Review of Economic Studies (2026) **00**, 1–37 doi:10.1093/restud/rdag080
© The Author(s) 2026. Published by Oxford University Press on behalf of The Review of Economic Studies Limited.
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.
Advance access publication 29 July 2026

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

An Evaluation of Protected Area Policies in the European Union

Tristan Earle Grupp
University of Pennsylvania, USA
Prakash Mishra
California Institute of Technology, USA
Mathias Reynaert
Toulouse School of Economics, University of Toulouse Capitole, France
and
Arthur A. van Benthem
University of Pennsylvania, USA

First version received June 2024*; Accepted January* 2026 *(Eds.)*

The European Union (EU) 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.

Key words: Land protection, Protected areas, Conservation, Biodiversity, Deforestation, Vegetation cover, Nightlights, Staggered Difference-in-differences, Europe

JEL codes: Q23, Q24, Q57, R14

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 "30 × 30 target," which was the main outcome of the COP 15 Kunming-Montreal global biodiversity conference and is sometimes referred to as the "Paris Agreement for Nature," was ratified in

The editor in charge of this paper was Bård Harstad.

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 30×30 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 nightlights. 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 EU from 1985 to 2020. Our data include key outcome variables (vegetation cover, nightlights) at the 300×300 m or 1×1 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, Normalized Difference Vegetation Index (NDVI) derived from Landsat, 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 *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 *et al.*, 2026; Burgess *et al.*, 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.

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 *et al.*, 2010). Yet, this greening trend has slowed down substantially in more recent decades.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 (de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021; Sun and Abraham, 2021). To overcome this challenge, we apply the doubly-robust estimator of Callaway 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

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 *et al.*, 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 *et al.*, 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, 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 *et al.*, 2022; Cheng *et al.*, 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 *et al.*, 2023; Reynaert *et al.*, 2023). Our study is unique in that our data span four decades at high frequency and have 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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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, 2026), and mechanism design of US conservation markets (Aronoff and Rafey, 2026; Aspelund and Russo, Forthcoming). Our paper contributes to this literature by focusing on the environmental outcomes of a large-scale, regulation-based conservation programme 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 30 × 30 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 EU

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

3. The Auffhammer *et al.* (2021) and Bahrami *et al.* (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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 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 *et al.* (2011) and is titled "Dialogue for Natura 2000." Under this plan, the French state contacts the prefect of a region if its 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 *et al.*, 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 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

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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 programme 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 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 $\mathbb{I} = \{1, 2, \ldots, I\}$. 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

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 sources 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 modelling 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 $a_i(0) = 0$).

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:

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

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

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.

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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:

$$\max_{\{D_{it}\}_{i=1}^{I}} \sum_{0}^{\infty} \delta^t \sum_{i=1}^{I} [v_i\,(a_i(D_{it})) + e_i\,(a_i(D_{it})) - c\,(D_{it})] \quad (3)$$

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

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:

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

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

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

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: $S_{t+1} = S_t \backslash S^*$.

Repeat Steps 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, $\Delta e_i > 0$.

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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.

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

$$\max_{\{D_{it}\}_{i=1}^{I}} \sum_{0}^{\infty} \delta^t \sum_{i=1}^{I} \left[v_i \left(a_i(D_{it}) \right) + (1-\alpha_1) e_i \left(a_i(D_{it}) \right) - c \left(D_{it} \right) + \alpha_2 \left(D_{it} \right) \right] \tag{5}$$

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

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

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 $t = t^{\text{target}}$.

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 favour the representation of those affected by the losses from protection. As described in Section 2,

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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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

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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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.

3.3. *Siting and strictness under delegation*

Until now, we have modelled 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 costs. 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:

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

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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 EU between 1985–2019.[15] The most granular dataset (to analyse vegetation cover) divides Europe into 117 million equal-sized grids of 300 m by 300 m. The second dataset (to analyse 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 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 *et al.*, 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 been 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 animal counts, including insects, but we lack evidence on how vegetation cover

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 *et al.* (2024) and McDonald *et al.* (2024) for an analysis of the impacts of marine protected areas under the 30×30 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.

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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 Programme (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 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 harmonizing 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 *et al.* (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

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/.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 EU.

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 increases in NDVI between 1985 and 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.

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 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]

