights. Nightlights vary from 0 to 68 units of luminosity, while treatment effects are generally within 0.5 units. The EU-wide average treatment effect is -0.14 (0.05). For reference, moving from a bright country such as Belgium, with a luminosity of 30.5, to a dark country such as Bulgaria, with a luminosity of 9.3, implies a 20-point difference in the measure. Our results suggest a zero impact on human activity in areas set aside for protection for at least two decades (left panel). In addition, the effect of land protection is close to zero regardless of the year of first protection. This corroborates our NDVI results (right panel).

Figure 7: Treatment effects on night-time lights over event time and by cohort



NOTES: Treatment effects on night-time lights aggregated by event-study period θ_e^{EU} in the left panel, and by cohort θ_g^{EU} in the right panel. Bottom panels indicate the number of treated grid cells for each event-time or cohort. Event-times in the left panel trimmed to $[-10, 20]$. Night-time lights vary from 0 to 68, with all event-time estimates being smaller than 0.5 units of luminosity. Both panels use the Callaway and Sant'Anna doubly-robust estimator with bootstrapped 95% confidence bands. Appendix Tables A.15 and A.16 present the regression results.

It is conceivable that land protection increased biodiversity in ways neither measured by NDVI nor nightlights. The lack of comprehensive, granular, consistent, and reliable species count data prevents us from testing this hypothesis causally. Appendix E contains some limited, non-causal, descriptive results using the BioTIME data. This event-study analysis shows no clear increase in species counts following CDDA establishments near the BioTIME study locations.

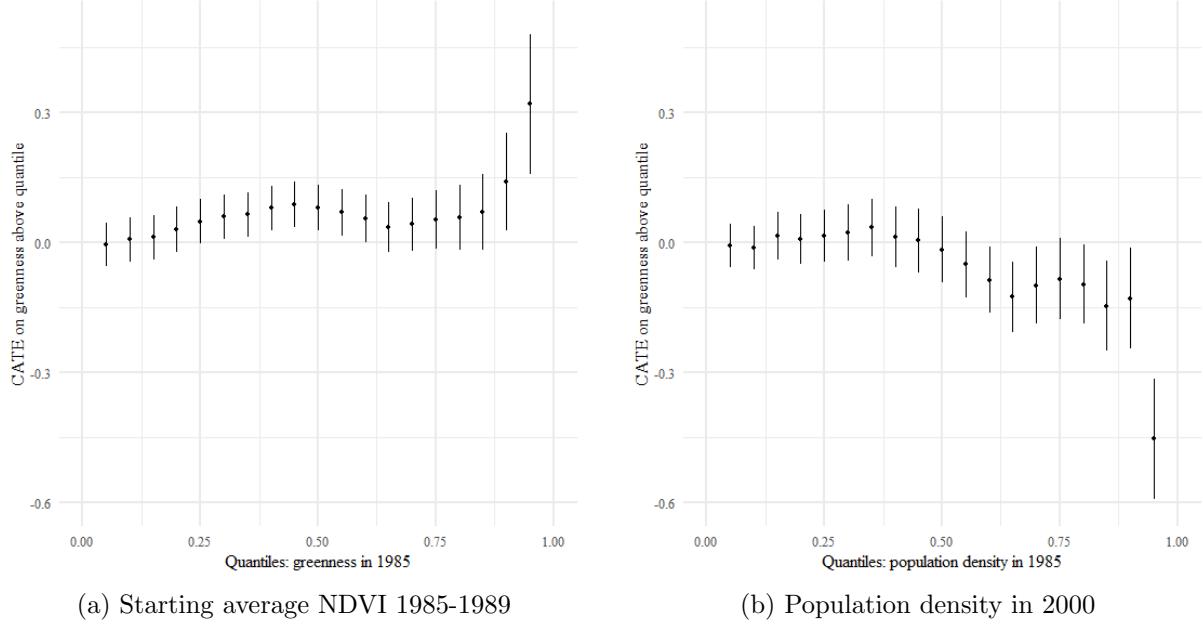
Finally, we confirm that our main results are consistent with evidence from discrete land-use data. Despite the limitations of discrete classifications of land use (see Section 4), we use decadal data from HILDA in Table 2 to compare land-use transition probabilities between protected and never-protected areas. We limit the transitions to our study period, 1980 to 2010, to facilitate comparison with our previous empirical exercises. To focus on a consistent 30-year time window, we limit treated observations to areas protected between 1970 and 1980. Controls are matched on 1970 land-use and time-invariant observables: slope angle, slope steepness, solar radiation, long-run precipitation, and distance to a shoreline. The forest transition shares are very similar for protected and control units, which is consistent with our null results for continuous vegetation greenness. The largest discrepancy is a 3% larger transition of cropland into grassland in protected areas. However, the grassland-cropland classification is particularly difficult for discrete land-use measures and may be confounded by pasture.

Table 2: Comparison of discrete land-use transitions in protected areas and never-protected areas, 1980-2010

1980 / 2010			Cropland	Forest	Grassland	Other	Settlement	Water
Control	Cropland	% row	84.8	1.6	13.2	0.0	0.5	0.0
	Forest	% row	0.2	97.5	2.2	0.0	0.1	0.0
	Grassland	% row	3.8	10.8	85.1	0.0	0.3	0.0
	Other	% row	0.0	0.0	0.0	100.0	0.0	0.0
	Settlement	% row	1.1	0.3	0.5	0.0	98.2	0.0
	Water	% row	0.1	0.1	0.4	0.0	0.2	99.2
Treated	Cropland	% row	80.7	2.7	16.1	0.0	0.5	0.0
	Forest	% row	0.3	98.2	1.3	0.0	0.1	0.0
	Grassland	% row	5.4	10.1	84.3	0.0	0.2	0.0
	Other	% row	0.0	0.0	0.0	100.0	0.0	0.0
	Settlement	% row	1.3	0.4	0.4	0.0	97.9	0.0
	Water	% row	0.0	0.0	0.0	0.0	0.0	100.0
Total (2010)		% row	27.3	35.3	28.7	1.5	4.2	2.9

NOTES: Table reports land-use transitions relative to the base year of 1980 in 2010. We tabulate transitions by treatment status: treated units are treated in 1970-80 and have at least 50% of landmass in a protected area. Controls are matched on pre-1970 observables. Transitions are defined based on the HILDA land cover data, which classifies land into one of six land categories in each decade from 1900 to 2010.

Figure 8: Distribution of the conditional average treatment effect for initial greenness and population density



NOTES: Both panels plot estimates of the conditional average treatment effect across 20 ventiles of a key explanatory variable. Bar at ventile $q \in \{0.05, 0.1, \dots, 0.95\}$ corresponds to the CATE of plot i conditional on a realization of X above that quantile, e.g., $\mathbb{E}[\tau_i | X > x(q)]$. Confidence bands are 95% confidence intervals constructed through 30 bootstraps of a 1% stratified sample of the EU.

7.2 Testing for treatment-effect heterogeneity

As discussed in Section 6.2, our average null finding does not reject the existence of a (small) tail of high-impact conservation efforts. We investigate whether there is meaningful underlying heterogeneity in treatment effects by estimating conditional average treatment effects and by correlating the country-level estimates from Figure 5 with explanatory variables related to the political economy of our setting.

To estimate conditional average treatment effects, we apply the Wager and Athey (2018) estimator on a random sample of the data (whose construction is discussed in Appendix B.2). The sample average treatment effect on NDVI is 0.01 (0.26). Our data reject the existence of any meaningful heterogeneity. While our random forest detects statistically significant heterogeneity in the test of Chernozhukov et al. (2024), economically meaningful heterogeneity is limited. Splitting the treatment effect distribution at the median, the CATE in the bottom half of the CATE distribution is 1.75 (0.76) NDVI units smaller than that in the top half. We thus estimate limited heterogeneity around zero across the whole covariate distribution, not just for average land.

In Figures 8a and 8b, we highlight two important dimensions on which we can reject meaningful heterogeneity. These figures plot conditional average treatment effects at ventiles (percentiles in increments of 5%) of key covariates. Figure 8a shows more sparsely (or barely) vegetated areas,

as measured in the first years of the sample, do not green faster due to land protection than areas with higher initial greenness. Conditional on any initial greenness, average treatment effects are not significantly positive, except in the very last ventiles, where the CATEs equal 0.15-0.3. Figure 8b plots conditional average treatment effects across population density ventiles, a proxy for land pressure. We find no evidence that densely populated land greens as a consequence of protection: if anything, they green less than comparable controls without protection (this is consistent with the evidence presented in Appendix Figure A.12). Together, Figures 8a and 8b imply that less green, higher population density lands that likely face stronger economic development pressure—such as areas around urban centers—are not re-greening faster after protection relative to similar, unprotected areas. Additionally, the conditional average treatment is approximately zero across the covariate range.

In the appendix, we also find no significant variation in treatment effects when considering covariates linked to high-quality agricultural land. The left panel of Appendix Figure A.8 shows CATEs along the distribution of the soil phosphorus content. Significant soil phosphorus deposits can aid the development of “roots and shoots” and often also mark prior agricultural fertilizer use. The findings show no notable difference in greening patterns between areas with different soil phosphorus content. Thus, protecting observably more valuable agricultural land does not lead to measurably richer vegetation. The right panel shows the CATE by growing season length in 1985. Longer growing seasons indicate climatic conditions that are more conducive to growing common field crops. If anything, protected areas with longer growing seasons saw somewhat lower treatment effects on NDVI than comparable controls.

Next, we explore the average treatment effect as a function of the strictness of the land-protection regime. We calculate the average treatment effect using the staggered difference-in-differences estimator, restricting treated units to each IUCN category (see our prior discussion of Figure 3) in Table 3.³⁵ We find that the estimated average treatment effect is close to zero across all IUCN categories. Estimates are significantly discernible; the largest positive effect on NDVI is among type Ib wilderness areas, which are the second strictest group of IUCN codes. However, this effect is 0.55 NDVI points, an economically small gain in greenness. We conclude there is no clear gradient in treatment effects across this strictness measure.

Finally, we correlate the country-level treatment effects from Figure 5 with a range of explanatory variables in Table 4. While most country-level treatment effects are close to zero, the gap between the smallest (Serbia) and the largest treatment effect (Croatia) amounts to about four NVDI points. We explore if this heterogeneity can be explained by different economic, political, or land endowment factors. We present correlation coefficients from a bivariate regression using a country’s GDP per capita, a corruption perception index, environmental expenditures as a percent of GDP and in levels, a farmland bird index, urbanized land shares, and the number of NGO actions as controls. Overall, we find statistically noisy coefficients with unexpected signs. For

³⁵ For memory management, controls are a 1% random sample of never-treated land.

Table 3: ATE in each IUCN category

Code	Description	ATE (SE)	CDDAs
Ia	Strict nature reserve	-0.69 (0.04)	5,829
Ib	Wilderness area	0.55 (0.06)	777
II	National Park	0.26 (0.07)	431
III	Natural monument	-0.05 (0.07)	3,669
IV	Species management area	-0.16 (0.02)	39,472
V	Protected landscape	0.28 (0.01)	11,372
VI	Protected area with sustainable use of natural resources	-0.43 (0.05)	1,215

NOTES: Table summarizes the Callaway and Sant'anna estimator for the average treatment effect as a function of protected area strictness. Controls are a 1% sample of untreated units within each country, while treated units comprise the universe of each IUCN category. Excludes protected areas with missing, not assigned, or not reported IUCN categories.

example, while enforcement of land-use policy is often imperfect, as seen in the Amazon rainforest (Keles, Pfaff, and Mascia 2022), we expect treatment effects to vary between countries with varying degrees of institutional strength. However, our results show small effects in the vast majority of EU countries, and there is no clear pattern suggesting weaker effects in countries with a history of limited enforcement. While we cannot directly measure enforcement, we test for correlation between the country-level ATEs and an index of apparent corruption from www.transparency.org. The linear regression of the average treatment effect on the log of this index, measured from 0 to 100, has a coefficient of 0.44 (0.46). We thus find no evidence of a relationship between a measure of state enforcement capacity and land-protection outcomes. Likewise, we count the number of environmental NGO campaigns targeting firms and addressing biodiversity-related issues in each country. We find a near-zero effect of NGO pressure on protection.

Table 4: Bivariate decomposition of country-level average treatment effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log GDP per capita	0.10 (0.27)							
Log corruption perceptions		0.44 (0.46)						
Log environmental expenditures (%GDP)			0.74* (0.42)					
Log environmental expenditures (m. EUR)				-0.13 (0.11)				
Log farmland bird index					0.24 (0.66)			
Log urbanized share						0.14 (0.56)		
Log num. NGO actions							0.04 (0.09)	
CATE								0.09 (0.11)
R^2	0.004	0.03	0.12	0.05	0.007	0.002	0.008	0.03
Observations	33	33	26	26	22	31	30	26
Independent variable mean	40,576	62	0.43	2,857	75	0.39	56	0.60

NOTES: We regress country-level ATEs in levels from Table A.10 on a set of country-level variables. GDP per capita from the Penn World Table, v10.01. Corruption perceptions index from Transparency.org data (www.transparency.org). Environmental expenditures, farmland bird index, and urbanized share are supplied by EUROSTAT. NGO actions data come from the SigWatch database, and focus on the total biodiversity-related actions conducted 2011-2024, by country (www.sigwatch.com). CATEs from our own estimation. All are measured in the most recent year in which there is a non-missing value for the given country. Standard errors are IID (country-level). Regressions in (1)-(7) are weighted by the inverse standard error of country-level ATEs. Regression in (8) are weighted by the inverse variance of country-level CATE and ATTs, assuming zero covariance.

We also correlate the country-level CATE estimates with the country-level doubly-robust difference-in-differences estimates. Both methods use different identification assumptions. We find that a positive CATE is not correlated with a positive treatment effect in Figure 5, again pointing to an absence of economically meaningful heterogeneity across countries and this limited heterogeneity being explained by statistical noise. Our estimates support a zero effect throughout the distribution of protection and land covariates, and we find little evidence that any of the economic, political, or land endowment factors we include in Table 4 explain the modest differences in treatment effects across countries.

7.3 Discussion of the results

We now offer several potential explanations for the zero treatment effect of land protection. We first interpret our empirical results in the context of the model and predictions in Section 3. Our model can rationalize the zero results from the empirical analysis under a specific set of circumstances.

Prediction 1 says that a social planner will only protect land when the ecological treatment effect is positive and outweighs the economic losses; infra-marginal land not under pressure of development would never be protected. Our empirical results in Figure 5 contradict this setting of a social planner engaging in unconstrained, optimal protection. Prediction 2 does not rule out that treatment effects may be constant across cohorts (as in our empirical results) but, by Prediction 1, the cohort-level treatment effects should be positive—this is not the case in our empirics (Figure 6).

The presence of preference misalignment, green glow, or a binding quantity target by itself, covered in Prediction 3, does not explain our results either. Each of these deviations from the planner’s problem would imply that we observe positive treatment effects before the local policymaker resorts to protecting infra-marginal land. Yet, we find zero effects from the first to the final cohort.

A combination of preference misalignment and delegated implementation can rationalize the observed null treatment effects in the EU. Consistent with Prediction 3, when local decision makers place little weight on ecological outcomes but derive benefits from designation itself (green glow) or are required to meet area-based targets, they optimally select land with very low baseline land-use intensity. Protecting such infra-marginal land minimizes economic and enforcement costs, but generates little or no change in land use and hence null ecological effects, even in early protection cohorts.

Prediction 4 highlights a complementary channel through which local authorities comply with EU mandates, operating via protection stringency. Even when protected areas are designated on land facing substantial development pressure, local authorities may choose weak protection regimes that impose only limited restrictions on economic activity. This expands the set of plots on which protection is effectively of little impact, allowing local policymakers to meet area-based targets while avoiding large local economic losses.

Empirically, Figure 4 shows that strict protection on low pressure land accounts for only about

4% of protected areas. In contrast, 54% of protected areas are subject to weaker protection and are located on land facing higher development pressure. Taken together, these patterns suggest that local authorities respond to the EU’s push for expanded protection primarily by designating areas and choosing protection regimes that do not severely restrict local economic activity.

We also consider explanations for the absence of a protection effect outside of our model. One possibility is that protected plots were under threat of *future* development rather than contemporaneous economic activity. Another is that authorities do not enforce the restrictions of economic activity. Empirically, these explanations seem implausible. Our long panel allows us to identify treatment effects up to 30 years after the first treatment in 1991. Even for the earliest cohorts, we do not see any positive treatment effects decades later. As we discussed in Section 7.2, there is no evidence of treatment-effect heterogeneity that lines up with enforcement heterogeneity. Any alternative explanation needs to explain that treatment effects are close to zero across every period, country, and the whole covariate space.

We conclude that the explanations underlying Predictions 3 and 4 in our model are the most likely: the zero effect can be explained by a combination of targeting land not at risk of economic development and choosing weak protection levels for land in areas that face development pressure. Our descriptive evidence in Section 5 confirms that land protection targets greener areas across both sparsely and densely populated regions, and that strict protection is more likely on sparsely populated and initially green land. The latter aligns with the political economy of the EU setting. EU countries use a multi-stakeholder process by which new protected areas are added—consensus will be easier to reach when economic opportunity costs are low. The EU Commission is explicitly imposing area-based protection targets on member states, yet protection decisions are often delegated to local jurisdictions whose preferences for biodiversity protection are much weaker and who might engage in protection for green-glow reasons more so than to increase broader ecological value.

When interpreting our results, it is important to mention that NDVI has important strengths as well as limitations. As discussed in Section 4, NDVI has been found to correlate well with measures of biodiversity beyond vegetation cover, such as bird diversity. That said, NDVI will not be able to detect other types of biodiversity improvements. For example, species management areas may improve insect counts, which may not always correlate with vegetation cover. However, our overall zero result is quite pervasive—across almost all countries, time periods, areas of different greenness and land use types—and provides compelling evidence that land protection has not generally been effective, even if we cannot rule out potential improvements in certain biodiversity measures on particular types of land. One would need much higher-resolution data on species counts over long periods of time to test the causal effects on biodiversity beyond vegetation.

8 Conclusions

Policies to protect a quarter of the EU’s landmass have not led to a change in vegetation cover measured by NDVI or human activity measured by nightlights. We find no meaningful heterogeneity in treatment effects across event time, protection cohorts, strictness of protection, population density, land, soil, or climate characteristics. Modest differences in treatment effects across countries cannot be explained by political economy variables. Overall, control plots show equal vegetation cover trends as treated plots. This finding is new and important, as it helps evaluate whether the siting decisions and stringency of protected areas are preventing land development and ecosystem degradation. It also provides the first large-scale evidence of the effects of land protection in advanced economies, which are critically important to meet global protection targets but exhibit very different land-use and enforcement dynamics than in tropical forest countries, which is where the literature has focused until now.