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

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

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

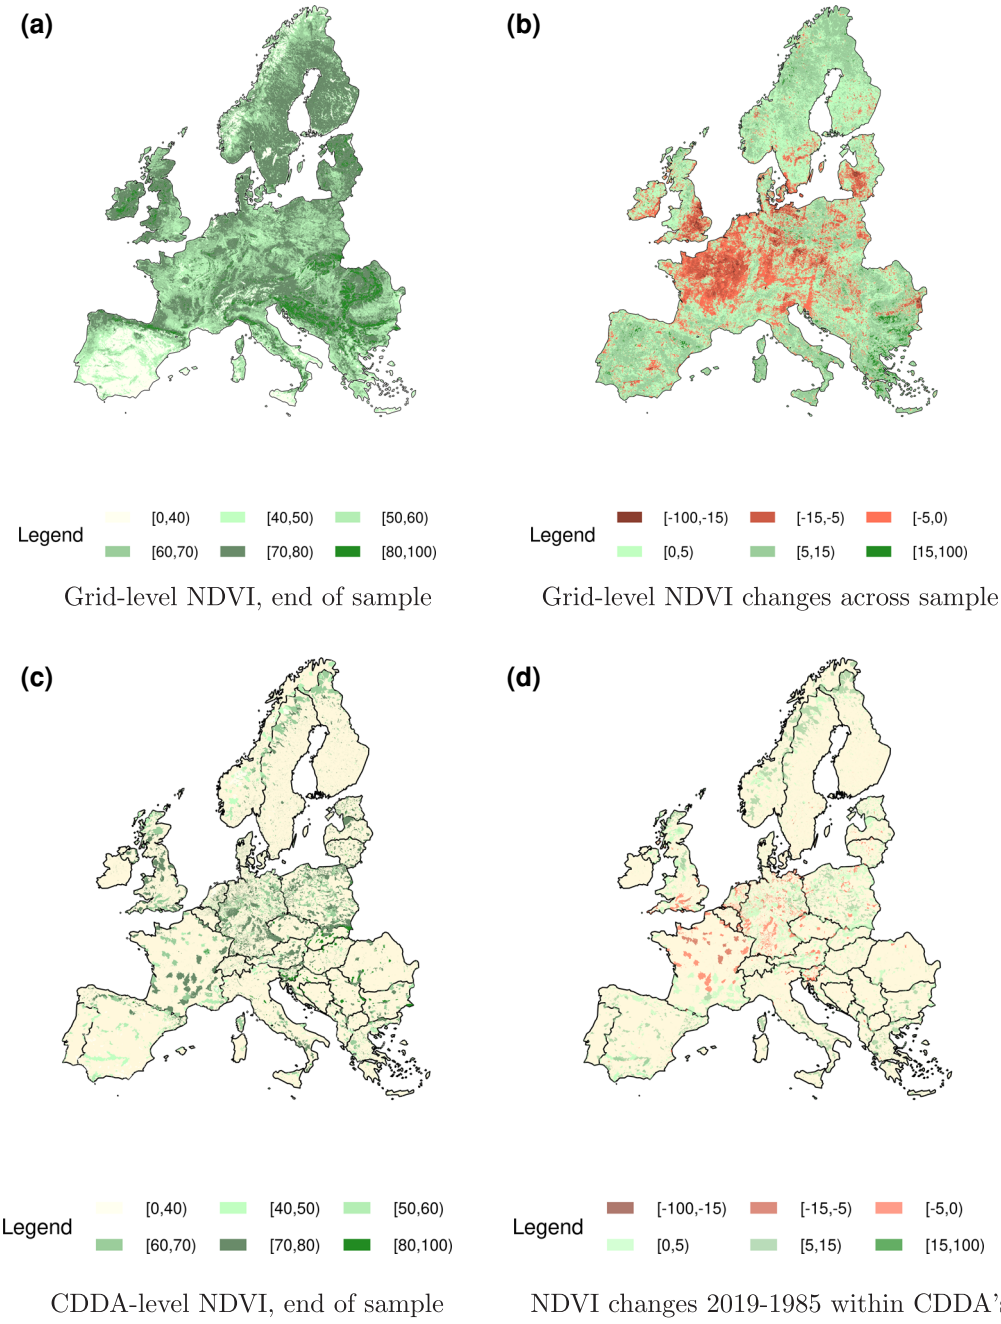


FIGURE 1

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 2015–9 and 1985–9. Panel (c) plots NDVI within each CDDA for 2015–2019. Panel (d) plots the change in NDVI within each CDDA between 2015–9 and 1985–9.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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	6.2	9.3	11.0	12.2
Percent of grid protected	21.8	39.3	27.1	41.3
Population density, 2000	46.3	142.1	106.3	271.9
Slope steepness	2.3	2.8	2.2	2.8
Solar radiance	10.7	2.2	11.4	2.0
Precipitation	762.9	267.2	799.2	273.6
Crop suitability	6.2	1.3	6.3	1.3
Forest suitability	3.5	1.4	4.1	1.5
Grassland suitability	5.6	1.5	5.6	1.6
Potassium	147.0	55.6	174.3	70.6
Nitrogen	2.2	1.0	2.3	1.0
Growing season length	221.2	66.7	257.9	56.7

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 2000–10.

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 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.

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

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

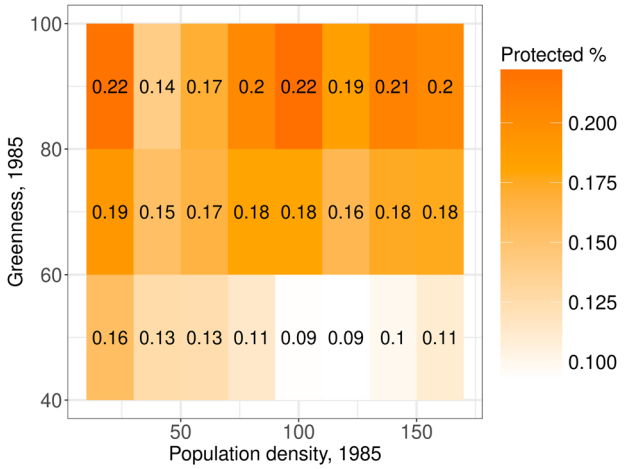


FIGURE 2

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 m. Grid cells are considered protected if any non-zero share of the grid cell is 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.

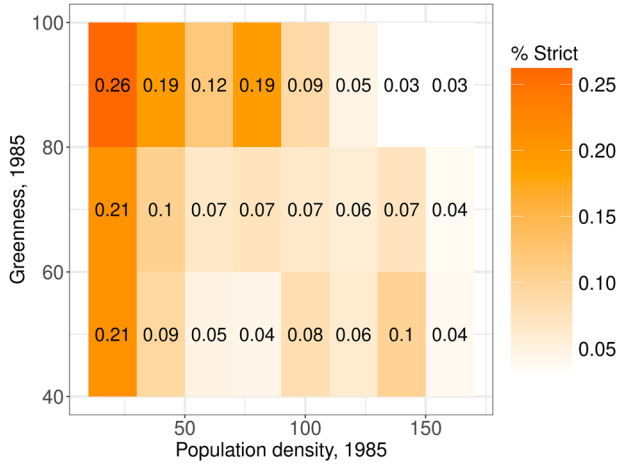


FIGURE 3

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 (darker) values of the strictness index.

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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

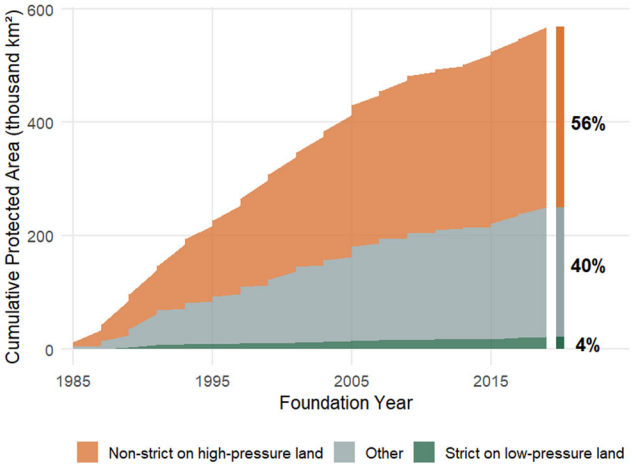


FIGURE 4
Cumulative share of land protection by strictness and land pressure

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 and 2019.

In summary, the EU 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, sparsely-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 $\mathbb{T} = \{1985, 1987, \ldots, 2019\}$ 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: $D_{it} = \mathbb{1}[t \geq g(i)]$.

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:

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

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 and 1989. To 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.

27. We use a difference-in-differences approach instead of a regression discontinuity design (RDD) as in Turner *et al.* (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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 *et al.*, 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:

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

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:

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

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:

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

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:

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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

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.* (2025) 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:

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

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

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 neighbours. 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*