We attribute the absence of detectable treatment effects to the political economy of EU land protection, in particular to the interaction between area-based protection targets and the delegation of siting and implementation decisions to local authorities who place limited weight on broader ecological benefits. In practice, strict protection has been disproportionately directed toward land facing little development pressure, while land under substantial pressure has typically been designated under weak protection regimes that impose few constraints on economic activity. Although policy debates in the EU and beyond emphasize achieving area-based targets—such as protecting 30% of land area—it is the choice of specific sites and their counterfactual trajectories in the absence of protection that ultimately determine policy effectiveness. Our findings do not imply that protection will never restrict economic activity in the long run, but it does suggest that Europe’s current land-protection regime has been poorly targeted and has favored low-impact designations in economically contested areas.

References

Abman, Ryan. 2018. “Rule of Law and Avoided Deforestation from Protected Areas.” *Ecological Economics* 146:282–289.

Alix-Garcia, Jennifer, and Daniel Millimet. 2022. “Remotely Incorrect? Accounting for Non-classical Measurement Error in Satellite Data on Deforestation.” *Journal of the Association of Environmental and Resource Economists* 10 (5): 1335–1367.

Andam, Kwaw S., Paul J. Ferraro, Alexander Pfaff, G. Arturo Sanchez-Azofeifa, and Juan A. Robalino. 2008. “Measuring the Effectiveness of Protected Area Networks in Reducing Deforestation.” *Proceedings of the National Academy of Sciences* 105 (42): 16089–16094.

Aronoff, Daniel, and Will Rafey. 2023. “Conservation Priorities and Environmental Offsets: Markets for Florida Wetlands.” NBER Working Paper No. 31495.

Aspelund, Karl M., and Anna Russo. 2024. “Additionality and Asymmetric Information in Environmental Markets: Evidence from Conservation Auctions.” Working Paper.

Assunção, Juliano, Clarissa Gandour, and Romero Rocha. 2023. “DETER-ing Deforestation in the Amazon: Environmental Monitoring and Law Enforcement.” *American Economic Journal: Applied Economics* 15 (2): 125–156.

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

Auffhammer, Maximilian, Eyal Frank, David McLaughlin, Beia Spiller, and David Sundig. 2021. “The Cost of Species Protection: The Land Market Impacts of the Endangered Species Act.” Working Paper, https://conference.nber.org/conf_papers/f157199.pdf.

Bahrami, Golnaz, Matthew Gustafson, and Eva Steiner. 2024. “The Biodiversity Protection Discount.” Available at SSRN.

Balboni, Clare, Robin Burgess, and Benjamin A. Olken. 2023. “The Origins and Control of Forest Fires in the Tropics.” Working Paper, https://www.robinburgess.com/s/Balboni_etal_2020_The-Origins-and-Control-of-Forest-Fires-in-the-Tropics.pdf.

Barbier, Edward B., Joanne C. Burgess, and Alan Grainger. 2010. “The Forest Transition: Towards a More Comprehensive Theoretical Framework.” *Land Use Policy* 27 (2): 98–107.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. “How Much Should We Trust Differences-in-Differences Estimates?” *Quarterly Journal of Economics* 119 (1): 249–275.

Bickenbach, Frank, Eckhardt Bode, Peter Nunnenkamp, and Mareike Söder. 2016. “Night Lights and Regional GDP.” *Review of World Economics* 152:425–447.

Bonthoux, Sébastien, Solenne Lefevre, Pierre-Alexis Herrault, and David Sheeren. 2018. “Spatial and Temporal Dependency of NDVI Satellite Imagery in Predicting Bird Diversity over France.” *Remote Sensing* 10 (7): 1136.

Börner, Jan, Dario Schulz, Sven Wunder, and Alexander Pfaff. 2020. “The Effectiveness of Forest Conservation Policies and Programs.” *Annual Review of Resource Economics* 12 (1): 45–64.

Burgess, Robin, Francisco J.M. Costa, and Benjamin A. Olken. 2023. “National Borders and the Conservation of Nature.” Working Paper, https://www.robinburgess.com/s/Burgess-Costa_Olken_2019_The-Brazilian-Amazons-Double-Reversal-of-Fortune.pdf.

Burgess, Robin, Matthew Hansen, Benjamin A. Olken, Peter Potapov, and Stefanie Sieber. 2012. “The Political Economy of Deforestation in the Tropics.” *Quarterly Journal of Economics* 127 (4): 1707–1754.

Callaway, Brantly, and Pedro H.C. Sant’Anna. 2021. “Difference-in-Differences with Multiple Time Periods.” *Journal of Econometrics*, Themed Issue: Treatment Effect 1, 225 (2): 200–230.

Chaisemartin, Clément de, and Xavier D’Haultfœuille. 2020. “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects.” *American Economic Review* 110, no. 9 (September): 2964–96.

Chen, Xi, and William Nordhaus. 2011. “Using Luminosity Data as a Proxy for Economic Statistics.” *Proceedings of the National Academy of Sciences* 108 (21): 8589–8594.

Cheng, Audrey, Katharine R.E. Sims, and Yuanyuan Yi. 2023. “Economic Development and Conservation Impacts of China’s Nature Reserves.” *Journal of Environmental Economics and Management* 121 (September): 102848.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val. 2024. “Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments.” Forthcoming, *Econometrica*.

Donaldson, Dave, and Adam Storeygard. 2016. “The View from Above: Applications of Satellite Data in Economics.” *Journal of Economic Perspectives* 30 (4): 171–198.

Druckenmiller, Hannah, and Charles A. Taylor. 2022. “Wetlands, Flooding, and the Clean Water Act.” *American Economic Review* 112 (4): 1334–1363.

Einhorn, Catrin. 2022. “Nearly Every Country Signs on to a Sweeping Deal to Protect Nature.” *The New York Times* (December).

European Commission. 2011. *Our Life Insurance, Our Natural Capital: An EU Biodiversity Strategy to 2020*. Document 52011DC0244, Accessed: 2024-05-22. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52011DC0244>.

European Environment Agency. 2020. *State of Nature in the EU*. Technical report 10/2020. <https://www.eea.europa.eu/publications/state-of-nature-in-the-eu-2020>.

European Parliament. 2023. *Report on the Proposal for a Regulation of the European Parliament and of the Council on Nature Restoration*. Document A9-0220/2023, Accessed: 2024-05-22. https://www.europarl.europa.eu/doceo/document/A-9-2023-0220_EN.html.

European Union. 2009. *Directive 2009/147/EC of the European Parliament and of the Council of 30 November 2009 on the Conservation of Wild Birds*. Accessed: 2024-05-22. <https://eur-lex.europa.eu/eli/dir/2009/147/oj>.

Eurostat. 2021. *Forests, Forestry and Logging*. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Forests,_forestry_and_logging.

———. 2022. *Protected Areas: Over a Quarter of EU Land*. <https://ec.europa.eu/eurostat/web/products-eurostat-news/-/edn-20220521-1>.

Frank, Eyal G. 2024. “The Economic Impacts of Ecosystem Disruptions: Costs from Substituting Biological Pest Control.” *Science* 385 (6713): eadg0344.

Fuchs, Richard, Peter H. Verburg, Jan G.P.W. Clevers, and Martin Herold. 2015. “The Potential of Old Maps and Encyclopaedias for Reconstructing Historic European Land Cover/Use Change.” *Applied Geography* 59:43–55.

Gibson, John, Susan Olivia, Geua Boe-Gibson, and Chao Li. 2021. “Which Night Lights Data Should We Use in Economics, and Where?” *Journal of Development Economics* 149:102602.

Goodman-Bacon, Andrew. 2021. “Difference-in-Differences with Variation in Treatment Timing.” Themed Issue: Treatment Effect 1, *Journal of Econometrics* 225 (2): 254–277.

Hancock, Alice, and Sam Jones. 2024. “EU Approves Nature Law in Face of Austrian Legal Threat.” Accessed 11 July 2025, *Financial Times* (June). <https://www.ft.com/content/7175b191-04a2-406f-8149-cc1974d5b308>.

Harstad, Bård. 2023. “The Conservation Multiplier.” *Journal of Political Economy* 131 (7): 1731–1771.

Harstad, Bård, and Torben K. Mideksa. 2017. “Conservation Contracts and Political Regimes.” *Review of Economic Studies* 84 (4): 1708–1734.

Hassan Souheil, Danielle Boivin, Laurent Germain, and Robert Douillet. 2011. “Guide Néthodologique d’Elaboration des Documents d’Objectifs Natura 2000,” no. 82, <https://www.natura2000.fr/documentation/references-bibliographiques/guide-methodologique-elaboration-documents-objectifs>.

Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi. 2014. “Growth and Structural Transformation.” *Handbook of economic growth* 2:855–941.

Hobi, Martina L., Maxim Dubinin, Catherine H. Graham, Nicholas C. Coops, Murray K. Clayton, Anna M. Pidgeon, and Volker C. Radeloff. 2017. “A Comparison of Dynamic Habitat Indices Derived from Different MODIS Products as Predictors of Avian Species Richness.” *Remote Sensing of Environment* 195:142–152.

Hsiao, Allan. 2021. “Coordination and Commitment in International Climate Action: Evidence from Palm Oil.” Working Paper, Department of Economics, MIT, http://allanhhsiao.com/files/Hsiao_palmoil.pdf.

Joppa, Lucas, and Alexander Pfaff. 2010. “Reassessing the Forest Impacts of Protection: The Challenge of Nonrandom Location and a Corrective Method.” *Annals of the New York Academy of Sciences* 1185:135–149.

Joppa, Lucas N., and Alexander Pfaff. 2009. “High and Far: Biases in the Location of Protected Areas.” *PLOS ONE* 4 (12): 1–6.

Kedward, Katie, Josh Ryan-Collins, and Hugues Chenet. 2023. “Biodiversity Loss and Climate Change Interactions: Financial Stability Implications for Central Banks and Financial Supervisors.” *Climate Policy* 23 (6): 763–781.

Keles, Derya, Alexander Pfaff, and Michael Mascia. 2022. “Does the Selective Erasure of Protected Areas Raise Deforestation in the Brazilian Amazon?” *Journal of the Association of Environmental and Resource Economists* 10 (4): 1121–1147.

Knittel, Christopher R., and Samuel Stolper. 2021. “Machine Learning about Treatment Effect Heterogeneity: The Case of Household Energy Use.” *AEA Papers and Proceedings* 111 (May): 440–44. <https://doi.org/10.1257/pandp.20211090>. <https://www.aeaweb.org/articles?id=10.1257/pandp.20211090>.

Lassiter, Allison. 2022. “Identifying Causal Changes in Landscape Greenness with Very High-Resolution Airborne Multispectral Imagery and a Panel Data Model.” *Urban Forestry & Urban Greening* 67:127380.

Leite Mariante, Gabriel, and Veronica Salazar Restrepo. 2024. “Does Conservation Work in General Equilibrium?” Working Paper, <https://vsalazarr.github.io/jobmarketpaper/>.

Li, Xuecao, Yuyu Zhou, Min Zhao, and Xia Zhao. 2020. “A Harmonized Global Nighttime Light Dataset 1992–2018.” *Scientific Data* 7:168.

Liang, Yuanning, Ivan Rudik, and Eric Zou. 2023. “The Environmental Effects of Economic Production: Evidence from Ecological Observations,” https://static1.squarespace.com/static/56034c20e4b047f1e0c1bfca/t/647a7f669c2ddd18a1e933e7/1685749629464/LRZ_biodiversity-2023-06.pdf.

Maxwell, Sean L., Victor Cazalis, Nigel Dudley, Michael Hoffmann, Ana S.L. Rodrigues, Sue Stolton, Piero Visconti, et al. 2020. “Area-Based Conservation in the Twenty-First Century.” *Nature* 586 (7828): 217–227.

McDonald, Gavin, Jennifer Bone, Christopher Costello, Gabriel Englander, and Jennifer Raynor. 2024. “Global Expansion of Marine Protected Areas and the Redistribution of Fishing Effort.” *PNAS* 121 (29): e2400592121.

Metrick, Andrew, and Martin L. Weitzman. 1998. “Conflicts and Choices in Biodiversity Preservation.” *Journal of Economic Perspectives* 12 (3): 21–34.

Nieto, Sebastian, Pedro Flombaum, and Martin F. Garbulsky. 2015. “Can Temporal and Spatial NDVI Predict Regional Bird-Species Richness?” *Global Ecology and Conservation* 3:729–735.

Oceana Europe. 2022. *As EU Celebrates 30 years of Natura 2000, NGOs Call for These Areas to be Actually ‘Protected’ and for an EU-Wide Trawl Ban in Them.* <https://europe.oceana.org/press-releases/eu-celebrates-30-years-natura-2000-ngos-call-these-areas-be-actually/>.

Pfaff, Alexander, Juan Robalino, Catalina Sandoval, and Diego Herrera. 2015. “Protected Areas’ Impacts on Brazilian Amazon Deforestation: Examining Conservation – Development Interactions to Inform Planning.” *PLOS ONE* 370 (1681): e0129460.

Rambachan, Ashesh, and Jonathan Roth. 2023. “A More Credible Approach to Parallel Trends.” *Review of Economic Studies* 90 (5): 2555–2591.

Reynaert, Mathias, Eduardo Souza-Rodrigues, and Arthur A. van Benthem. 2023. “The Environmental Impacts of Protected Area Policy.” Forthcoming, *Regional Science and Urban Economics*, https://souza-rodrigues.economics.utoronto.ca/wp-content/uploads/RSUE_land-protection-Paper.pdf.

Ribeiro, Ines, Vânia Proença, Pere Serra, Jorge Palma, Cristina Domingo-Marimon, Xavier Pons, and Tiago Domingos. 2019. “Remotely Sensed Indicators and Open-Access Biodiversity Data to Assess Bird Diversity Patterns in Mediterranean Rural Landscapes.” *Scientific Reports* 9 (6826).

Rico-Straffon, Jimena, Zhenhua Wang, Stephanie Panlasigui, Colby J. Loucks, Jennifer Swenson, and Alexander Pfaff. 2023. “Forest Concessions and Eco-Certification in the Peruvian Amazon: Deforestation Impacts of Logging Rights and Restrictions.” *Journal of Environmental Economics and Management* 118:102780.

Robalino, Juan A., Alexander Pfaff, and Laura Villalobos. 2017. “Heterogeneous Local Spillovers from Protected Areas in Costa Rica.” *Journal of the Association of Environmental and Resource Economists* 4 (3): 795–820.

Roth, Jonathan, and Pedro H.C. Sant’Anna. 2023. “When Is Parallel Trends Sensitive to Functional Form?” *Econometrica* 91 (2): 737–747. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA19402>.

Schlenker, Wolfram, and Michael Roberts. 2009. “Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change.” *Proceedings of the National Academy of Sciences* 106 (37): 15594–15598.

Sims, Katharine R.E. 2010. “Conservation and Development: Evidence from Thai Protected Areas.” *Journal of Environmental Economics and Management* 60 (2): 94–114.

Sims, Katharine R.E., and Jennifer M. Alix-Garcia. 2017. “Parks versus PES: Evaluating Direct and Incentive-Based Land Conservation in Mexico.” *Journal of Environmental Economics and Management* 86:8–28.

Souza-Rodrigues, Eduardo. 2019. “Deforestation in the Amazon: A Unified Framework for Estimation and Policy Analysis.” *Review of Economic Studies* 86 (6): 2713–2744.

Sun, Liyang, and Sarah Abraham. 2021. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.” Themed Issue: Treatment Effect 1, *Journal of Econometrics* 225 (2): 175–199.

Taylor, M. Scott, and Rolf Weder. 2024. “On the Economics of Extinction and Possible Mass Extinctions.” *Journal of Economic Perspectives* 38, no. 3 (August): 237–59. <https://doi.org/10.1257/jep.38.3.237>. <https://www.aeaweb.org/articles?id=10.1257/jep.38.3.237>.

The Council of the European Communities. 1992. *Council Directive 92/43/EEC of 21 May 1992 on the conservation of natural habitats and of wild fauna and flora*. Accessed: 2024-05-22. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A31992L0043>.

_____. 1997. *Council Directive 79/409/EEC of 2 April 1979 on the conservation of wild birds*. Accessed: 2024-05-22. <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A31979L0409>.

Torchiana, Adrian L., Ted Rosenbaum, Paul T. Scott, and Eduardo Souza-Rodrigues. 2023. “Improving Estimates of Transitions from Satellite Data: A Hidden Markov Model Approach.” Forthcoming, *The Review of Economics and Statistics*, https://www.ptscott.com/papers/hmm_error_correction.pdf.

Turner, Matthew A., Andrew Haughwout, and Wilbert van der Klaauw. 2014. “Land Use Regulation and Welfare.” *Econometrica* 82 (4): 1341–1403.

Villasenor-Derbez, Juan Carlos, Christopher Costello, and Andrew J. Plantinga. 2024. “A Market for 30x30 in the Ocean.” *Science* 384 (6701): 1177–1179.

Wager, Stefan, and Susan Athey. 2018. “Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests.” *Journal of the American Statistical Association* 113 (523): 1228–1242.

Weitzman, Martin L. 1998. “The Noah’s Ark Problem.” *Econometrica* 66 (6): 1279–1298.

Winkler, Klaus, Richard Fuchs, Mark Rounsevell, et al. 2021. “Global Land Use Changes Are Four Times Greater than Previously Estimated.” *Nature Communications* 12:2501. <https://doi.org/10.1038/s41467-021-22702-2>.

Wolf, Christopher, Taal Levi, William J. Ripple, Diego A. Zarrate-Charry, and Matthew G. Betts. 2021. “A Forest Loss Report Card for the World’s Protected Areas.” *Nature Ecology & Evolution* 5:520–529.

Appendix

A Data

In this data appendix, we describe how we consolidate a variety of publicly available data sources to create country-level data files containing outcome and control variables. Table A.1 lists all data sources.