7.1.1. 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

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

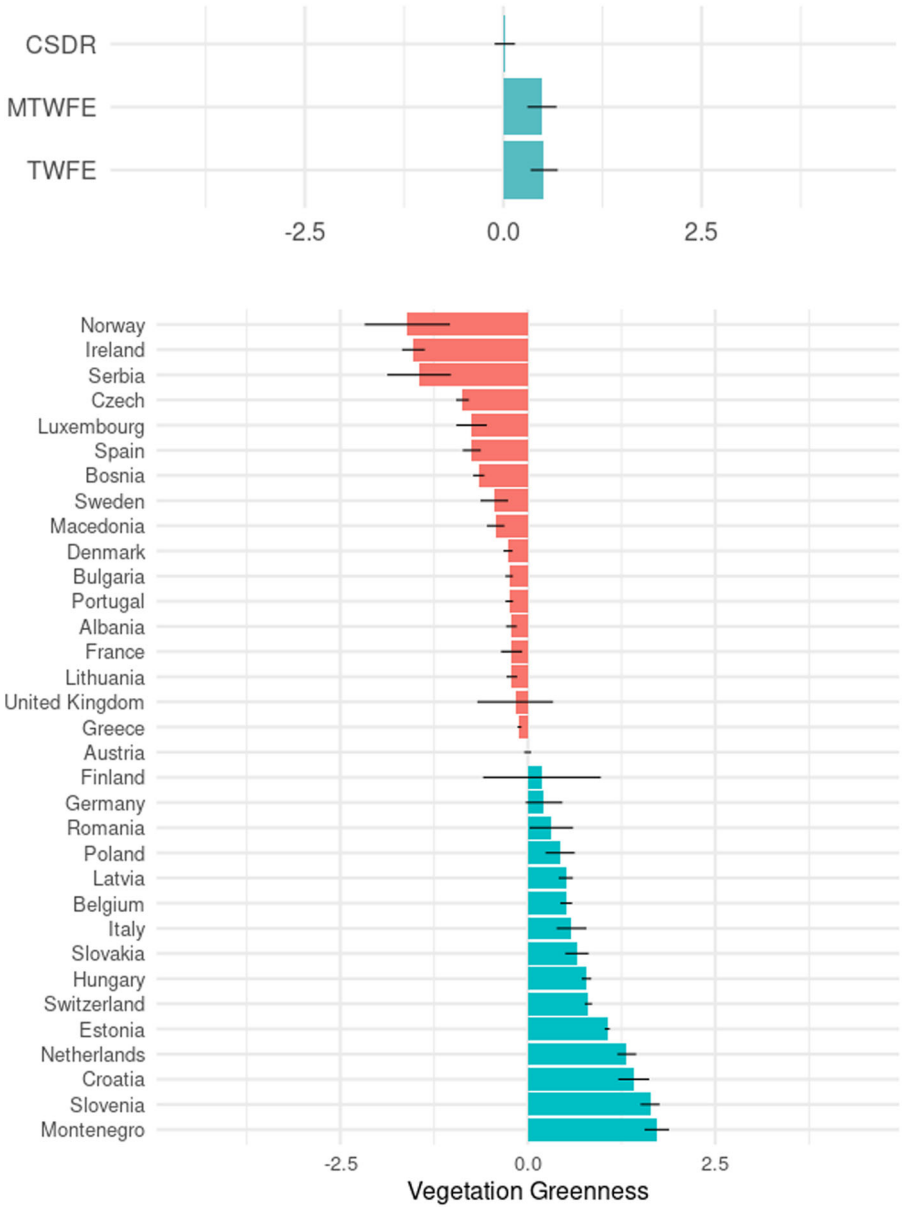


FIGURE 5
Treatment effects on NDVI, EU-wide and by country

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. Horizontal black lines indicate 95% confidence intervals. Greenness varies from 0 to 100. Appendix Table A.10 presents the regression results.

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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 of the 32 countries, the treatment effect is less than a 0.5 point change in the NDVI index, which is small relative to the 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 (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.

7.1.2. 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 3, 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 estimates θ_{cg} in an early-treated and late-treated group and test for a difference in the treatment effect size between

33. Because our results indicate a zero treatment effect of protection on greenness, we are not concerned about spillovers. Robalino *et al.* (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.

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).

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

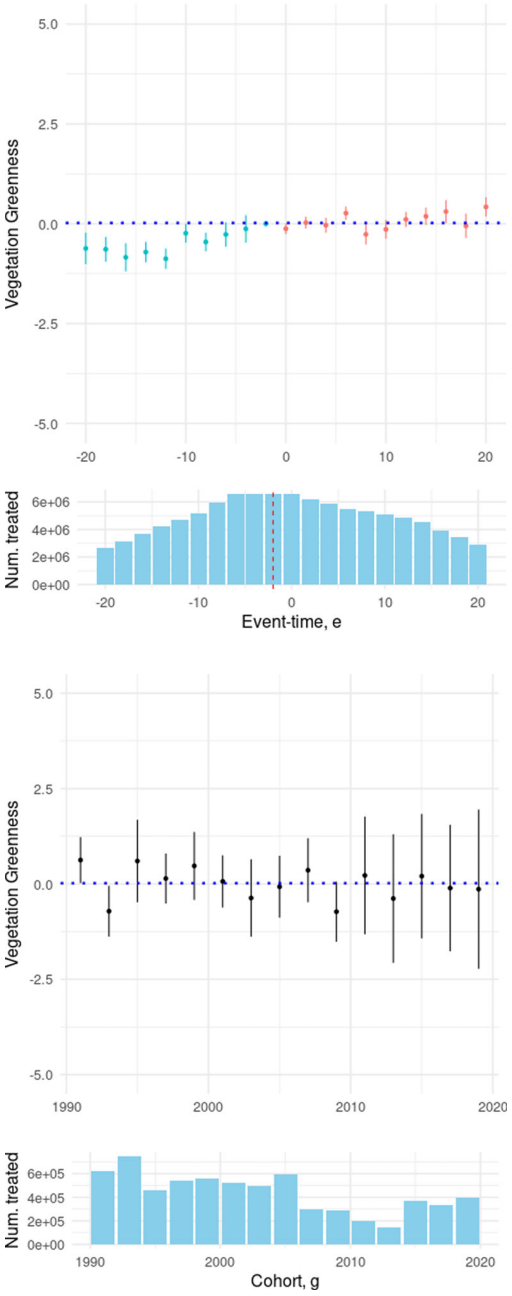


FIGURE 6
Treatment effects on NDVI over event time and by cohort

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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

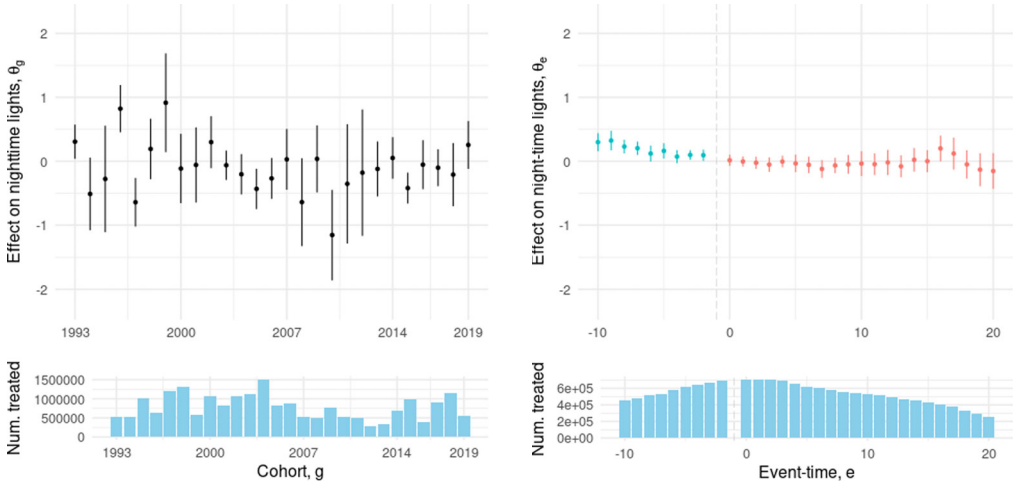


FIGURE 7
Treatment effects on night-time lights over event time and by cohort

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.

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.

7.1.3. Average treatment effects on other outcomes. Having established a mostly-precise zero 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).

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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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	84.7	1.8	13.0	0.0	0.5	0.0
	Forest	% row	0.2	97.5	2.1	0.0	0.1	0.0
	Grassland	% row	3.8	10.9	85.0	0.0	0.3	0.0
	Other	% row	0.0	0.1	0.1	99.9	0.0	0.0
	Settlement	% row	1.1	0.3	0.7	0.0	97.9	0.0
	Water	% row	0.2	0.3	0.6	0.0	0.2	98.6
Treated	Cropland	% row	80.2	2.9	16.3	0.0	0.5	0.1
	Forest	% row	0.3	98.1	1.4	0.0	0.1	0.0
	Grassland	% row	5.2	9.5	85.0	0.0	0.3	0.0
	Other	% row	0.0	0.0	0.1	99.9	0.0	0.0
	Settlement	% row	1.5	0.4	0.6	0.0	97.5	0.0
	Water	% row	0.0	0.2	0.1	0.0	0.1	99.5
	Total (2010)	% row	27.3	35.6	28.8	1.4	4.1	2.8

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.

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%-point 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.

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 to 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.* (2025), 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 Figure 8a and b, 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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

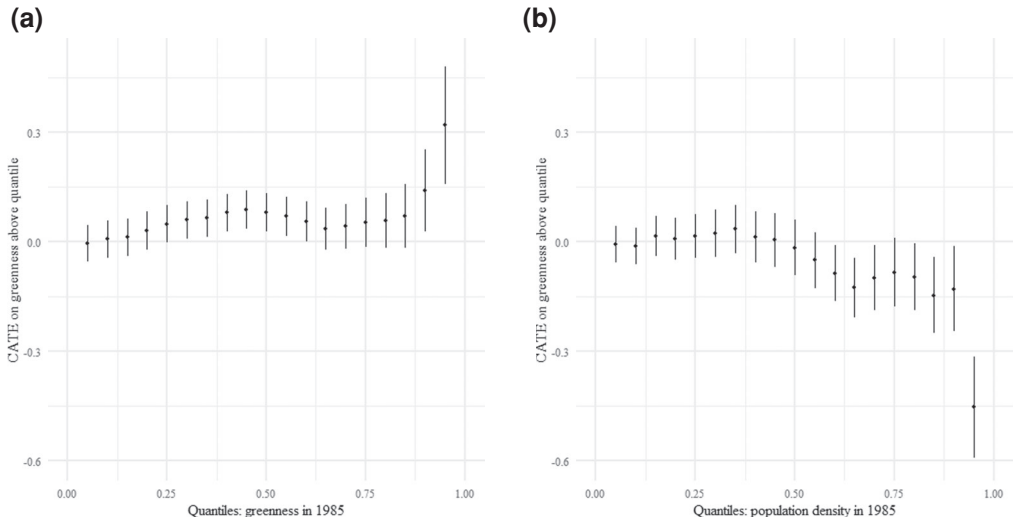


FIGURE 8

Distribution of the conditional average treatment effect for initial greenness and population density. (a) Starting average NDVI 1985–9. (b) Population density in 2000

Notes: Both panels plot estimates of the conditional average treatment effect across 20 ventiles of a key explanatory variable. Bar at ventile $q \in \{0.05, 0.1, \ldots, 0.95\}$ 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.

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 effect 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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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.

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 NDVI 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 example, while enforcement of land-use policy is often imperfect, as seen in the Amazon rainforest (Keles *et al.*, 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.

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 that this limited heterogeneity is 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.