Table A.1: Data sources

	Name	Description	Data source
[1]	CDDA	Location and foundation date of PAs	Link to source
[2]	Eurostat	National boundaries	Link to source
[3]	EEA	Shorelines of Europe	Link to source
[4]	Landsat 5, 7, 8	NDVI/Greenness	Link to source
[5]	EEA	Biogeographical regions of Europe	Link to source
[6]	EOBS (Cornes et al. 2018)	European climate surface observations	Link to source
[7]	ESDAC (Ballabio et al. 2019)	Soils and topography	Link to source
[8]	ESDAC (Günther et al. 2014)	Climate-physical regions	Link to source
[9]	SEDAC (Warszawski et al. 2017)	Gridded population data	Link to source
[10]	ECMWF	Annual total precipitation	Link to source
[11]	HILDA (Winkler et al. 2020)	Discrete land-use data	Link to source
[12]	World Clim (Fick and Hijmans 2017)	Solar radiation and precipitation	Link to source
[13]	Li et al. 2020a, Li et al. 2020b	Harmonized global night time lights	Link to source
[14]	BioTIME (Dornelas et al. 2018)	Species count data	Link to source
[15]	ESDAC (Panagos et al. 2015)	Slope steepness and elevation	Link to source
[16]	ESDAC (Tóth and Hermann 2016)	Soil suitability	Link to source

A.1 Creation of grids

We define a unit of observation as a square grid cell. Grid cells divide geographic areas into evenly spaced areas with corners given by latitude and longitude coordinates. The grids are constant across time. For our vegetation greenness sample, the grid cell is 300 meters by 300 meters in resolution. We also generate a 1km by 1km grid for the nightlights analysis.

Grids are spatially joined with vector data that are spatially explicit (e.g., data that come in the form of a shapefile or other geodatabase) using ArcGIS. Grids which intersect more than one geometry are assigned the characteristics of the geometry with the largest intersection.

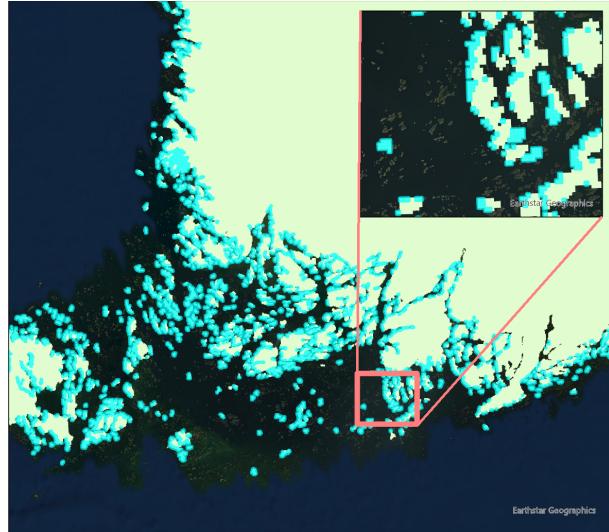
A.2 Bio-geographical regions and climate zones

We add bio-geographical regions and climate-physical zones from the European Environment Agency (EEA) and the European Soil Data Centre (ESDAC). Bio-geographical regions describe distributions and patterns of terrestrial life. The EEA data delineate these bio-geographical regions to show

distinct habitats across Europe. ESDAC produces climate-physical regions.¹ Climate-physical regions are essentially Köppen climate zones with adjustments for high mountains.

By nature, the 300m by 300m grids are not spatially fine enough to pick up on complex geographies. When merging with bio-geographical data and climate-physical regions, the grids miss craggy islands or coastlines of countries such as Norway and Finland. See Figure A.1 for an example. Of the 3,777,950 grids in Finland, fewer than 0.05% were null for bio-geographical regions and 0.5% were null for climate-physical regions. There are similar numbers for Norway. Simpler coastlines have fewer missings. For example, of Poland's 3,473,457 grids, only 114 are null for bio-geographical regions and 3,594 for climate-physical zones. Table A.2 summarizes the different bio-geographical regions and climate zones in Europe.

Figure A.1: Example of grids outside the climate-physical region boundary (Finland)



NOTES: The grids highlighted in cyan are null for climate-physical zones. The country geometry for Finland is more complicated than the climate-physical zone geometry. It captures more islands and coast edges. Grids were generated over these more complex geometries. These grids do not intersect with the more simplified climate-physical regions, so therefore are null.

A.3 Protected sites

We add spatial data on the protected sites from the Common Database on Designated Areas (CDDA), an inventory of European protected areas for 38 nations. The database, maintained by the EEA, includes the location and foundation dates of protected areas established as early as 1800.

The foundation dates of 1,447 designated areas in France are missing. We manually search for the dates of the 127 largest CDDAs with missing foundation dates. We list the dates here with a URL source and notes. The remaining CDDAs with missing foundation dates did not return search

1. The climate-physical regions are based on an intersection of Köppen climate zones with NORDREGIO mountain classification deduced from GTOPO30 information.

Table A.2: Distribution of bio-geographical regions and climate zones in Europe

Bio-geographical region	Land coverage percentage
Continental	28
Mediterranean	18
Boreal	18
Atlantic	17
Alpine	13
Pannonic	3
Arctic	2
Steppic	1
Black Sea	0.2
Climate zone	Land coverage percentage
Cold climate, warm summer	29
Temperate climate without dry period	22
Arid/temperate climate	14
Cold climate, cold summer	13
Polar/cold climate	12
Arid/temperate climate, dry summer	9
Coastal area	2

results. We mainly retrieve the missing foundation dates from site management documents, the conservation pages of provinces, press articles of site establishment or purchase, on the "reserves naturelles" directory, on EEA site factsheets, and tourism pages for CDDAs with recreational and educational uses. The links provide examples to each of these different types of sources. Modern CDDAs may have been the result of multiple staggered land acquisitions rather than a single act of protection. We assign treatment year based on the largest additional acquisition. As a tiebreaker, we assign the average of the purchase dates. See the notes here for detail.²

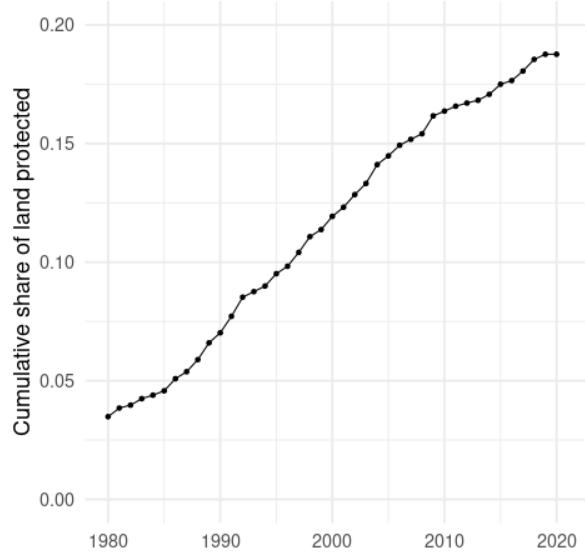
Many of the protected area polygons provided by the EEA have topology errors. Self-intersections are common in the CDDA dataset. These are corrected in ArcGIS.

After these additions and corrections, we relate the CDDA information to the grids via a spatial join in ArcGIS. For each CDDA we add the foundation year and unique CDDA ID so that we can match non-spatial information of the CDDA below. We will also calculate the area of overlap of each CDDA with the grid(s) it covers. Some grids do not fall entirely within CDDAs. Such partial coverage of grids is important when determining which grids are protected vs. treated. We define any plot that overlaps a CDDA as treated.

Our large data collection effort gives us the unique opportunity to collate novel descriptive statistics, both about vegetation growth and conservation policy. Our data in Figure A.2 show

2. One example is Les Pelouses de Blere. In 2003, 14.34 hectares were acquired by Le Conservatoire d'Espaces Naturels Centre-Val de Loire. In 2005, the municipality gave the conservatory 63.58 hectares to manage. The second land acquisition was more than four times greater than the original land acquisition. Because the second land acquisition was larger than the first, we chose the date of the second acquisition for the foundation year. Another site, Les Fiches Des Parterres, was acquired "par le Conservatoire de 22.87 ha entre 1995 et 1999" For this site, we chose 1997 for the foundation date.

Figure A.2: Fraction of landmass in the study area which was protected in or before a given year



NOTES: Data aggregate up from the 300 meter grid-cell level. Any 300 meter grid cell with non-zero overlap of a protected area is defined as treated. The y-axis thus reports the percent of grid cells which overlap any protected area founded on or after the year reported on the x-axis.

that in the study period, by 2020, around 20% of European landmass was protected (as of 2023, this percentage had increased to 26%). Further, relative to the start of our study period in 1985, this is a 300% increase in protected landmass. Our sample period thus captures the epoch with a large percentage point increase in protected area growth.

There is variation in protection regimes. Some areas restrict all human access, while others allow industrial and agricultural activities. While we do not observe the realized level of enforcement, the IUCN categorizes protected areas by their relative strictness. For example, categories Ia and Ib indicate the strictest level of protection, where due to wildlife preservation human activity is strictly limited to either indigenous communities with prior claims to that land or approved research activities. Table A.3 shows that these strict reserves comprise around 8.7% of treated land mass and 7.7% of protected areas. 47% of protected areas by count are habitat and species management areas, though these are a small share of protected area. 14% of parks are protected landscapes, 12% are nature reserves, national parks, and natural monuments, and 26% lack an IUCN category. Species management areas (IUCN category IV) are designed to encourage the propagation of a particular species in the region but are not necessarily as limited in their economic uses. Notably, areas which are protected but are not assigned an IUCN category are on average greener than the areas with the strictest protection.

Table A.3: Average greenness in 1985 by IUCN protected area category

Code	IUCN category	Percent land	Percent CDDAs	Mean greenness
Ia, Ib	Strict nature reserve or wilderness area	8.58	7.80	59.19
II	National park	9.54	0.56	53.03
III	Natural monument	0.68	4.34	51.92
IV	Species management area	15.57	46.70	53.98
V	Protected landscape	44.87	13.65	61.98
VI	Protected area with sustainable use of natural resources	2.85	1.45	57.48
	No IUCN	17.92	25.50	63.30
Ia, Ib, II	Total strict	18.12	8.36	55.95

NOTES: Table summarizes the strictness of protection of CDDAs. Percent land shows the percent of total protected areas. Percent CDDAs shows the share of total counts of protected areas. Sample restricted to terrestrial protected areas with non-missing greenness. Excludes areas with missing (NA) IUCN data. Mean greenness is weighted by area.

A.4 Raster-derived variables

Raster data employ a matrix-based structure, where each cell (or pixel) in the matrix stores a value representing a particular attribute (such as NDVI, elevation, or rainfall). To relate raster data to our grids, we use the `exact_extract()` function in the `exact-extractr` R package to efficiently relate raster data to polygons. Table A.4 gives a description of each of the raster-derived variables added to the country-level grids data. The table provides a brief description of each variable, the units of the data, the spatial resolution, the frequency of the time series for time-series variables (annual, biennial, every 5 years, etc), and the source of the data.

Table A.4: Table of variables which come from a raster format

Variable name	Units	Spatial resolution	Time step, years	Data reference year(s)	URL
Climate					
Long-run precipitation	mm/year	0.5 degrees		average, 1970-2000	WorldClim
Growing season length	days	0.1 degree	1	1985-2019	Copernicus
Heating degree days	deg C	0.1 degree	1	1985-2019	Copernicus
Rainfall	meters	11 km	1	1985-2019	Google Earth Engine
Land characteristics data					
Topsoil potassium concentration	g/kg	500 meters		2019	ESDAC Website
Topsoil nitrogen concentration	mg/kg	500 meters		2019	ESDAC Website
Topsoil phosphorus concentration	mg/kg	500 meters		2019	ESDAC Website
Terrain measures					
Slope steepness	combined LS-factor	100 meters		2015	ESDAC Website
Elevation	meters	100 meters		2015	ESDAC website
Slope angle index	0-8	200 meters		2018	ESDAC Website
Soil suitability index	0-4	1000 meters		2016	ESDAC Website
Economic value measures					
High-value farmland indicator	binary	100 meters		2015	EEA Europa
Distance to shoreline	meters	300 meters		2017	EEA Europa
Population density	count per km ²	30 arc sec	5	1985-2019	SEDAC Website
Outcome measures					
Greenness, LANDSAT-5	index	300 meters	2	1985-2013	Google Earth Engine
Greenness, LANDSAT-8	percent	300 meters	2	2013-2019	Google Earth Engine

Here we provide further information on the raster data that we use in our study.

Climate data. We obtain both a long-term average precipitation and a time-varying measure of precipitation for each grid. We collect monthly long-term averages (long-term mean 1981-2010) of total rainfall from WorldClim. The month's values are an average precipitation for that month from 1981-2010. We average the monthly data to get a total annual figure.

We generate time-varying precipitation from the total precipitation band in the ERA5-Land Monthly Averaged - ECMWF Climate Reanalysis dataset. The climate measures in ERA-5 have a resolution of 11,132 meters. The total precipitation band is the depth of monthly precipitation in meters. ERA-5 total precipitation captures most precipitation but does not include fog, dew, or precipitation which evaporates before reaching the earth.

We use annual total precipitation, rather than summer precipitation that would coincide with our NDVI imagery, because of the importance of year-round rainfall for summer vegetation growth. Pasho et al. (2012) found that larger tree-ring width—an indication that the tree grew more during that year—is related to the amount of autumn and winter rain that had recharged the soils. Dannenberg, Wise, and Smith (2019) and Vieira, Nabais, and Campelo (2021) found that trees in the United States had decreased radial growth and higher mortality risk when winter and summer precipitation were lower. These two articles demonstrate the importance of winter rainfall in vegetation health and justify our use of a yearly total rainfall measure.

To obtain the time-varying climate data from “cold indices” on E-OBS indices, we select “annual” in growing season length and “annual” in heating degree days.

Land characteristics data. The chemical and physical land characteristics are sourced from ESDAC. ESDAC conducted a large survey with approximately 20,000 topsoil samples of soils in Europe to produce a coherent pan-European physical and chemical topsoil database, which can serve as baseline information for an EU wide harmonized soil monitoring. We extract nitrogen, potassium, and phosphorus levels from the LUCAS 2009/2012 topsoil database. We extract soil biomass productivity variables from the EEA 2006 classification. The soil suitability score was created in 2016. We also extract elevation, slope angle, and steepness of slopes from ESDAC.

Economic value measures. The Center for International Earth Science Information Network (CIESIN) of NASA’s Socioeconomic Data and Applications Center (SEDAC) provides gridded population density rasters. CIESIN estimates population density every 5 years to be consistent with national censuses. These numbers are scaled to match UN country-level totals. The data are available at 30 arc sec (1km x 1km) spatial resolution, slightly coarser than the grids. We interpolate the population density data linearly across time.

We compute the distance of each grid to the closest shoreline from polyline data available from the EEA. We compute the Euclidean Distance in ArcGIS at 300 meter pixel resolution and store the data as a raster.

The EEA has created a binary image of high nature value farmland (HNVF). This HNVF measure indicates the potential biodiversity value of existing farms. A value of 1 represents farmland

of high nature value and 0 represents low nature value farmland. We use the data as a proxy of the counterfactual value of agricultural land.

Solar radiation data are available from WorldClim for the period 1950-2020 and is based on ground-based on-site observations and a CERES global radiation satellite product. The data report average monthly solar radiation levels. From these data, we compute an annual average solar radiation raster.

A.5 NDVI

A.5.1 Background

NDVI (normalized difference vegetation index) is an index with values between -1 and 1 representing the level of “greenness” of land cover. It is a well-established index for vegetation monitoring, indicated level of greenness, and plant health. Negative values of the index correspond to water. Low values (0.1 to 0.2) correspond to barren areas, settlements, snow and clouds. Values between 0.2 and 1 correspond to vegetation. NDVI is the ratio between the red (R) and near infrared (NIR) bands:

$$NDVI = (NIR - R) / (NIR + R)$$

We use Google Earth Engine (GEE) to transform Landsat surface reflectance imagery to NDVI and export the NDVI image. Landsat is a satellite imagery program for the entire earth. The early versions, Landsat 1-4, are very similar, are not of sufficient quality, and do not correspond to later versions. Landsat 5-8 have much higher resolution and contain the necessary visual information to capture NDVI consistently across each sensor. We, therefore, use the Landsat 5 data starting in 1985.

We use two-year periods for obtaining NDVI images. Annual data suffer greater missing data issues related to cloud cover. Because accurate NDVI measurement requires leaf-on conditions, we limit the sample of images to summer months. We center our search for images around July as it is the height of summer greenness (Peled et al. 2010; Van Oijen et al. 2014). Ideally, we would produce a single cloudless image using only a two week period—corresponding to the complete earth imaging time of 16 days of the Landsat satellites—in July. However, clouds are often present, requiring that the range of months to search for images be extended to either June through August or, in exceptional circumstances, May through September, until the percentage of missing images falls below 5 percent. We discuss the details of the procedure for producing a cloud-free mosaic in our data repository.

When we compute average NDVI by the discrete land use categories in HILDA, we find that forests are the greenest land use category with an average NDVI of 63 in 1985. More than 25% of grids classified as forest have an NDVI above 73. Grassland has a somewhat lower NDVI that is mostly between 43 and 69. Cropland has the lowest NDVI, with an average value of 49 in 1985. NDVI variation thus matches our priors about land use, but we also find differences between biomes

and climate regions to be important. Forests in arid regions of Spain are characterized by shrubs and greater spacing of vegetation and have lower NDVI. The 25th percentile of the distribution of NDVI on forested land is 58 (still higher than most cropland, but lower than the median of grassland at 59). We estimate treatment effects separately by country and include climate, weather, and soil controls to ensure valid treatment-control pairs.

A.5.2 Landsat 7 scan line correction

The Landsat 7 satellite was launched on April 15, 1999. It collected quality images until May 31, 2003 when the Scan Line Corrector (SLC) in the Enhanced Thematic Mapper (ETM+) instrument failed. Landsat 7 images after this date are not usable for our analysis. The 1999-2003 Landsat 7 operation period overlaps with Landsat 5 (April 1, 1984 - June 5, 2013). During this time frame, we use Landsat 7 images and fill in any missing pixels with Landsat 5 data.

A.5.3 Landsat 8 OLI to Landsat 5 ETM+ spectral response correction

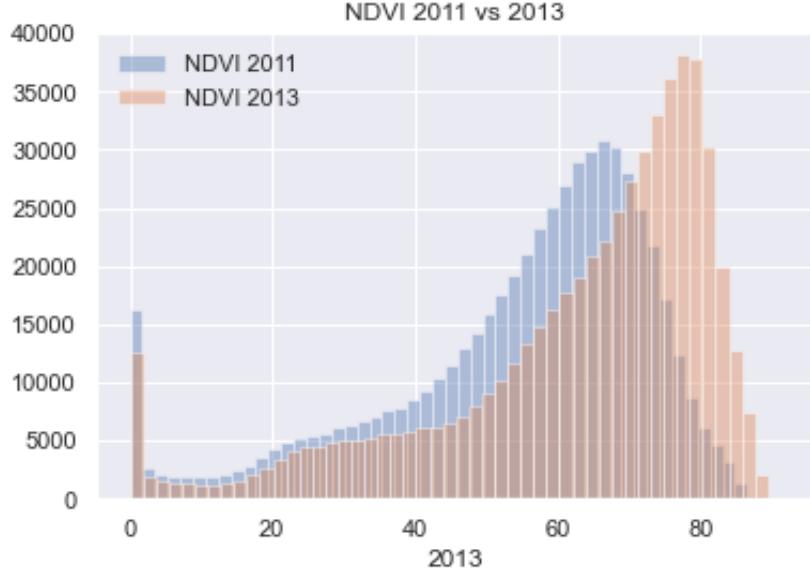
Landsat 8 and Landsat 5/7's sensors are largely comparable; they have the same spatial resolution and 16-day revisit time. However, the spectral response functions of the two sensors differ. Landsat 8's Operational Land Imager (OLI) is an improvement on Landsat 5/7's Enhanced Thematic Mapper Plus (ETM+). However, we need to correct the spectral response function of Landsat 8 to make NDVI directly comparable throughout the panel.

The differences in the spectral response functions of the two sensors lead to a “brighter” NDVI for Landsat 8 imagery than for Landsat 5. This is clear from the histogram of NDVI for all European grids in 2011 (Landsat 5 ETM+) and in 2013 (Landsat 8 OLI) in Figure A.3. There is a noticeable skewing towards higher values of the Landsat 8 NDVI in a brief period of just two years. Because the differences in NDVI in Figure A.3 appear to come largely from measurement, we use a correction to harmonize the NDVI measure across each satellite. Roy et al. (2016) provide the coefficients to apply to each band to harmonize Landsat 8 to Landsat 5. We chose to harmonize OLI to ETM+ because the thematic mapper makes up the majority of our NDVI imagery (1985 to 2011, 13 two-year images).

A.6 Discrete land-use data

The HILDA dataset provides over 100 years of land-cover maps at a 1 kilometer grid cell resolution. As we discuss in Section 4, it has severe limitations as a main outcome variable for measurement of biodiversity. However, HILDA does provide descriptive insights into long-run land-cover trends since 1900. HILDA omits a subset of countries which are included in our analysis sample. These are Albania, North Macedonia, Montenegro, Croatia, Bosnia and Herzegovina, Serbia, Norway, and Iceland. We omit HILDA reporting on several smaller nations which are not in our main analysis sample: Andorra, Monaco, Jersey, Guernsey, Isle of Man, and Faroe.

Figure A.3: Spectral response affecting NDVI 2011-2013



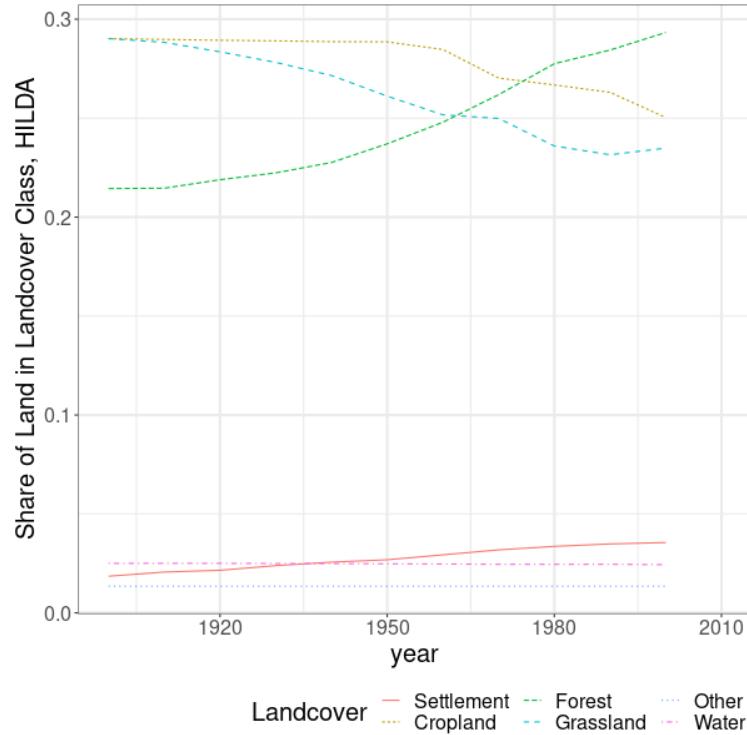
NOTES: The histogram shows a strong skewing towards higher NDVI values from 2011 to 2013 across all of Europe. This is not attributable to large scale land cover shifts; this is the result of the difference in the spectral response functions of ETM+ and OLI. Between 2011 and 2013, Landsat 5 was decommissioned and Landsat 8 became operational.

HILDA categorizes land into one of six discrete classes: cropland, forest, grassland, other (such as mountains or barren land surface), settlement, and water. In Figure A.4 we plot the share of each land use in the HILDA data. The data are at a decadal frequency since 1900. The data confirm a century-long growth in forest cover across the EU. HILDA reports a forest share in the EU of around 22% in 1900 and a 29% forest share in 2010.

In Table A.5, we report a transition matrix over the full breadth of the HILDA data. Looking at the diagonal entries, it is clear that cropland and grassland experienced the greatest land-cover shifts in percentage terms: 30% and 40% point of the 1900 landcover has transitioned to other land uses. The chief beneficiary of both appears to have been forest, though there is non-negligible transition between grassland and cropland. This latter transition can be both a true return of cropland to natural use (or vice versa), or can encompass measurement error as cropland and grassland are less readily discerned by classifiers, or finally can also indicate conversion between crop and pasture (which is not a dimension of biodiversity we are particularly interested in). However, overall, Europe became more forested over this long time horizon: 10% of cropland in 1900 appears to be forest in 2010, and 27% of grassland in 1900 is forest by 2010. This outweighs the 9% conversion from forests to grasslands. Settled areas also transition to forest—some subset of land returned to natural use.

Finally, we re-create Table 1 in two other time periods in Table A.6. The top panel shows a balance table across land deforested and not deforested 1990-2010, and the bottom does the same for land across 2000-2010. Compared to full-sample deforestation, land deforested more recently

Figure A.4: Land share in the EU by decade and land use category



NOTES: The HILDA landcover data classify land into one of six land areas in each decade from 1900 to 2010. We omit a small percentage of areas which are classified as NA or missing values.

Table A.5: Discrete land use transitions for the entirety of Europe between 1900-2010

Land use in 1900 / 2010		Land-use transition probabilities					
		Cropland	Forest	Grassland	Other	Settlement	Water
Cropland	% row	68.5	9.6	17.7	0.0	4.1	0.0
Forest	% row	0.8	89.7	8.8	0.0	0.8	0.0
Grassland	% row	11.7	27.2	58.8	0.0	2.3	0.0
Other	% row	0.0	0.0	0.0	100.0	0.0	0.0
Settlement	% row	7.4	5.0	6.8	0.0	80.8	0.0
Water	% row	1.2	0.2	0.9	0.0	0.4	97.2
Land-use shares							
Land use in 2010		Cropland	Forest	Grassland	Other	Settlement	Water
Total (2010)	% row	27.7	35.2	28.4	1.6	4.2	2.9

NOTES: Table reports land use transitions relative to base year of 1900 in 2010. Transitions are defined based on the HILDA landcover data, which classifies land into one of 6 land areas in each decade from 1900 to 2010. We omit a small percentage of areas which are classified as NA or missing values.

tends to be on more populated land, corroborated by a slightly higher measure of luminosity in the nightlights data. Deforested land is less likely to be in protected areas than the full-sample case, Deforested land also has more potassium in the soil: in Europe, this is likely to be driven by agricultural or urban water runoff.

B Econometrics

B.1 Estimation details: Staggered difference-in-differences design

We use the ‘did’ package in the R programming language.³ We select all observations for which there are no missing variables and keep grids that cross the greenness threshold of NDVI 40 at least once in the sample. Plots that fully or partially overlap a CDDA are considered treated units for our results on NDVI. We construct our estimator in the syntax of the ‘att_gt’ command. We treat the data as a panel, using the rounded year of protection as the cohort definition. The command constructs bootstrapped confidence intervals. We keep all computational defaults regarding bootstrap size and sampling procedures. Our estimator is calculated separately for each country.

For large countries, the did package is too slow to produce estimates in a reasonable timeframe. Therefore, we conduct a subsampling procedure of the estimator, where we draw a 5% stratified subsample at the cohort-identifier level. While not strictly necessary for the estimator, we also stratify at the biogeographical region level to maintain a consistent set of matching covariates relative to small countries. We do this for Finland, France, Germany, Italy, Poland, Romania, Spain, Sweden, and the United Kingdom. The subsampling procedure mirrors the bootstrap of the estimator by ensuring that every cohort and event time is represented in each subsample, so that the full set of θ_{gt} is estimable for each subsample. We subsample with replacement and define new plot identifiers (e.g., a new grid cell-level ID) to ensure the command runs with potential duplicates in the data. We sample 50 clustered 5% draws for each country, bootstrapping the entire matrix θ_{cgt} , θ_{ce} , θ_{cg} , and θ_c as well, obtaining bootstrapped standard errors. Because we already find the estimator to be slow with this bootstrapping scheme, we do not apply the Callaway and Sant’Anna wild bootstrap to compute uniform confidence bands. These could, in principle, widen our standard errors as they account for covariance between dynamic treatment effects, but given the level of precision at which our treatment effects are computed, we do not anticipate these change the interpretation of our results.

Estimates for Cyprus, Malta, and Liechtenstein are not computed due to missing data in the time-varying weather patterns. There are significant missing shares of time-invariant variables for Switzerland, Luxembourg, and Montenegro. A subset of countries have a very high ($> 99\%$) share of one or more matching variables: these are matched on the remaining covariates. Countries lacking slope sleepness and soil suitability measures are Switzerland, Serbia, Bosnia, Albania, and

3. See: <https://www.rdocumentation.org/packages/did/versions/2.1.2>.

Table A.6: Average of key variables among land classified as forest in 1900 (top) and 2000 (bottom) by whether that land was deforested

	Not deforested, 1990-2010		Deforested, 1990-2010	
	Mean	Std. dev.	Mean	Std. dev.
Nightlights in 2010	6.6	9.8	11.9	15.3
Percent of grid protected	24.1	46.6	18.2	41.0
Population density in 2000	50.8	169.2	179.0	659.4
Crop suitability	6.1	1.4	6.3	1.4
Forest suitability	3.6	1.5	4.0	1.6
Grassland suitability	5.4	1.5	5.6	1.5
Slope steepness	2.2	3.0	2.0	2.6
Solar radiance	11.1	2.5	11.4	2.2
Precipitation	759.2	271.8	763.8	243.1
Potassium	157.5	67.1	164.2	61.3
Nitrogen	2.1	0.9	2.2	0.8
Growing season length	236.4	72.5	262.1	61.1

	Not deforested, 2000-2010		Deforested, 2000-2010	
	Mean	Std. Dev.	Mean	Std. Dev.
Nightlights in 2010	6.8	10.0	11.8	14.9
Percent of grid protected	24.0	46.6	13.6	36.3
Population density in 2000	52.0	169.3	110.5	384.2
Crop suitability	6.1	1.4	5.9	1.4
Forest suitability	3.6	1.5	3.8	1.5
Grassland suitability	5.4	1.5	5.4	1.6
Slope steepness	2.2	2.9	1.1	1.8
Solar radiance	11.1	2.5	12.0	29.2
Precipitation	757.7	270.8	652.6	168.5
Potassium	159.0	69.3	192.2	100.1
Nitrogen	2.1	0.9	1.8	0.7
Growing season length	237.9	72.9	265.9	77.3

NOTES: Table presents balance of several key variables against an indicator based on the discrete land-use classifications provided by the HILDA data. Top panel: not deforested indicates 1 km × 1 km grid cells which were coded as forest in 1990 and were still coded as forest in 2010. Deforested indicates areas which were forest in 1990 but were coded as any other category in 2010. Bottom panel does the same for grid cells coded as forest in 2000. Percent of grid protected indicates the percent of the 1 square kilometer grid which contains protected areas, regardless of their designation year and CDDA designation. Means of time-varying variables are calculated in a specific cross-section, as indicated in the table. Units for all variables are indicated in Appendix Table A.4.

Montenegro. Malta lacks rainfall data on 60% of its landmass. We report results for estimators that omit time-invariant variables from the matching procedure.

The ‘did’ package computes clustered standard error at default by clustering at the plot level, allowing for serial correlation in the error terms at the observation level. Additionally, we cluster at the protected area unit level (CDDA number) to account for potential spatial correlation. This two-way clustering makes standard errors robust to correlation in the time and cross-sectional (spatial) dimensions.

We maintain the same matching variables and technical specifications for our nightlights analysis. There are two key exceptions. Rather than matching on pre-period greenness, we match on pre-period nightlights. We require a single pre-period of 1992 as the nightlights data do not cover the same breadth of data as the greenness data. Because nightlight imagery is taken at night, it is less frequently missing than greenness and we obtain an annual panel. We therefore have treated cohorts from 1993 through 2019 (27 total potential values of g). Our nightlights data are also collected at a lower spatial resolution (1 km grid cells, see Appendix A), so we reconsider our treatment definition. As our second deviation from the greenness strategy, treated grid cells must overlap at least 50% (in area terms) with a protected area.

B.2 Estimation details: Conditional average treatment effects (CATE)

We use the CATE estimator provided in the R package ‘grf’, generalized random forest.⁴ Our data match the package’s description of a medium-sized dataset. We follow the package documentation to construct a suitable estimator for our dataset. Computationally, we target an accurate calibration that passes the test in Chernozhukov et al. (2024) and precisely separates the top and bottom quartiles of estimated CATEs. This latter point ensures heterogeneity is rejected with precision if no significant heterogeneity exists.

We apply the Wager and Athey (2018) estimator to a sample of plots. Samples are stratified at the country-foundation year-biogeoregion-climate zone level. Samples are taken if and only if a stratum has at least 30 units to allow for inference on CATES. In addition to never-treated land, our sample consists of treated land which enters protection after 1991 to ensure sufficient variation in matching variables. Each sample consists of two periods: a pre-period $g - 1$ and 2019 (the universal post-period as it is the last period in our data).

We supply all potential exogenous variables to compute the forest. We include rich interactions and third-order splines in continuous variables as well as differences in slopes across categorical variables. Exogenous variables are used for matching classification trees and heterogeneity regression trees, but each tree is trained on separate sub-samples (“honesty”, in the language of Wager and Athey (2018)). We omit missing observations. Categorical variables are converted to indicator variables. Soil classes, continuously graded from 0 to 4, are rounded to increments of 0.1.

4. See: grf-labs.github.io/grf/

With a matrix of observables, an outcome vector of greenness, and a treatment indicator, we split the data into a training and a test dataset. We retain 30% of the data for training, and the remainder for testing. Training data are used to calibrate the random forest. The algorithm statistically tests the calibrated random forest on a separate, withheld test dataset.

Propensity scores and the outcome regression adjustment are calibrated using package default parameters. In our final step, we change some computational parameters from package defaults to maintain accuracy: we grow 10,000 trees per cohort, with a larger sample fraction of 0.5 to avoid cutting the dataset too finely, and we require 50 observations for the random forest to generate a new leaf. This last step prevents the forest from slicing the data too finely, which can lead to overfitting.

We report the results of the Chernozhukov et al. (2024) test of meaningful heterogeneity using the ‘test_calibration’ command. For all computed conditional average treatment effects, we ensure that they are estimated on the subsample of data with overlap (the ‘target.sample’ parameter). Reported values are doubly robust. CATEs are thus calculated using the ‘average_treatment_effect’ command, subsetting to the portion of data of interest.

C Additional tables and figures

C.1 Propensity score weighting and balance tables

The main identification challenge we face is the nonrandom selection of land for treatment. While the aggregate picture in the EU suggests (Table A.7) that there may be overlap between treated and never-treated units, zooming into countries suggests significant imbalances. We demonstrate this imbalance in Table A.8 for the case of France (chosen because it is relatively large and has highly-heterogeneous protected areas). Protected areas are in less populated areas (lower population density), and steeper and less accessible areas (elevation and slope). To obtain balance, the algorithm of Callaway and Sant’Anna (2021) applies cohort-specific propensity score weighting. This is important in our setting as we expect that later-treated cohorts consist of different land than early-treated cohorts. In each cohort, the procedure develops a propensity weight based on variables that appear in the vegetation greenness production function: elevation, slope steepness, soil quality, solar radiance, growing season length, heating degree days, rainfall, and starting greenness. In each cohort, the resulting estimator weighs control units by their similarity on vegetation-relevant observables to treated protected areas. Table A.9 shows an example of how the inverse propensity weighting shrinks the difference in means of variables between treated and untreated units in the year before treatment, in France and for the cohort treated in 2005.

C.2 Treatment effect estimates: NDVI

Here we report the results of the Callaway and Sant’Anna estimators in more detail. Table A.10 lists the average treatment effect aggregated over both event time and cohort for each country in

Table A.7: Balance table for the entire European Union

	Never-treated		Treated	
	N	Mean (Std. dev.)	N	Mean (Std. dev.)
Greenness	41,964,510	56.77 (17.57)	7,958,603	59.49 (16.97)
Growing season length	41,964,510	247.98 (66.12)	7,958,603	235.23 (66.69)
Heating degree days	41,964,510	3,035.79 (1,667.98)	7,958,603	3,331.62 (1,679.45)
High-value farmland fraction	41,964,510	0.17 (0.33)	7,958,603	0.21 (0.36)
Population density	41,964,510	109.66 (446.32)	7,958,603	102.61 (310.24)
Rainfall (mm)	41,964,510	740.51 (280.68)	7,958,603	771.78 (304.12)
Slope angle	41,964,510	2.13 (1.52)	7,958,603	2.53 (1.80)
Elevation (m)	41,964,510	348.07 (378.37)	7,958,603	449.48 (488.35)
Solar radiance	41,964,510	11,839.17 (2,537.39)	7,958,603	11,434.40 (2,552.27)
Distance to shore (m)	41,964,510	151,739.42 (136,316.84)	7,958,603	156,095.52 (143,080.76)
Soil phosphorus (mg/kg)	41,279,623	26.72 (13.12)	7,928,447	26.66 (13.52)
Soil nitrogen (mg/kg)	41,279,623	2.08 (0.94)	7,928,447	2.28 (1.07)
Soil potassium (mg/kg)	41,279,623	198.79 (95.35)	7,928,447	173.17 (85.24)

NOTES: Table compares observables for ever- vs. never-treated units. We enforce that land has non-missing values of controls to be included in the table. Plots which were protected before 1990 are trimmed to ensure we can compare several periods' worth of pre-trends. Time-varying variables (heating degree days, rainfall, and growing season length) are averaged over 1985-1989 levels. Statistics are weighted by the number of observations in each country.