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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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) is weighted by the inverse variance of country-level CATE and ATEs, assuming zero covariance. *Indicates signifiance at the 90% confidence level.

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 from 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

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

(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 on 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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

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 they do suggest that Europe's current land-protection regime has been poorly targeted and has favoured low-impact designations in economically contested areas.

Acknowledgments. We thank Adam Streff, Danial Syed, and Xiaomeng Wu for excellent research assistance. 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, Matthew Wibbenmeyer, and numerous seminar and conference participants 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 no. 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 EU nor the granting authority can be held responsible for them. Reynaert acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir programme). 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 Programme, and Analytics at Wharton for generous support.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability

The data underlying this article, together with the replication code, are available at: https://doi.org/10.5281/zenodo.18849775. Data for which redistribution permission was obtained are included in the replication package. Some data are owned by third parties and could not be redistributed; for these datasets, the replication package provides detailed instructions on how they can be accessed from the original providers.

REFERENCES

ABMAN, R. (2018), "Rule of law and Avoided Deforestation from Protected Areas", *Ecological Economics*, **146**, 282–289.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

ALIX-GARCIA, J. and MILLIMET, D. (2022), "Remotely Incorrect? Accounting for Nonclassical Measurement Error in Satellite Data on Deforestation", *Journal of the Association of Environmental and Resource Economists*, **10**, 1335–1367.
ANDAM, K. S., FERRARO, P. J., PFAFF, A., *et al.* (2008), "Measuring the Effectiveness of Protected Area Networks in Reducing Deforestation", *Proceedings of the National Academy of Sciences*, **105**, 16089–16094.
ARONOFF, D. and RAFEY, W. (2026), "Conservation Priorities and Environmental Offsets: Markets for Florida Wetlands", *American Economic Review*, **116**, 1723–1764.
ASPELUND, K. M. and RUSSO, A. (Forthcoming), "Additionality and Asymmetric Information in Environmental Markets: Evidence from Conservation Auctions", *American Economic Review*.
ASSUNÇÃO, J., GANDOUR, C. and ROCHA, R. (2023), "DETER-ing Deforestation in the Amazon: Environmental Monitoring and Law Enforcement", *American Economic Journal: Applied Economics*, **15**, 125–156.
ASSUNÇÃO, J., MCMILLAN, R., MURPHY, J., *et al.* (2022), "Optimal Environmental Targeting in the Amazon Rainforest", *Review of Economic Studies*, **90**, 1608–1641.
AUFFHAMMER, M., FRANK, E., MCLAUGHLIN, D., *et al.* (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.
BÖRNER, J., SCHULZ, D., WUNDER, S., *et al.* (2020), "The Effectiveness of Forest Conservation Policies and Programs", *Annual Review of Resource Economics*, **12**, 45–64.
BAHRAMI, G., GUSTAFSON, M. and STEINER, E. (2024), "The Biodiversity Protection Discount", Available at SSRN.
BALBONI, C., BURGESS, R. and OLKEN, B. A. (2026), "The Origins and Control of Forest Fires in the Tropics", *The Review of Economic Studies*, **93**, 2348–2389. "The Origins and Control of Forest Fires in the Tropics", (Working Paper) https://doi.org/10.1093/restud/rdaf088.
BARBIER, E. B., BURGESS, J. C. and GRAINGER, A. (2010), "The Forest Transition: Towards a More Comprehensive Theoretical Framework", *Land Use Policy*, **27**, 98–107.
BERTRAND, M., DUFLO, E. and MULLAINATHAN, S. (2004), "How Much Should We Trust Differences-In-Differences Estimates?", *Quarterly Journal of Economics*, **119**, 249–275.
BICKENBACH, F., BODE, E., NUNNENKAMP, P., *et al.* (2016), "Night Lights and Regional GDP", *Review of World Economics*, **152**, 425–447.
BONTHOUX, S., LEFEVRE, S., HERRAULT, P.-A., *et al.* (2018), "Spatial and Temporal Dependency of NDVI Satellite Imagery in Predicting Bird Diversity over France", *Remote Sensing*, **10**, 1136.
BURGESS, R., COSTA, F. J. and OLKEN, B. A. (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, R., HANSEN, M., OLKEN, B. A., *et al.* (2012), "The Political Economy of Deforestation in the Tropics", *Quarterly Journal of Economics*, **127**, 1707–1754.
CALLAWAY, B. and SANT'ANNA, P. H. (2021), "Difference-in-Differences with Multiple Time Periods", *Journal of Econometrics*, **225**, 200–230.
CHEN, X. and NORDHAUS, W. (2011), "Using Luminosity Data as a Proxy for Economic Statistics", *Proceedings of the National Academy of Sciences*, **108**, 8589–8594.
CHENG, A., SIMS, K. R. and YI, Y. (2023), "Economic Development and Conservation Impacts of China's Nature Reserves", *Journal of Environmental Economics and Management*, **121**, Article 102848.
CHERNOZHUKOV, V., DEMIRER, M., DUFLO, E., *et al.* (2025), "Fisher–Schultz Lecture: Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, With an Application to Immunization in India", *Econometrica*, **93**, 1121–1164.
DE CHAISEMARTIN, C. and D'HAULTFŒUILLE, X. (2020), "Two-way Fixed Effects Estimators with Heterogeneous Treatment Effects", *American Economic Review*, **110**, 2964–2996.
DONALDSON, D. and STOREYGARD, A. (2016), "The View from Above: Applications of Satellite Data in Economics", *Journal of Economic Perspectives*, **30**, 171–198.
DRUCKENMILLER, H. and TAYLOR, C. A. (2022), "Wetlands, Flooding, and the Clean Water act", *American Economic Review*, **112**, 1334–1363.
EINHORN, C. (2022), "Nearly Every Country Signs on to a Sweeping Deal to Protect Nature", *The New York Times*.
European Commission (2011), "Our Life Insurance, Our Natural Capital: An Eu Biodiversity Strategy to 2020" https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52011DC0244. Document 52011DC0244, Accessed 22 May 2024.
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" https://www.europarl.europa.eu/doceo/document/A-9-2023-0220_EN.html. Document A9-0220/2023, Accessed 22 May 2024.
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" https://eur-lex.europa.eu/eli/dir/2009/147/oj. Accessed 22 May 2024.
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.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