Table A.8: Balance table for France

	Never-treated	Treated	Diff. in Means
	N = 4883391	N = 584464	
NDVI, 1991	56.82 (14.41)	61.50 (14.51)	-4.68 (0.02)
NDVI, 1985	60.25 (13.13)	63.75 (14.33)	-3.5 (0.02)
NDVI change, 1985-1989	-0.05 (11.61)	1.13 (10.45)	-1.18 (0.01)
Population density, 1985	111.79 (525.22)	59.36 (189.36)	52.43 (0.34)
Heating degree days, 1985	2,268.47 (661.06)	2,409.66 (712.05)	-141.19 (0.98)
Growing season length, 1985	286.10 (33.73)	280.54 (37.02)	5.56 (0.05)
Soil suitability index	3.66 (0.75)	3.39 (0.89)	0.27 (0)
Slope angle	1.93 (1.36)	2.67 (1.80)	-0.74 (0)
Slope steepness	1.38 (2.25)	2.45 (3.11)	-1.07 (0)
Solar radiance	12,288.29 (1133.66)	12,851.02 (1356.51)	-562.73 (1.85)
High nature-value farmland (%)	12.20 (28.67)	23.82 (36.97)	-11.62 (0.05)

NOTES: Treated data are aggregated across all cohorts. The sample selection procedure for this table is the same as the rest of the paper. Plots which were protected before 1990 are trimmed to ensure we can compare several periods' worth of pre-trends. Plots with a missing covariate are omitted. Time-varying variables are captured in 1989, the last pre-period year in the sample. Standard errors computed assuming independent populations.

Table A.9: Balance table for France, specific to the cohort treated in 2005 after propensity-score matching

	Never-treated	Treated	Diff. in Means	IPW Diff. in Means
	N = 1301749	N = 53017		
Growing season length, 1985	286.10 (33.73)	257.05 (38.39)	29.05 (0.17)	11.17 (0.05)
Heating degree days, 1985	2,268.47 (661.06)	,2966.42 (623.52)	-697.95 (2.77)	-241.99 (0.83)
High nature-value farmland (%)	12.20 (28.67)	31.68 (37.94)	-19.48 (0.17)	-18.75 (0.05)
NDVI change, 1985-1989	-0.05 (11.61)	2.50 (7.46)	-2.55 (0.03)	-3.21 (0.01)
NDVI, 1985	60.25 (13.13)	70.25 (11.76)	-10 (0.05)	0.61 (0.01)
NDVI, 1991	53.89 (16.02)	68.74 (11.66)	-14.85 (0.05)	-2.3 (0.01)
Population density, 1985	111.79 (525.22)	40.59 (121.71)	71.2 (0.7)	16.04 (0.12)
Slope angle	1.93 (1.36)	2.95 (1.64)	-1.02 (0.01)	0.22 (0)
Slope steepness	1.38 (2.25)	2.72 (3.09)	-1.34 (0.01)	0.63 (0)
Soil suitability index	3.66 (0.75)	3.38 (0.89)	0.28 (0)	-0.23 (0)
Solar radiance	12,288.29 (1133.66)	12,964.72 (1242.97)	-676.43 (5.49)	-233.44 (1.39)

NOTES: Balance table calculated in first pre-treatment panel year, 2003. IPW = inverse probability weight matching. Matching variables are long-run precipitation, elevation, solar radiance, slope steepness, a soil suitability index, slope angle, rainfall, heating degree days, and a cubic polynomial in greenness. These variables are taken in the pre-period. Additionally, we included a three-year average and variance of rainfall, greenness, and heating degree days between 1985-1989. The sample here is trimmed for balance: propensity scores lower than 0.01 and higher than 0.99 are removed.

the European Union. We report bootstrapped standard errors using the Callaway and Sant'Anna methodology. We also report the number of unique individual plot identifiers remaining in our data for each country as well as the number of unique protected areas (CDDAs). Some countries are under-represented due to missing data. Others are under-represented because of a lack of sufficient overlap in some of the control variables. Most countries are very precisely estimated.

We also report examples of dynamic treatment effect plots θ_{ce} here. These dynamic effects are most useful as visual tests of the parallel trends assumption. It would be exponentially overwhelming to report visual evidence for all countries; instead, we show two examples (Poland and Spain) that are representative of an overall absence of (trends in) treatment effects.

We start with the results for Poland, a large country with lots of standing old-growth forest. Figure A.5 shows the estimated dynamic treatment effects $\hat{\theta}_{ce}$ for Poland. The treatment effect for the base period -2 is fairly close to 0, suggesting that we have decent claim to a conditional parallel trends assumption. We see that up to 20 years after protection, all treatment effects are within the $[-1,1]$ range.

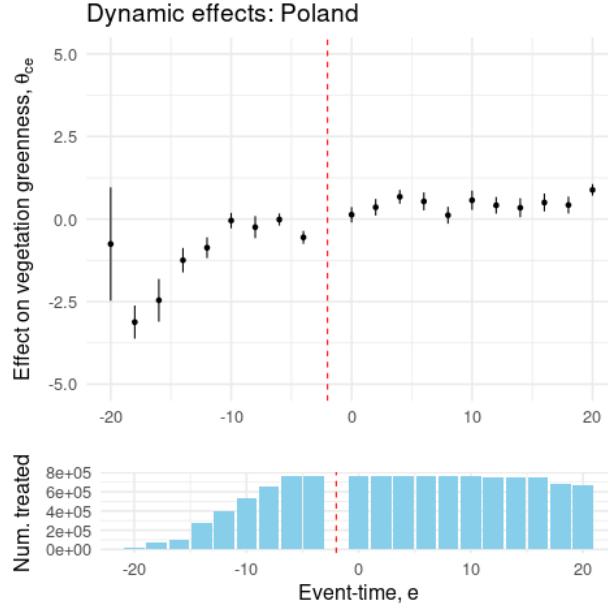
We provide a second example of parallel trends. Spain is in a different bio-geographical region and climate zone than Poland. It is also less densely vegetated. In Figure A.6, the trend up to 20 years prior to protection is flat and close to 0. We see that up to 20 years after protection, treatment effects are around zero, and negative on average. Note that both this warmer and sunnier climate and the densely-forested Poland demonstrate similar conclusions: no evidence for a positive impact on greenness up to 20 years after the date of land protection. This long time horizon shows neither

Table A.10: Estimated average treatment effect on vegetation greenness for each country, aggregated over cohort and event time

	ATT estimate (Std. error)	Observations	Num. CDDAs
CSDR	0.02 (0.07)	98,712,954	74,727
MTWFE	0.49 (0.09)	35,985,945	74,727
TWFE	0.52 (0.09)	36,411,204	74,727
Albania	-0.22 (0.04)	303,589	53
Austria	-0.00 (0.02)	757,581	709
Belgium	0.51 (0.04)	319,704	1,468
Bosnia	-0.65 (0.04)	554,491	31
Bulgaria	-0.25 (0.03)	1,146,230	387
Croatia	1.41 (0.11)	513,541	103
Czech	-0.87 (0.04)	711,577	1,516
Denmark	-0.26 (0.03)	415,875	303
Estonia	1.06 (0.02)	461,819	8,358
Finland	0.19 (0.40)	5,562,252	10,669
France	-0.22 (0.07)	10,520,352	3,006
Germany	0.22 (0.12)	18,593,748	11,128
Greece	-0.11 (0.01)	1,010,673	520
Hungary	0.78 (0.03)	907,934	243
Ireland	-1.53 (0.08)	575,616	172
Italy	0.58 (0.10)	4,111,506	695
Latvia	0.51 (0.05)	634,185	411
Lithuania	-0.21 (0.04)	669,488	397
Luxembourg	-0.75 (0.10)	26,692	79
Macedonia	-0.43 (0.06)	260,411	45
Montenegro	1.72 (0.08)	143,056	13
Netherlands	1.32 (0.06)	353,777	232
Norway	-1.61 (0.29)	1,418,454	821
Poland	0.43 (0.10)	13,794,318	1,005
Portugal	-0.25 (0.03)	904,108	153
Romania	0.31 (0.15)	2,582,370	644
Serbia	-1.45 (0.22)	933,000	245
Slovakia	0.66 (0.08)	480,830	977
Slovenia	1.63 (0.06)	203,836	711
Spain	-0.75 (0.06)	12,907,512	1,575
Sweden	-0.45 (0.09)	11,439,774	12,745
Switzerland	0.81 (0.02)	402,941	8,904
UK	-0.17 (0.26)	5,091,714	6,409

NOTES: Average treatment effect θ_c of conservation on vegetation greenness (an index varying between 0 and 100) in Equation (9) estimated within each country in the European Union using data from 1985-2019 on a biannual basis. Top three rows report the Callaway and Sant'anna, doubly-robust estimator (CSDR), matched two-way fixed effects (MTWFE), and two-way fixed effects without matching (TWFE), respectively. Observations are at a 300 meter resolution. To ensure adequate pre-period variation, treated units are limited to those units protected in or after 1991. “Non-missing, matched grid cells” represents the number of panel observations ($N \times T$) that were identified with non-missing matching variables across foundation years between 1991 and 2019.

Figure A.5: Plot of dynamic treatment effects on vegetation greenness for Poland

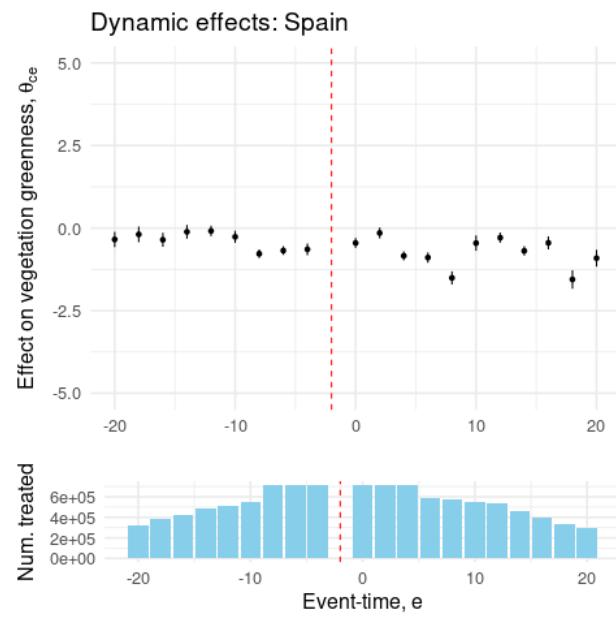


NOTES: Top panel shows estimates and confidence bands for θ_{ce} for Poland. Bottom panel shows the number of treated grid cells for each event time. Event-times are trimmed at $[-20, 20]$. Confidence bands are based on a bootstrapping procedure discussed in Appendix B.

a change in greenness nor evidence of a trend in greenness which might indicate more gradual vegetation growth. If anything, treatment effects are negative.

The overall EU dynamic treatment effects, shown in the main text in Figure 6 (left panel), are presented numerically here in Table A.11. There is no meaningful trend in EU-wide dynamic treatment effects, rejecting that protection has led to long-term vegetation recovery or (re-)growth for the average protected area.

Figure A.6: Plot of dynamic treatment effects on vegetation greenness for Spain



NOTES: Top panel plots estimates and confidence bands for θ_{ce} for Spain. Bottom panel shows the number of treated grid cells for each event time. Event-times are trimmed at $[-20, 20]$. Confidence bands are based on a bootstrapping procedure discussed in Appendix B.

Table A.11: Dynamic treatment effects on vegetation greenness across the EU

	Mean (Std. error)	Num. treated grids	Num. CDDAs
-20	-0.62 (0.20)	2,625,150	39,090
-18	-0.64 (0.16)	3,123,475	46,821
-16	-0.84 (0.18)	3,648,715	51,406
-14	-0.71 (0.13)	4,203,856	54,414
-12	-0.88 (0.13)	4,740,919	57,300
-10	-0.24 (0.12)	5,198,154	60,381
-8	-0.46 (0.12)	5,945,125	63,389
-6	-0.27 (0.16)	6,571,274	66,144
-4	-0.13 (0.18)	6,571,274	68,414
-2	0.00 (0.00)	6,571,274	70,949
0	-0.12 (0.07)	6,570,363	73,310
2	0.03 (0.08)	6,175,759	65,964
4	-0.04 (0.09)	5,840,098	62,277
6	0.27 (0.08)	5,470,259	59,377
8	-0.27 (0.13)	5,324,282	56,093
10	-0.14 (0.12)	5,129,335	49,864
12	0.11 (0.10)	4,840,454	45,746
14	0.19 (0.11)	4,541,966	40,749
16	0.31 (0.15)	3,946,124	34,251
18	-0.05 (0.16)	3,447,799	26,520
20	0.42 (0.12)	2,922,559	21,935

NOTES: Dynamic treatment effect θ_t^{EU} of conservation on vegetation greenness (an index varying between 0 and 100) estimated for the entire European Union. Estimates are aggregated from country-level estimates of $\theta_c(g, t)$. Underlying dataset spans 1985-2019 on a biannual basis. Observations are at a 300 meter resolution. The earliest included foundation year period is 1991, meaning the latest event-time in the sample is 28. Similarly, the last valid foundation year is 2019. Number of treated grids refers to the number of unique 300 square meter grid cells in our sample which were protected at a time such that $e = t - g$ was observed (e.g., for $e = -34$, this will be the number of CDDAs founded in 2019, the last year of the data). To convert to area in square kilometers, multiply the number of grid cells by 0.09.

We also report the overall EU cohort-level effects in Figure 6 (right panel) as Table A.12 here. There are 15 cohorts. The treatment effects again center around zero, with no discernible trend in treatment effects moving from early- to late-treated cohorts.

Table A.12: Cohort-level treatment effects on vegetation greenness across the EU

	Mean (Std. error)	Num. treated grids	Num. CDDAs
1991	0.63 (0.31)	626,149	2,755
1993	-0.72 (0.34)	746,971	3,008
1995	0.60 (0.55)	457,235	3,081
1997	0.14 (0.33)	537,063	2,886
1999	0.47 (0.45)	555,141	3,008
2001	0.07 (0.35)	525,240	4,585
2003	-0.37 (0.52)	498,325	7,731
2005	-0.07 (0.41)	595,842	6,498
2007	0.36 (0.43)	298,488	4,997
2009	-0.73 (0.40)	288,881	4,118
2011	0.22 (0.79)	194,947	6,229
2013	-0.38 (0.86)	145,977	3,284
2015	0.20 (0.83)	369,839	2,900
2017	-0.11 (0.85)	335,661	3,687
2019	-0.14 (1.07)	394,604	7,346

NOTES: Cohort-level treatment effect θ_{cg}^{EU} of conservation on vegetation greenness (an index varying between 0 and 100) estimated for the entire European Union. Estimates are aggregated from country-level estimates of $\theta_c(g, t)$. We sum across all available event-times. Underlying dataset spans 1985-2019 on a biannual basis. Observations are at a 300 meter resolution. Treated units are excluded if they had a foundation year earlier than 1991 to ensure at least 2 periods of parallel trends (3 data observations: 1985, 1987, and 1989).

Finally, Table A.13 presents cohort-level effects aggregated to the country level (rather than at the EU level in Table A.12). We use these country-level aggregations to test for selection of protected areas over time and to compare any potential selection across regimes. The first column indicates the number of cohorts for which treatment effects are calculated in the range 1991-2019. The second column reports a trendline from regressing θ_{cg} , the country-level analog of Equation (12), on cohorts $g - 1991$. We difference out 1991 so changes are interpreted as the effect of being treated two years later. Trends are economically small, suggesting that the change in treatment effect across successive cohorts varies by as little as 0.01% of maximum NDVI. The largest positive trend is in Estonia (0.19). Estonia's trend suggests that on average, plots treated in the last cohorts in 2019 experienced a treatment effect 2.7 units higher than those treated in 1991. The third column constitutes a less parametric approach to estimating trends. These "split-difference" estimators compare average treatment effects θ_{cg} in the back half of the study period to those in the first half. The estimators find similar results with less extrapolation involved. Overall, these results demonstrate that the selection of protected areas does not appear to manifest in significant

trends in cohort-level effects of land protection on greenness.

Two-way fixed effects

For completeness, we report the dynamic effects estimated by a classic matched, two-way fixed effects estimator. The estimating equation includes grid i and time t fixed effects:

$$Y_{it} = \beta D_{it} + \lambda_i + \lambda_t + \epsilon_{it} \quad (\text{A.1})$$

Unlike the Callaway and Sant'Anna estimator, the classic matched two-way fixed effects estimator uses a single matching function rather than cohort-specific matching functions. Thus, the TWFE matching cannot address between-cohort variation in selection.

Figure A.7 plots the event-study coefficients. We make two remarks. First, the two-way fixed effects estimates indicate a long-run pre-trend compared to the Callaway and Sant'Anna methodology. This pre-trend is addressed in the Callaway and Sant'Anna methodology through (1) cohort-specific matching and (2) avoiding forbidden comparisons (Goodman-Bacon 2021). In our setting, absent a treatment effect after protection, we believe (1) is the primary channel through which the conventional TWFE estimator introduces bias. Second, the two-way fixed effects estimates indicate a minimal EU-level treatment effect of protection. Of course, this number is difficult to causally interpret in the presence of a pre-trend and statistical imprecision.

C.3 Treatment effect estimates: nightlights

Nightlight outcomes are measured on a luminosity scale ranging from 0 to 68. Across the EU in 2021, the unconditional average luminosity is 9.9. However, 30% of our sample has 0 luminosity (at a 1-kilometer grid cell resolution). A 0 requires neither indoor nor outdoor electricity: these 0s correspond to clear “undisturbed” land area. Conditional on having positive luminosity, the average luminosity is 14.3. Countries vary in average nightlights from very bright in Belgium (conditional on a positive luminosity, mean of 30.5) and the Netherlands (26.2) to relatively less so in Ireland (7.9) or Bulgaria (9.3).

Classically, long-run changes in nightlights have been associated with long-run GDP growth (Donaldson and Storeygard 2016; Gibson et al. 2021). When regressing GDP growth on changes in nightlights, Henderson, Storeygard, and Weil (2012) report that a 1% change in nightlights is associated with a 0.3% increase in GDP growth over a 20-year panel. , but this approach has also received criticism about unstable relationships and data inconsistencies (Chen and Nordhaus 2011; Bickenbach et al. 2016). In this paper, we do not interpret nightlights as a measure of economics activity; we are interested in its direct measurement of human presence.

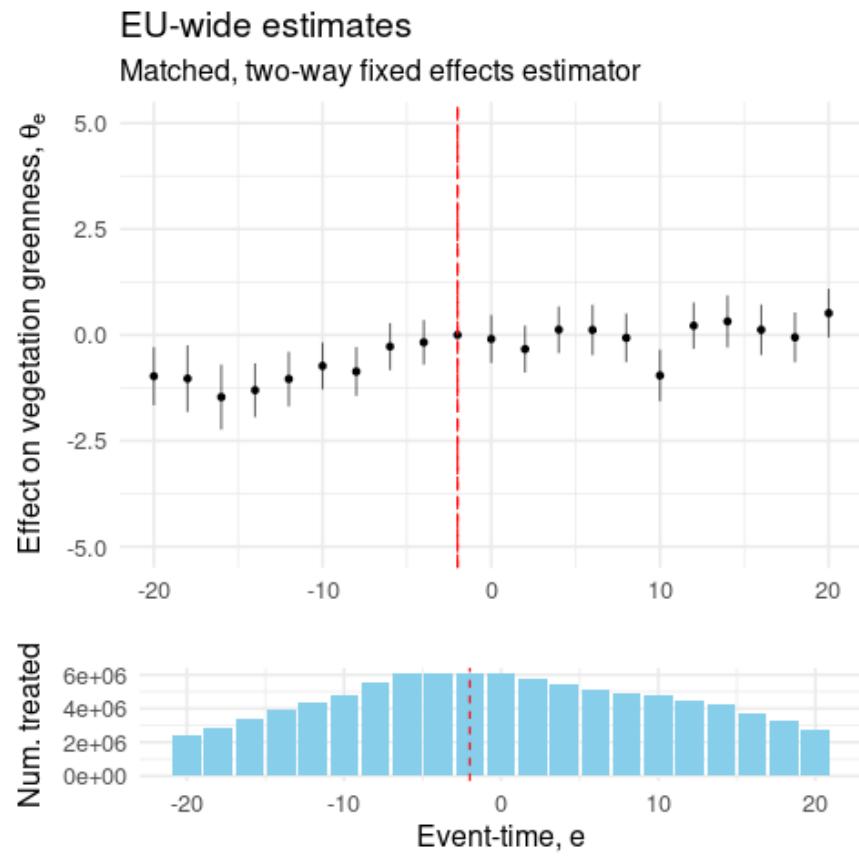
We provide some details regarding land protection given these data. Appendix Table A.14 constructs a linear probability model of land protection during our sample years, 1992-2019. We use the linear probability model to lend some simple interpretation to our coefficients, though the

Table A.13: Testing for differences in treatment effects across cohorts

	Number of cohorts	Trend:	Split difference:
		Estimate (Std. error)	Estimate (Std. error)
Albania	11	0.10 (0.0083)	0.99 (1.31)
Austria	15	0.01 (0.0010)	-0.58 (0.85)
Belgium	15	-0.06 (0.0004)	-1.06 (0.63)
Bosnia	10	0.02 (0.0082)	0.33 (2.25)
Bulgaria	8	0.11 (0.0008)	1.03 (0.63)
Croatia	12	-0.11 (0.0052)	-3.43 (1.64)
Czech	15	0.01 (0.0008)	-0.74 (0.93)
Denmark	8	-0.23 (0.0163)	-2.36 (2.66)
Estonia	13	0.19 (0.0075)	2.47 (1.22)
Finland	14	0.01 (0.0005)	0.16 (0.36)
France	15	-0.02 (0.0006)	-0.52 (0.59)
Germany	15	-0.07 (0.0006)	-1.07 (0.57)
Greece	10	0.06 (0.0027)	0.67 (0.69)
Hungary	12	0.01 (0.0081)	1.24 (1.19)
Ireland	4	-0.14 (0.0075)	-1.26 (1.56)
Italy	10	-0.06 (0.0097)	-0.13 (1.05)
Latvia	8	-0.04 (0.0041)	0.07 (1.52)
Lithuania	8	-0.06 (0.0015)	-1.88 (1.02)
Luxembourg	5	-0.03 (0.0024)	-0.04 (1.30)
Macedonia	10	-0.04 (0.0093)	0.52 (1.94)
Montenegro	5	-0.22 (0.0452)	-0.65 (3.09)
Netherlands	12	-0.05 (0.0168)	1.48 (2.72)
Norway	14	0.01 (0.0028)	0.56 (1.28)
Poland	13	-0.06 (0.0078)	-0.93 (0.50)
Portugal	11	-0.04 (0.0007)	-1.11 (0.78)
Romania	9	0.04 (0.0074)	-0.82 (1.32)
Serbia	11	0.02 (0.0015)	0.68 (1.06)
Slovakia	13	-0.01 (0.0020)	0.43 (0.94)
Slovenia	12	-0.01 (0.0015)	-0.35 (0.79)
Spain	15	0.09 (0.0007)	0.75 (0.67)
Sweden	15	0.04 (0.0003)	0.15 (0.52)
Switzerland	15	0.02 (0.0017)	0.66 (0.55)
UK	15	-0.04 (0.0043)	0.93 (1.10)

NOTES: Table reports the number of cohorts in each country (note that the maximum number here is 15 as our sample contains biannual cohorts from 1991-2019). The trend estimates a linear regression of the cohort level treatment effect θ_{cg} on the cohort itself $g - 1991$. The split difference estimates a difference of means in θ_{cg} in the later half of the treated cohorts relative to the first half of the treated cohorts. Countries require at least 3 cohorts' worth of data to be included.

Figure A.7: Dynamic estimates from matched two-way fixed effects estimation



NOTES: Figure plots estimates from Equation (A.1). Standard errors are clustered at the unit-by-foundation year level to be comparable with Callaway and Sant'Anna results. Bottom panel shows the number of treated grid cells for each event time. The dashed red line indicates the omitted pre-period, $e = -2$.

actual propensity score is generated via a logistic regression. The outcome is an indicator variable indicating any protection after 1992, creating a cross-sectional propensity score model. Column (1) focuses only on the land use recorded in the most recent decade (1990). Protection tends to occur at higher rates among land which starts out as forest, grassland, or “other,” with settled and agricultural areas having statistically significantly fewer protected areas founded after 1992. Columns (2) and (3) instead focus on key (logged) covariates, many of which may correlate with underlying land use. From column (2), an apparent contradiction emerges: settlement and cropland are less likely to get protected than “natural” land use, but nightlights and growing season length seem to increase the probability of protection. Composition effects explain the contradiction: e.g., nightlights, pooled across Europe, will reflect not only comparisons between rural and urban areas within countries but also comparisons between the more developed west and the less “luminous” east. Thus, column (3) re-runs the exercise in (2) with a country fixed effect. Notably, these two variables change sign, suggesting that holding fixed the protection regime, more economically valuable land generates a lower probability of protection.

Next, we provide tables that summarize the treatment effects depicted in Figure 7. We identify a grid cell as treated if it is at least 50% covered by a protected area. Table A.15 indicates the average treatment effects from the dynamic aggregation θ_e^{EU} , corresponding to the left panel of Figure 7. Standard errors are bootstrapped. At event-time 18, the estimates drop slightly, but this effect is inconsistent over the remaining periods, rising to -0.38 in event time 20 yet dropping to -0.07 by the last event time. This suggests that if there is a drop in nightlights, it is (1) not sustained at its initial levels and (2) occurs 20 years after treatment, making it difficult to attribute to protection alone.

Table A.14: Linear probability model predicting land protection after 1992

	(1)	(2)	(3)
Constant	0.1318*** (0.0004)	0.1602*** (0.0050)	
Settlement in 1990	0.0047*** (0.0009)		
Cropland in 1990	-0.0377*** (0.0005)		
Forest in 1990	0.0305*** (0.0005)		
Grassland in 1990	0.0167*** (0.0005)		
Other in 1990	0.0821*** (0.0016)		
Water in 1990	0.0648*** (0.0012)		
Nightlight luminosity, 1992	0.0035*** (0.0002)	-0.0009*** (0.0002)	
Elevation	-0.0054*** (0.0002)	0.0050*** (0.0003)	
Distance to shore	0.0110*** (0.0002)	-0.0171*** (0.0002)	
Solar radiance	-0.2013*** (0.0016)	-0.0686*** (0.0035)	
Growing season length in 1992	0.0637*** (0.0011)	-0.0105*** (0.0020)	
R^2	0.00732	0.01289	0.10225
Observations	5,123,083	2,285,588	2,285,588
Country fixed effects			✓

NOTES: Land protection after 1992 defined as at least 50% of a 1 square kilometer grid cell being protected by the end of the study period in 2019. Standard errors indicate heteroskedasticity-robust errors. Land-use covariates come from the HILDA landcover dataset. Covariates are expressed in logs, so all coefficients are interpretable as percent changes.

Table A.15: Dynamic treatment effects on nightlight luminosity across the EU

	Mean (Std. error)	Num. treated grids	Num. CDDA		Mean (Std. error)	Num. treated grids	Num. CDDAs
-28	0.66 (0.65)	105	28	1	-0.03 (0.02)	44,801	12,733
-27	0.35 (0.21)	1,292	384	2	-0.03 (0.02)	44,696	12,705
-26	0.20 (0.16)	2,688	799	3	-0.06 (0.03)	43,509	12,349
-25	0.19 (0.14)	4,062	1,200	4	-0.06 (0.03)	42,113	11,934
-24	0.09 (0.13)	5,418	1,532	5	-0.05 (0.04)	40,739	11,533
-23	-0.04 (0.12)	7,015	1,940	6	-0.15 (0.04)	39,383	11,201
-22	0.06 (0.11)	8,514	2,419	7	-0.17 (0.04)	37,786	10,793
-21	0.13 (0.10)	10,082	2,924	8	-0.13 (0.04)	36,287	10,314
-20	0.17 (0.09)	11,525	3,373	9	-0.08 (0.05)	34,719	9,809
-19	0.20 (0.09)	13,015	3,836	10	-0.15 (0.06)	33,276	9,360
-18	0.25 (0.08)	14,497	4,317	11	-0.12 (0.06)	31,786	8,897
-17	0.21 (0.08)	16,599	4,809	12	-0.21 (0.06)	30,304	8,416
-16	0.18 (0.08)	18,088	5,255	13	-0.23 (0.07)	28,202	7,924
-15	0.18 (0.07)	19,835	5,792	14	-0.22 (0.08)	26,713	7,478
-14	0.06 (0.06)	21,597	6,308	15	-0.20 (0.08)	24,966	6,941
-13	0.15 (0.06)	23,686	6,928	16	-0.14 (0.08)	23,204	6,425
-12	0.18 (0.06)	25,893	7,552	17	-0.12 (0.09)	21,115	5,805
-11	0.18 (0.05)	27,775	8,190	18	-0.16 (0.10)	18,908	5,181
-10	0.18 (0.05)	29,611	8,639	19	-0.29 (0.10)	17,026	4,543
-9	0.14 (0.04)	31,458	9,258	20	-0.38 (0.11)	15,190	4,094
-8	0.16 (0.04)	33,101	9,651	21	-0.39 (0.12)	13,343	3,475
-7	0.11 (0.03)	35,079	10,124	22	-0.33 (0.13)	11,700	3,082
-6	0.16 (0.03)	36,816	10,590	23	-0.36 (0.14)	9,722	2,609
-5	0.12 (0.03)	38,678	10,967	24	-0.25 (0.14)	7,985	2,143
-4	0.07 (0.03)	40,382	11,481	25	-0.34 (0.16)	6,123	1,766
-3	0.03 (0.02)	42,024	11,873	26	-0.15 (0.17)	4,419	1,252
-2	-0.00 (0.02)	43,483	12,351	27	0.36 (0.19)	2,777	860
0	-0.02 (0.02)	44,801	12,733	28	-0.14 (0.26)	1,318	382

NOTES: Dynamic treatment effect θ_t^{EU} of conservation on nightlight luminosity (an index varying between 0 and 68) estimated for the entire European Union. Treatment effect defined at the grid-cell observation level, 300 by 300 meters. Treatment requires grid cells overlap with a protected area over at least 50% of their area. Estimator is detailed in Appendix B. Underlying dataset spans 1992-2019 on an annual basis. The earliest included foundation year period is 1993, meaning the latest event time in the sample is 26.

Table A.16 indicates the average treatment effects from the cohort-level aggregation θ_g^{EU} , corresponding to the right panel of Figure 7. As with the previous table, we find little evidence of persistently-positive treatment effects, although there are some larger positive and negative estimates. Compositonally, we find no heterogeneity in treatment effects across any observables in later exercises, making these two cases outliers. Indeed, one is nearly as likely to find a positive effect on nightlights as a negative one. There is no discernible pattern in selection that indicates any systematic variation in these nightlights effects, either.

Table A.16: Cohort-level treatment effects on nightlight luminosity across the EU

	Mean (Std. error)	Num. treated grids	Num. CDDAs
1993	-0.10 (0.13)	39,540	382
1994	0.53 (0.12)	43,770	478
1995	-0.63 (0.12)	49,260	392
1996	-0.84 (0.09)	51,120	514
1997	-0.17 (0.09)	55,860	377
1998	-0.40 (0.11)	52,110	466
1999	-0.21 (0.08)	59,340	473
2000	-0.35 (0.13)	49,290	393
2001	-0.38 (0.10)	55,410	619
2002	0.30 (0.10)	55,080	449
2003	0.59 (0.14)	56,460	638
2004	-0.18 (0.11)	66,210	624
2005	-0.22 (0.07)	62,670	620
2006	-0.06 (0.10)	52,860	516
2007	-0.21 (0.10)	52,410	537
2008	-0.15 (0.11)	44,670	446
2009	0.03 (0.10)	63,060	492
2010	0.00 (0.14)	44,460	481
2011	0.17 (0.19)	44,700	463
2012	0.12 (0.11)	43,290	449
2013	-0.19 (0.19)	47,040	505
2014	-0.08 (0.17)	44,970	479
2015	-0.05 (0.10)	47,910	408
2016	0.01 (0.11)	40,680	332
2017	0.18 (0.08)	41,220	401
2018	0.13 (0.07)	41,880	415
2019	0.10 (0.09)	35,610	356
2020	0.20 (0.25)	3,150	28

NOTES: Cohort-level treatment effect θ_g^{EU} of conservation on nightlight luminosity (an index varying between 0 and 68) estimated for the entire European Union. Treatment effect defined at the grid-cell observation level, 300 by 300 meters. Treatment requires grid cells overlap with a protected area over at least 50% of their area. Estimator is detailed in Appendix B. Underlying dataset spans 1992-2019 on an annual basis. The earliest included foundation year period is 1993.

C.4 Heterogeneous treatment effects

Here we describe in more detail the heterogeneous treatment effects obtained via the random forest method of Wager and Athey (2018). We present results estimated on 30 different 1% stratified random samples of the data as described in Appendix B. In the sample, the average treatment effect is 0.01 (0.26).

The first result is a test of heterogeneity in treatment effects. We use the test in Wager and Athey (2018), which amounts to a random forest implementation of Chernozhukov et al. (2024). The results are shown in Table A.17. The test reports two coefficients in a regression. Observed greenness Y is projected onto the average treatment effect estimated by the random forest and the conditional average treatment effect estimated by the random forest:

$$Y_{it} = \beta_0 ATE_{it} + \beta_1 (CATE_{it} - ATE_{it}) + \epsilon_{it}$$

Intuitively, the true average treatment effect should contribute a coefficient of exactly 1 as a one-unit increase in the average treatment effect drives a 1 unit increase in expected counterfactual outcomes. Thus, the test for the coefficient on the ATE is a two-sided test for whether the coefficient is statistically different from 1. We find an estimate of 1.08 (0.10). Similarly, the coefficient on the CATE should be at least 1: if the CATE changes by 1 unit, we should expect the outcome itself to change by at least this much if the CATE is meaningful. The test on the CATE is one-sided, determining if the coefficient is greater than 1: rejection implies that the CATE predicted by the random forest is not driving deviations from the average treatment effect. Our results indicate our measured CATE is meaningful with coefficient 1.54 (0.07). Thus, our random forest has picked up on statistically significant deviations from the average treatment effect.

Table A.17: Test of the random forest model calibration

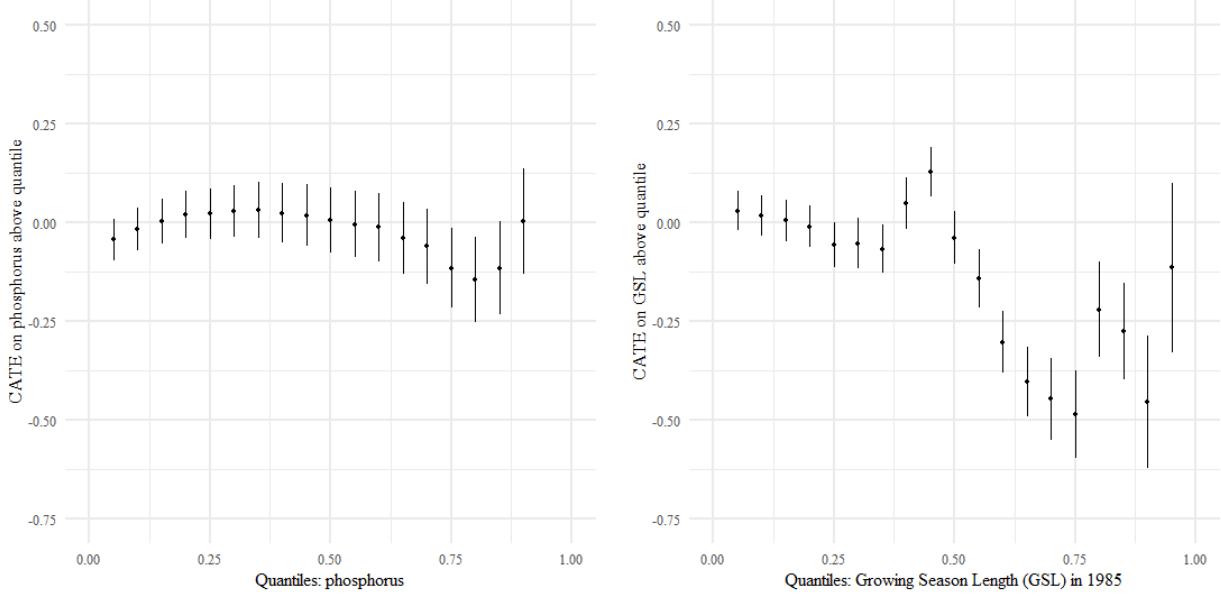
	Coefficient	Standard error
Mean forest prediction	1.01	(0.07)
Differential forest prediction	1.54	(0.07)

NOTES: Table presents the test in Wager and Athey (2018) for heterogeneous treatment effects. Sample selection discussed in B. The first coefficient describes whether the model captures the mean forest prediction. As it is statistically indistinguishable from 1, the random forest appears to be fit well. The second coefficient describes whether the model is able to find heterogeneity in calibrated treatment effects. The coefficient is robustly greater than 1, confirming that we have found salient heterogeneity. Standard errors are bootstrapped across samples.