FRANK, E. G. (2024), "The Economic Impacts of Ecosystem Disruptions: Costs from Substituting Biological Pest Control", *Science*, **385**, Article eadg0344.
FUCHS, R., VERBURG, P. H., CLEVERS, J. G., *et al.* (2015), "The Potential of Old Maps and Encyclopaedias for Reconstructing Historic European Land Cover/Use Change", *Applied Geography*, **59**, 43–55.
GIBSON, J., OLIVIA, S., BOE-GIBSON, G., *et al.* (2021), "Which Night Lights Data Should We Use in Economics, and Where?", *Journal of Development Economics*, **149**, Article 102602.
GOODMAN-BACON, A. (2021), "Difference-in-differences with Variation in Treatment Timing", *Journal of Econometrics*, **225**, 254–277. Themed Issue: Treatment Effect 1.
HANCOCK, A. and JONES, S. (2024), "Eu Approves Nature Law in Face of Austrian Legal Threat", *Financial Times* https://www.ft.com/content/7175b191-04a2-406f-8149-cc1974d5b308. Accessed 11 July 2025.
HARSTAD, B. (2023), "The Conservation Multiplier", *Journal of Political Economy*, **131**, 1731–1771.
HARSTAD, B. and MIDEKSA, T. K. (2017), "Conservation Contracts and Political Regimes", *Review of Economic Studies*, **84**, 1708–1734.
HASSAN SOUHEIL, D. B., GERMAIN, L. and DOUILLET, R. (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, B., ROGERSON, R. and VALENTINYI, A. (2014), "Growth and Structural Transformation", *Handbook of Economic Growth*, **2**, 855–941.
HOBI, M. L., DUBININ, M., GRAHAM, C. H., *et al.* (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, A. (2026), "Coordination and Commitment in International Climate Action: Evidence from Palm Oil", *Econometrica*, **94**, 1–33. https://doi.org/10.3982/ECTA20608.
JOPPA, L. N. and PFAFF, A. (2009), "High and far: Biases in the Location of Protected Areas", *PLoS One*, **4**, 1–6.
—— —— (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.
KEDWARD, K., RYAN-COLLINS, J. and CHENET, H. (2023), "Biodiversity Loss and Climate Change Interactions: Financial Stability Implications for Central Banks and Financial Supervisors", *Climate Policy*, **23**, 763–781.
KELES, D., PFAFF, A. and MASCIA, M. (2022), "Does the Selective Erasure of Protected Areas Raise Deforestation in the Brazilian Amazon?", *Journal of the Association of Environmental and Resource Economists*, **10**, 1121–1147.
KNITTEL, C. R. and STOLPER, S. (2021), "Machine Learning about Treatment Effect Heterogeneity: The Case of Household Energy Use", *AEA Papers and Proceedings*, **111**, 440–444.
LASSITER, A. (2022), "Identifying Causal Changes in Landscape Greenness with Very High-Resolution Airborne Multispectral Imagery and a Panel Data Model", *Urban Forestry & Urban Greening*, **67**, Article 127380.
LEITE MARIANTE, G. and SALAZAR RESTREPO, V. (2024), "Does Conservation Work in General Equilibrium?" (Working Paper) https://vsalazarr.github.io/jobmarketpaper/.
LI, X., ZHOU, Y., ZHAO, M., *et al.* (2020), "A Harmonized Global Nighttime Light Dataset 1992–2018", *Scientific Data*, **7**, 168.
LIANG, Y., RUDIK, I. and ZOU, E. (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, S. L., CAZALIS, V., DUDLEY, N., *et al.* (2020), "Area-Based Conservation in the Twenty-First Century", *Nature*, **586**, 217–227.
MCDONALD, G., BONE, J., COSTELLO, C., *et al.* (2024), "Global Expansion of Marine Protected Areas and the Redistribution of Fishing Effort", *PNAS*, **121**, Article e2400592121.
METRICK, A. and WEITZMAN, M. L. (1998), "Conflicts and Choices in Biodiversity Preservation", *Journal of Economic Perspectives*, **12**, 21–34.
NIETO, S., FLOMBAUM, P. and GARBULSKY, M. F. (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, A., ROBALINO, J., SANDOVAL, C., *et al.* (2015), "Protected Areas' Impacts on Brazilian Amazon Deforestation: Examining Conservation—Development Interactions to Inform Planning", *PLoS One*, **370**, Article e0129460.
RAMBACHAN, A. and ROTH, J. (2023), "A More Credible Approach to Parallel Trends", *Review of Economic Studies*, **90**, 2555–2591.
REYNAERT, M., SOUZA-RODRIGUES, E. and VAN BENTHEM, A. A. (2023), "The Environmental Impacts of Protected Area Policy," *Regional Science and Urban Economics*, forthcoming https://souza-rodrigues.economics.utoronto.ca/wp-content/uploads/RSUE_land_protection-Paper.pdf.
RIBEIRO, I., PROENÇA, V., SERRA, P., *et al.* (2019), "Remotely Sensed Indicators and Open-Access Biodiversity Data to Assess Bird Diversity Patterns in Mediterranean Rural Landscapes", *Scientific Reports*, **9**, Article 6826.
RICO-STRAFFON, J., WANG, Z., PANLASIGUI, S., *et al.* (2023), "Forest Concessions and Eco-Certification in the Peruvian Amazon: Deforestation Impacts of Logging Rights and Restrictions", *Journal of Environmental Economics and Management*, **118**, Article 102780.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026