Despite the statistically significant heterogeneity in the data, treatment effects vary little in economic terms. To further test the economic significance of the CATE measures, we plot conditional average treatment effects along the distribution of selected control variables in Figure A.8.

We are interested, in particular, if variables associated with high-quality agricultural land may change the predicted CATE. If we see, for example, that high soil phosphorus content, indicating soil that is highly agriculturally productive, has a higher treatment effect, this suggests protection

Figure A.8: Distribution of the conditional average treatment effect for phosphorus content in soil and growing season length in 1985



NOTES: Both panels plot estimates of the conditional average treatment effect across 20 ventiles of a key explanatory variable. Bar at ventile $q \in \{0.05, 0.1, \dots, 0.95\}$ corresponds to the CATE of plot i conditional on a realization of X above that quantile, e.g., $\mathbb{E}[\tau_i|X > x(q)]$. Red dotted line highlights a zero CATE. Confidence bands are 95% confidence intervals constructed through a doubly-robust procedure.

causes plots more suitable for agriculture to green more than plots less suitable for agriculture. Such a result is consistent with a subset of land protection having positive treatment effects due to a valuable counterfactual land value. We present the phosphorus CATE plot in Figure A.8 (left panel). We reject any trend in CATEs based on the phosphorus content of soil.

As a second example, we illustrate the conditional average treatment effect as a function of the growing season length (GSL) in 1985 in the right panel of Figure A.8. Longer growing seasons indicate weather conditions which are more conducive to growing common fieldcrops (though not specialty crops or horticulture). Plots in the top 25% of the sample with respect to their 1985 GSL have a conditional average treatment effect of -0.5 , which is statistically more negative than the full-sample ATE of 0.01 . Land with longer growing seasons in 1985, and thus more conducive to field crops, was less likely to experience beneficial effects of land protection, though the difference is small in absolute terms. Control land with similar GSL and observables greened more in the intervening period.

D Robustness checks

In this section we discuss various robustness checks of the doubly-robust difference-in-differences estimator discussed in Appendix B.

Functional form: first differences. Our main specification considers differences in levels of greenness due to land protection. Here, we test a first difference. There are two advantages to this robustness check. First, as vegetation growth tends to be a slow process, there is an advantage to testing for a change in the first differences of vegetation rather than the levels. Where in levels it can take a very long time to convert to biodiversity-rich, green forest from a low-greenness land cover like grassland, in first differences a clear change in growth rates should be easier to detect earlier in the process of vegetation regrowth. Second, the first-differences specification allows us to check that our levels functional form does not drive the null result. By taking a first difference, we are removing the effect of any time-invariant contributors to greenness. This approach is thus robust to arbitrary time-invariant heterogeneity, at the cost of correctly specifying determinants of changes in greenness. We apply the exact same matching function with the exception of matching on pre-period average greenness in levels and in trends. The aggregate treatment effect θ^{EU} is -0.12 (0.11).

We report the EU-wide aggregated treatment effects in Figure A.9. The left panel illustrates the dynamic treatment effects on changes in vegetation greenness, θ_t^{EU} . Pre-trends are balanced around 0 for 20 years prior to treatment. There is no evidence of a kink in greenness at the time of treatment. Importantly, while treatment effects in levels may take time to appear, we should see an immediate kink in first differences if there is indeed vegetation growth occurring post-treatment, which was not present in the counterfactual. The right panel of Figure A.9 illustrates the cohort-level treatment effects. Aside from an outlier in 2001, cohort effects are flat.

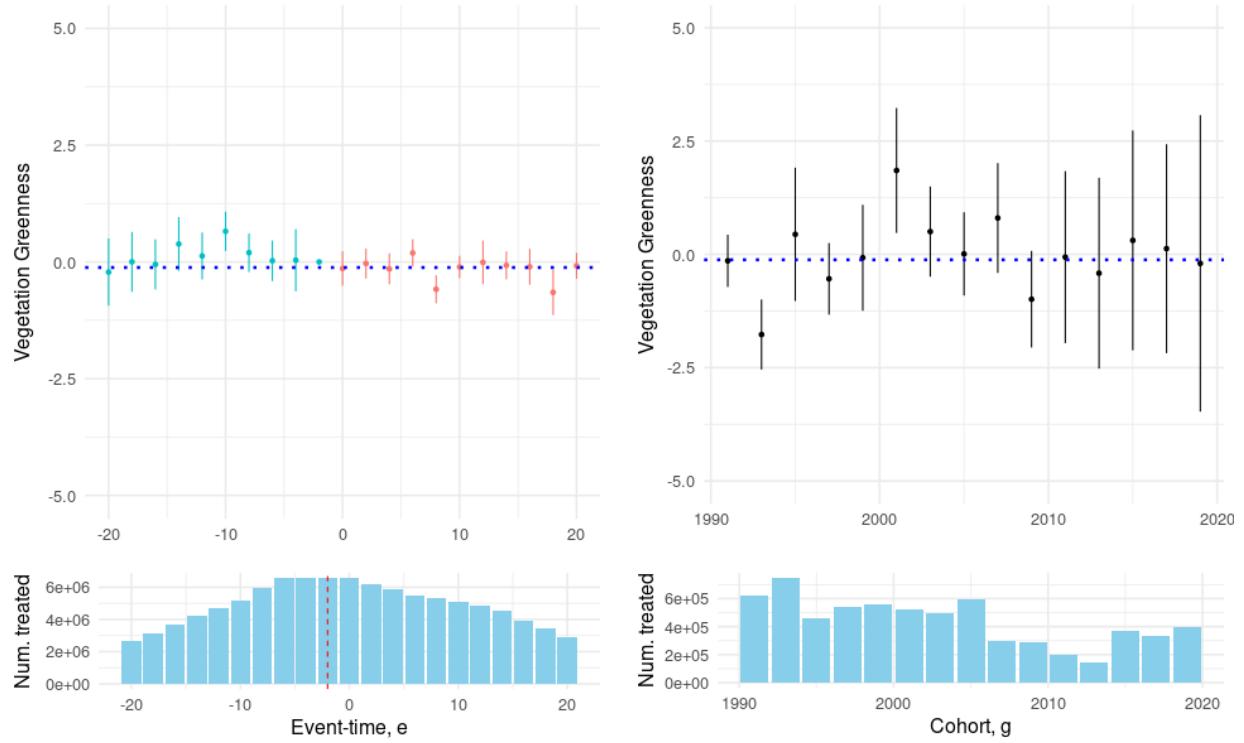
Collectively, our first-difference results emphasize the main zero treatment effect finding. There is evidence for neither a notch nor a kink in vegetation growth driven by protection.

Greenness thresholded sample. In this robustness check, we re-run our main estimator on a greener sub-sample of the dataset. For inclusion in this robustness check, a grid cell in our data must achieve a greenness of at least 40 at any point in the sample period 1985-2019. In this greener sample, we are largely dropping bare areas and dirt. Rocky regions are less likely to be protected with the intention to green, and greenness may be a less suitable metric of their ecological health. Importantly, a threshold of 40 is unlikely to remove meadows and wetlands.

The overall treatment effect in this sub-sample is -0.05 (0.06). Figure A.10 shows the dynamic (left) and cohort (right) effects. Both corroborate a 0 treatment effect on average. Greener grid cells have a slight positive pre-trend, but this pre-trend if anything suggests treated areas greened less post-protection (Rambachan and Roth 2023). We conclude our results are not driven by the inclusion of low greenness areas.

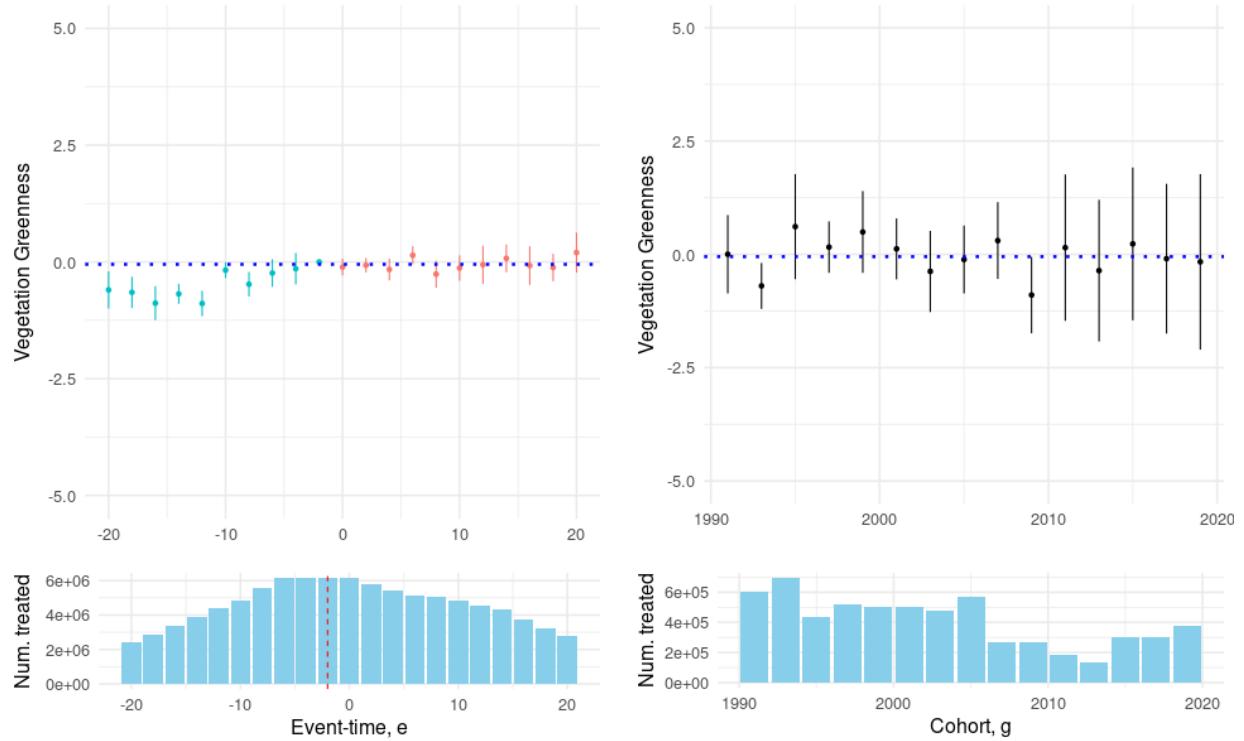
Not-yet-treated sample. Our main approach uses never-treated grids as controls for protected grid cells. An alternate approach leverages not-yet-treated protected areas as controls for earlier protected areas. Such an approach is desirable if there is significant selection of protected areas, rendering never-treated units an unsuitable control. On observables, we have significant over-

Figure A.9: Treatment effects on a first difference of vegetation greenness over event time and by cohort



NOTES: Treatment effects on a first difference of greenness aggregated by event-study period θ_e^{EU} in the left panel, and by cohort θ_g^{EU} in the right panel. Sample trims first calendar year, 1985, due to first differencing: we match on values from 1987 and 1989 for all cohorts treated in or after 1991. Both panels use the Callaway and Sant'Anna doubly-robust estimator with bootstrapped confidence bands.

Figure A.10: Treatment effects on NDVI: threshold 40 NDVI group

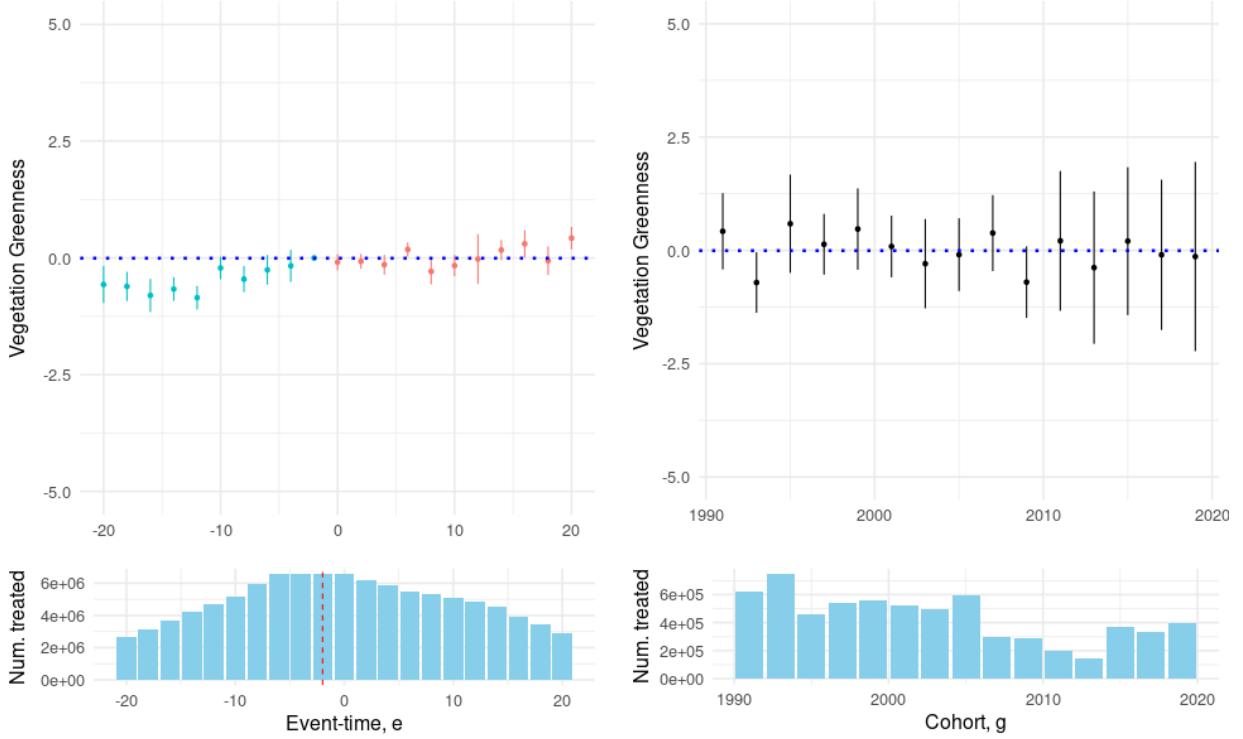


NOTES: Treatment effects on greenness aggregated by event-study period θ_e^{EU} in the left panel, and by cohort θ_g^{EU} in the right panel. Both panels use the Callaway and Sant'Anna doubly-robust estimator with bootstrapped confidence bands. Sample includes only grid cells with $NDVI \geq 40$ at least once between 1985-2019.

lap across treated and never-treated groups: that said, if there is sharp unobservable selection, the not-yet-treated group may be an appropriate alternative so long as protection is not anticipated.

Appendix Figure A.11 illustrates the dynamic (left) and cohort (right) effects of protection on greenness using the not-yet-treated group as a control. Pre-trends are quite flat, and we see no evidence of greening after treatment. Cohort effects are also similar in magnitude. There is one cohort with a significant, positive treatment effect in 2001, but there is also an equally large negative treatment effect in 1993. Overall, we do not conclude that not-yet-treated controls reveal an effect of land protection.

Figure A.11: Treatment effects on NDVI: not-yet-treated control group.



NOTES: Treatment effects on greenness aggregated by event-study period θ_e^{EU} in the left panel, and by cohort θ_g^{EU} in the right panel. Both panels use the Callaway and Sant'Anna doubly-robust estimator with bootstrapped confidence bands.

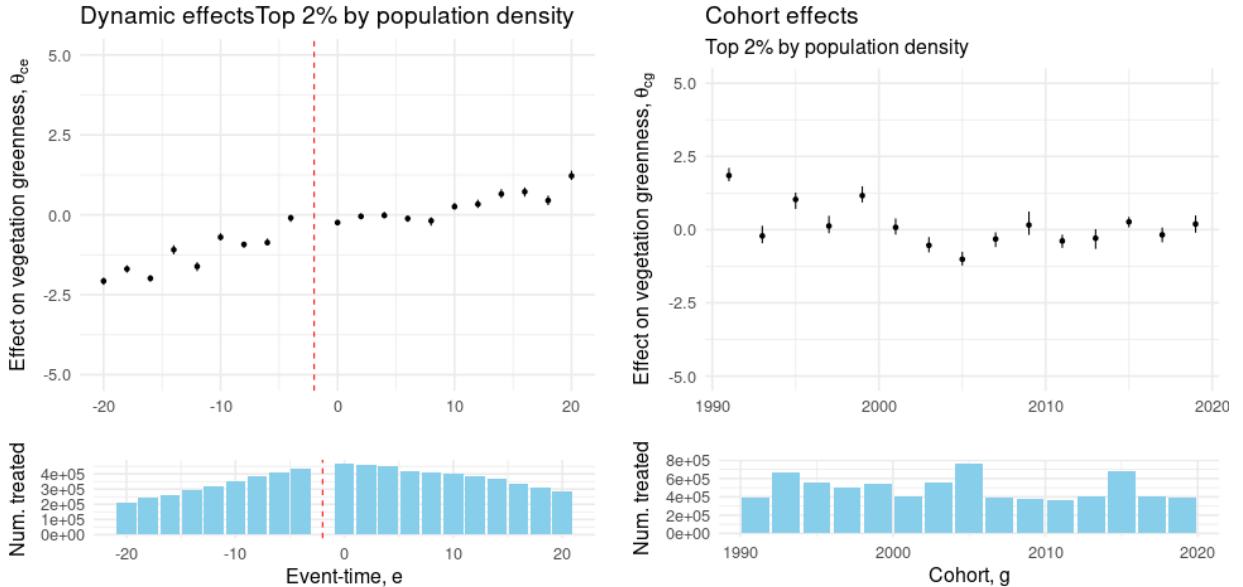
High population-density sample. Next, we test for robustness with respect to the sample definition. We re-run the Callaway and Sant'anna estimator for the top 2% of land by population density in the EU: land which has a population density of 913 people per sq. km. or more. These areas are highly urbanized in comparison to the full sample. Under our framework, high-population protected areas are likely to be protected less strictly or to yield much larger environmental dividends from protection.

A concern with taking a small cut of the data is that such a cut of the data may have very few protected areas. However, by country, we see significant protection in dense areas. Germany,

Denmark, and the Netherlands all have at least some protection in over 30% of their most dense grid cells. As a proportion of the total population dense sample, Germany, the UK, and Italy collectively comprise approximately half, followed by France comprising 10% of the sample.

In Appendix Figure A.12, we conclude that our approach is less well-suited to this highest population density land. We find that controls in populated areas are poor controls for protected areas in populated areas. Protected areas show remarkable greening trends for at least 3 decades prior to protection, consistent with pockets of long-run land abandonment or the elevation of existing parks to protected status. After protection, if anything, protected areas appear to green more slowly. This results in a *negative* treatment effect in aggregate. In the cohort effects, aside from a standout cohort in 1991, we do not conclude that there is a long-run trend in vegetation greening.

Figure A.12: Treatment effects on NDVI: top 2% of Europe by population density.



NOTES: Treatment effects on greenness aggregated by event-study period θ_e^{EU} in the left panel, and by cohort θ_g^{EU} in the right panel. Both panels use the Callaway and Sant'Anna doubly-robust estimator with bootstrapped confidence bands. Sample includes any 300×300 grid cell with population density above the 98th percentile of population density at the EU level measured in our 1×1 km data. Dashed vertical line on the dynamic (left) plot indicates the universal base period $e = -2$.

Selection on unobservables: spatial first differences. We next present a spatial first differences (SFD) strategy as a method to tackle selection on unobservables (Druckenmiller and Taylor 2022). This methodology takes a first difference of the data along a given spatial axis, thus creating comparisons between areas that are close by in space. The advantage of SFD is to eliminate spatially correlated sources of unobservable selection.⁵ For example, agricultural productivity

5. It does not explicitly treat violations of SUTVA, where control plots nearby CDDAs may be affected by protection. A conventional method for testing for such spatial spillovers is a “donut” difference-in-differences design in

of protected land or its unobserved development value may be unobservably lower. As long as these sources of heterogeneity are spatially correlated, differencing across space attenuates their importance in our treatment effect estimates. By spatially differencing the data, we compare areas which are nearby to one another, thus indirectly controlling for these unobservable drivers. We difference observations along the x -dimension so that areas with the same latitude are differenced against their neighbors. Introducing the index $i = (x, y)$ to identify a grid cell by the coordinates of the centroid of that grid cell, the outcome variable is:

$$\Delta Y_{xyt}^{LAT} = Y_{xyt} - Y_{x-1,yt}$$

The focus of the SFD is on estimating the greenness differences between treated units and their nearest neighbors, relative to matched controls. In principle, we could re-estimate SFD for many axes, such as a vertical or diagonal axis. When we estimate treatment effects, the SFD method estimates the difference-in-differences on the spatially-differenced outcome variable.

When calculating the spatial first difference across the x -dimension, we obtain an EU-wide treatment effect of -0.02 (0.14). Figure A.13 plots the event-study (left panel) and cohort-level aggregations (right panel). The left panel demonstrates a flat trend in the post-treatment periods. Cohort effects show no trend through 2015, with a slightly lower treatment effect in the 2017 and 2019 cohorts. Overall, the SFD estimates are an even more precise zero than our main specification.

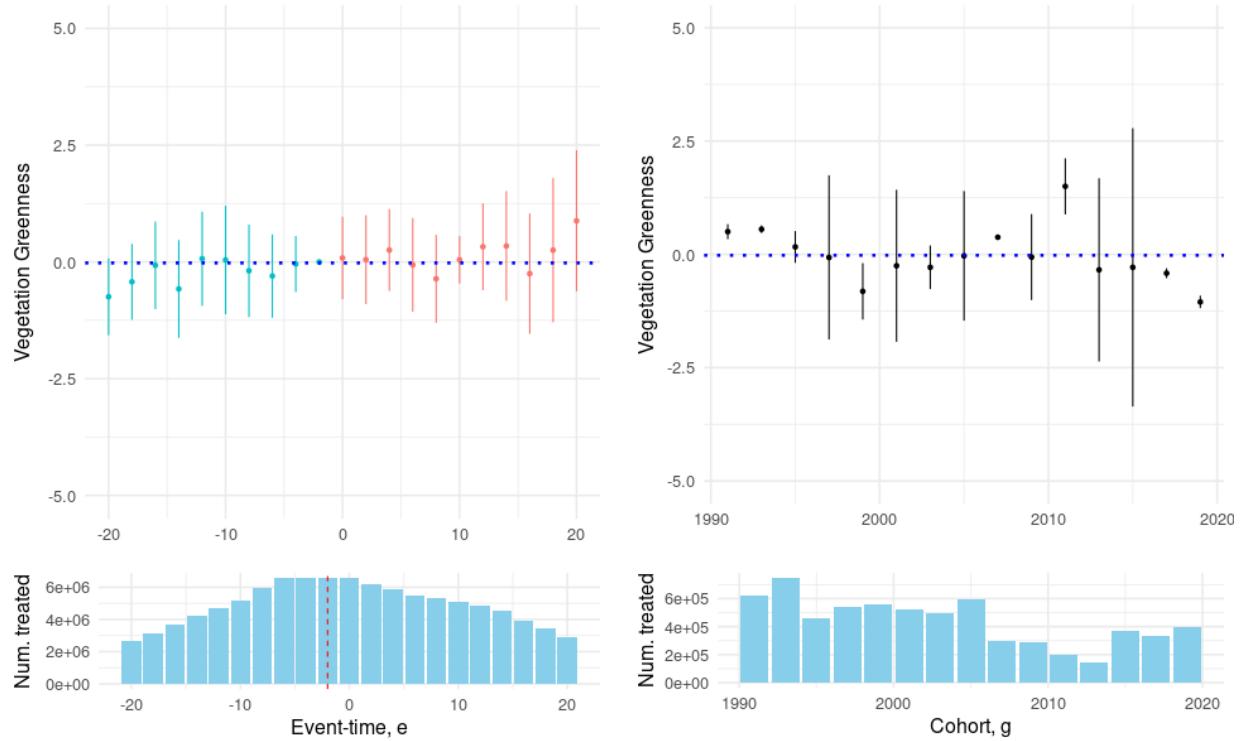
E Biodiversity outcomes from species counts

BioTIME data consist of a panel of animal species and vegetation flora biomass studies. Studies enter and exit the panel as they are conducted by biologists and ecologists. Records consist of a year, a species identifier, a study identifier, and either a count of species abundance or a measure of vegetation biomass. When focusing on the landmass of the European Union, there are a total of 58 studies, most of which focus on animal counts. We restrict our analysis in this section to species counts of animals as a measure of biodiversity.

Of the 58 BioTIME studies in the EU, 53 are within 5 kilometers of a protected area listed in the CDDA system. The remaining 5 studies are too few to serve as credible controls to establish causal effects as in our core analysis of greenness and nightlights. Instead, we construct event-study estimators of the impact of nearby CDDA openings on measured abundance. As a result, rather than leveraging variation relative to control units, we only look at within-study variation to determine whether structural breaks appear around the foundation of protected areas. The potential selection of these few study sites and the lack of a valid control group in the BioTIME

which the researcher discards potentially contaminated units nearest a protected area and recomputes treatment effects. In our application, the donut approach is prohibitively expensive from a computational perspective: it requires calculating an individual buffer for each of our over 100,000 treated areas, many of which may have overlapping buffers. Moreover, violations of SUTVA should bias our results away from zero, as protection may push economic activity to the area just outside of the CDDA boundary; despite that, we find zero treatment effects of protection.

Figure A.13: Treatment effects on NDVI: spatial first differences



NOTES: Treatment effects on greenness aggregated by event-study period θ_e^{EU} in the left panel, and by cohort θ_g^{EU} in the right panel. Outcomes are spatially differenced in the x direction. Both panels use the Callaway and Sant'Anna doubly-robust estimator with bootstrapped confidence bands.

data limit the external validity of these results. We thus present our specification as descriptive evidence rather than a causal design.

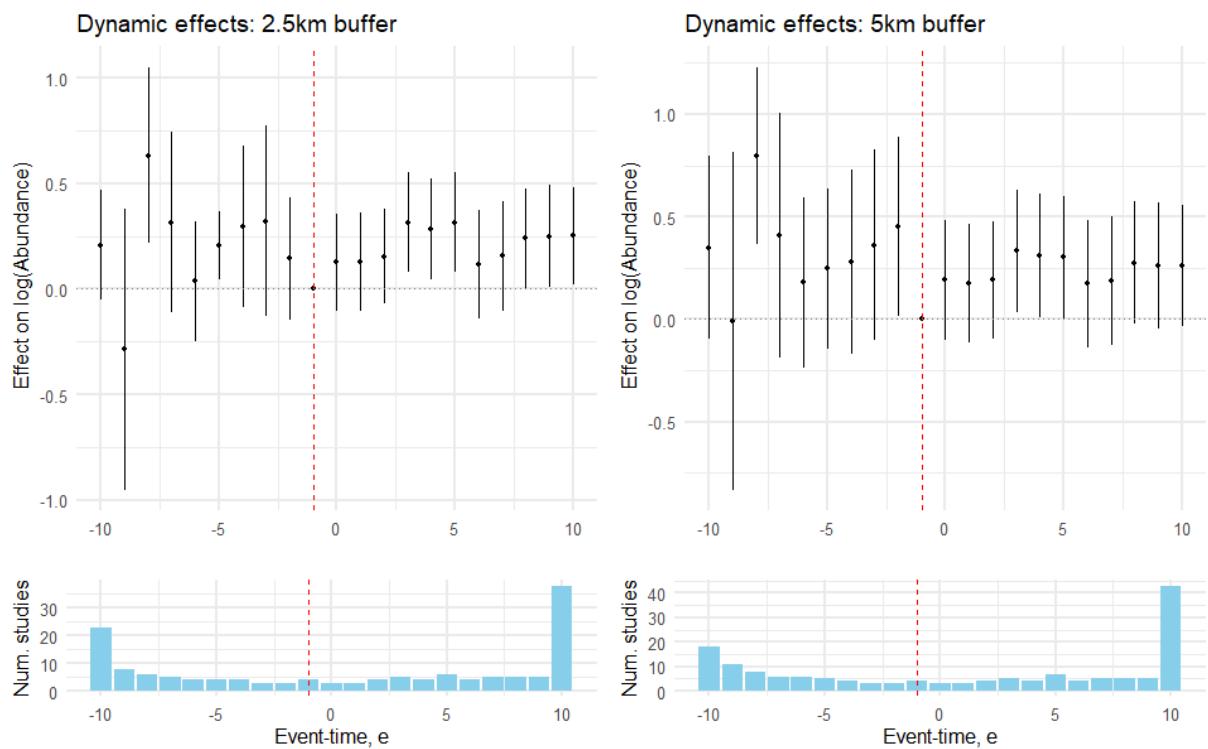
Our econometric specification considers a study i , species s , and year t . We construct a buffer of distance b kilometers around the study area, the neighborhood N_i^b . We assign treatment of the study area according to the minimum foundation year g of overlapping CDDAs. That is, the treatment indicator is $D_{it}^b = \mathbb{1}[t \geq \min_{N_i^b} g]$ with event time $e^b = t - \min_{N_i^b} g$. Then, the event-study design amounts to the following regression, where λ is used to denote a fixed effect:

$$Y_{ist} = \sum_{e^b=-10}^{10} \beta_{e^b}^b D_{it}^b + \lambda_i + \lambda_s + \epsilon_{ist} \quad (\text{A.2})$$

In Figure A.14 we plot the event-study coefficients $\beta_{e^b}^b$ for two buffer distances: 2.5 km and 5 km. We estimate the regression in logs. Event-study coefficients are indistinguishable from zero in the immediate time frame around the first foundation year in the vicinity. The average effect across post-periods was 0.013 (0.03) and 0.012 (0.03) at the 2.5 and 5 km buffers, respectively. Overall, we cannot reject the null of no structural break in these settings.

The BioTIME data are spatially concentrated in a few areas within Europe. This lack of spatial variation means many sites are close to one another, and close to many potential CDDAs. When expanding the treatment buffer from 2.5 km to 5 km, we move from 52 to 53 treated studies out of the original pool of 58. The jump in species counts represented is much smaller, increasing by less than 0.1%. Thus, there is a large density of species studies within the 5 km boundary. This indicates that the biodiversity data lack spatial breadth, making valid inference difficult.

Figure A.14: Event-study coefficients measuring species abundance with respect to nearby protected area foundation



NOTES: Estimates show the effect of founding the first protected area within 2.5 km (left) and 5 km (right) of a BioTIME study site. Estimates are a pure event study.

References for Appendix

Ballabio, Cristiano, Emanuele Lugato, Oihane Fernández-Ugalde, Alberto Orgiazzi, Arwyn Jones, Pasquale Borrelli, Luca Montanarella, and Panos Panagos. 2019. “Mapping LUCAS Topsoil Chemical Properties at European Scale Using Gaussian Process Regression.” *Geoderma* 355:113912.

Bickenbach, Frank, Eckhardt Bode, Peter Nunnenkamp, and Mareike Söder. 2016. “Night Lights and Regional GDP.” *Review of World Economics* 152:425–447.

Callaway, Brantly, and Pedro H.C. Sant’Anna. 2021. “Difference-in-Differences with Multiple Time Periods.” *Journal of Econometrics*, Themed Issue: Treatment Effect 1, 225 (2): 200–230.

Chen, Xi, and William Nordhaus. 2011. “Using Luminosity Data as a Proxy for Economic Statistics.” *Proceedings of the National Academy of Sciences* 108 (21): 8589–8594.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val. 2024. “Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments.” Forthcoming, *Econometrica*.

Cornes, Richard C., Gerard van der Schrier, Else J.M. van den Besselaar, and Philip D. Jones. 2018. “An Ensemble Version of the E-OBS Temperature and Precipitation Data Sets.” *Journal of Geophysical Research: Atmospheres* 123 (17): 9391–9409.

Dannenberg, Matthew P., Erika K. Wise, and William K. Smith. 2019. “Reduced Tree Growth in the Semiarid United States Due to Asymmetric Responses to Intensifying Precipitation Extremes.” *Science Advances* 5 (10): eaaw0667.

Donaldson, Dave, and Adam Storeygard. 2016. “The View from Above: Applications of Satellite Data in Economics.” *Journal of Economic Perspectives* 30 (4): 171–198.

Dornelas, Maria, Laura H. Antao, Faye Moyes, Amanda E. Bates, Anne E. Magurran, Dušan Adam, Asem A. Akhmetzhanova, Ward Appeltans, Jose Manuel Arcos, Haley Arnold, et al. 2018. “BioTIME: A Database of Biodiversity Time Series for the Anthropocene.” *Global Ecology and Biogeography* 27 (7): 760–786.

Druckenmiller, Hannah, and Charles A. Taylor. 2022. “Wetlands, Flooding, and the Clean Water Act.” *American Economic Review* 112 (4): 1334–1363.

Fick, Stephen E., and Robert J. Hijmans. 2017. “WorldClim 2: New 1-km Spatial Resolution Climate Surfaces for Global Land Areas.” *International journal of climatology* 37 (12): 4302–4315.

Gibson, John, Susan Olivia, Geua Boe-Gibson, and Chao Li. 2021. “Which Night Lights Data Should We Use in Economics, and Where?” *Journal of Development Economics* 149:102602.

Goodman-Bacon, Andrew. 2021. “Difference-in-Differences with Variation in Treatment Timing.” Themed Issue: Treatment Effect 1, *Journal of Econometrics* 225 (2): 254–277.

Günther, Andreas, Miet Van Den Eeckhaut, Jean-Philippe Malet, Paola Reichenbach, and Javier Hervás. 2014. “Climate-Physiographically Differentiated Pan-European Landslide Susceptibility Assessment Using Spatial Multi-Criteria Evaluation and Transnational Landslide Information.” *Geomorphology* 224:69–85.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil. 2012. “Measuring Economic Growth from Outer Space.” *American Economic Review* 102 (2): 994–1028.

Li, Xuecao, Yuyu Zhou, Min Zhao, and Xia Zhao. 2020a. “A Harmonized Global Nighttime Light Dataset 1992–2018.” *Scientific Data* 7:168.

———. 2020b. “Harmonization of DMSP and VIIRS Nighttime Light Data from 1992–2021 at the Global Scale.” *Scientific Data* 7:168.

Panagos et al., Panos. 2015. “A New European Slope Length and Steepness Factor (LS-Factor) for Modeling Soil Erosion by Water.” *Geosciences* 5 (2): 117–126.

Pasho, Edmond, J. Julio Camarero, Martín de Luis, and Sergio M. Vicente-Serrano. 2012. “Factors Driving Growth Responses to Drought in Mediterranean Forests.” *European Journal of Forest Research* 131:1797–1807.

Peled, E., Emanuel Dutra, Pedro Viterbo, and Alon Angert. 2010. “Technical Note: Comparing and Ranking Soil Drought Indices Performance over Europe, Through Remote-Sensing of Vegetation.” *Hydrology and Earth System Sciences* 14 (2): 271–277.

Rambachan, Ashesh, and Jonathan Roth. 2023. “A More Credible Approach to Parallel Trends.” *Review of Economic Studies* 90 (5): 2555–2591.

Roy, D.P., V. Kovalskyy, H.K. Zhang, E.F. Vermote, L. Yan, S.S. Kumar, and A. Egorov. 2016. “Characterization of Landsat-7 to Landsat-8 Reflective Wavelength and Normalized Difference Vegetation Index Continuity.” *Remote Sensing of Environment* 185:57–70.

Tóth, Gergely, and Tamas Hermann. 2016. “European Map of Soil Suitability to Provide a Platform for Most Human Activities (EU28),” <https://data.jrc.ec.europa.eu/dataset/jrc-esdac-42>.

Van Oijen, M., J. Balkovi, C. Beer, D.R. Cameron, P. Ciais, W. Cramer, T. Kato, et al. 2014. “Impact of Droughts on the Carbon Cycle in European Vegetation: A Probabilistic Risk Analysis Using Six vegetation Models.” *Biogeosciences* 11 (22): 6357–6375.

Vieira, Joana, Cristina Nabais, and Filipe Campelo. 2021. “Extreme Growth Increments Reveal Local and Regional Climatic Signals in Two *Pinus pinaster* Populations.” *Frontiers in Plant Science* 12:658777.

Wager, Stefan, and Susan Athey. 2018. “Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests.” *Journal of the American Statistical Association* 113 (523): 1228–1242.

Warszawski, L., K. Frieler, V. Huber, F. Piontek, O. Serdeczny, X. Zhang, Q. Tang, M. Pan, Y. Tang, Q. Tang, et al. 2017. “Gridded Population of the World, Version 4 (GPWv4): Population Density.” *Atlas of Environmental Risks Facing China Under Climate Change*, 228.

Winkler, Karina, Richard Fuchs, Mark D.A. Rounsevell, and Martin Herold. 2020. “HILDA+ Global Land Use Change between 1960 and 2019.” *Pangaea* 921846.