ROBALINO, J. A., PFAFF, A. and VILLALOBOS, L. (2017), "Heterogeneous Local Spillovers from Protected Areas in Costa Rica", *Journal of the Association of Environmental and Resource Economists*, **4**, 795–820.
ROTH, J. and SANT'ANNA, P. H. (2023), "When is Parallel Trends Sensitive to Functional Form?", *Econometrica: Journal of the Econometric Society*, **91**, 737–747.
SCHLENKER, W. and ROBERTS, M. (2009), "Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change", *Proceedings of the National Academy of Sciences*, **106**, 15594–15598.
SIMS, K. R. E. (2010), "Conservation and Development: Evidence from Thai Protected Areas", *Journal of Environmental Economics and Management*, **60**, 94–114.
SIMS, K. R. E. and ALIX-GARCIA, J. M. (2017), "Parks versus PES: Evaluating Direct and Incentive-Based Land Conservation in Mexico", *Journal of Environmental Economics and Management*, **86**, 8–28.
SOUZA-RODRIGUES, E. (2019), "Deforestation in the Amazon: A Unified Framework for Estimation and Policy Analysis", *Review of Economic Studies*, **86**, 2713–2744.
SUN, L. and ABRAHAM, S. (2021), "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects", *Journal of Econometrics*, **225**, 175–199. Themed Issue: Treatment Effect 1.
TAYLOR, M. S. and WEDER, R. (2024), "On the Economics of Extinction and Possible Mass Extinctions", *Journal of Economic Perspectives*, **38**, 237–260.
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" https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A31992L0043. Accessed 22 May 2024.
—— (1997), "Council Directive 79/409/eec of 2 April 1979 on the Conservation of Wild Birds" https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A31979L0409. Accessed 22 May 2024.
TORCHIANA, A. L., ROSENBAUM, T., SCOTT, P. T., *et al.* (2023), "Improving Estimates of Transitions from Satellite Data: A Hidden Markov Model Approach", *The Review of Economics and Statistics*, forthcoming https://www.ptscott.com/papers/hmm_error_correction.pdf.
TURNER, M. A., HAUGHWOUT, A. and VAN DER KLAAUW, W. (2014), "Land use Regulation and Welfare", *Econometrica: Journal of the Econometric Society*, **82**, 1341–1403.
VILLASENOR-DERBEZ, J. C., COSTELLO, C. and PLANTINGA, A. J. (2024), "A Market for 30 × 30 in the Ocean", *Science*, **384**, 1177–1179.
WAGER, S. and ATHEY, S. (2018), "Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests", *Journal of the American Statistical Association*, **113**, 1228–1242.
WEITZMAN, M. L. (1998), "The Noah's ark Problem", *Econometrica: Journal of the Econometric Society*, **66**, 1279–1298.
WINKLER, K., FUCHS, R., ROUNSEVELL, M., *et al.* (2021), "Global Land use Changes are Four Times Greater Than Previously Estimated", *Nature Communications*, **12**, Article 2501.
WOLF, C., LEVI, T., RIPPLE, W. J., *et al.* (2021), "A Forest Loss Report Card for the World's Protected Areas", *Nature Ecology & Evolution*, **5**, 520–529.

Downloaded from academic.oup.com/restud/advance-article/doi/10.1093/restud/rdag080/8746820 by California Institute of Technology user on 27 August 2026