Assistive AI and Data Manipulation in Citizen Science

Xinming Du*

March 8, 2026

Abstract

AI tools designed to help legitimate users may simultaneously lower the cost of fabrication. When the resulting fakes overwhelm fixed reviewer capacity, detection falls and fabrication rises further, a process I call a deterrence-erosion multiplier. I study this mechanism using eBird, the world's largest biodiversity database, exploiting the 2021 AI boom in Brazil. A triple difference design shows that rare species records in ecologically implausible areas rise by 22 percent while expert review rates fall by 9.7 percentage points. The surge concentrates at night, on weekends, and in areas with mobile coverage, pointing to casual AI-app users rather than professional guides. About 2,800 excess suspect records entered eBird's conservation-relevant database after 2021, including records of 60 threatened species whose occurrence data inform IUCN Red List threat classifications. When verification is costly and self-reports are valuable, assistive technology that lowers skill barriers can erode data quality through a deterrence-erosion multiplier.

JEL Classifications: D82, Q57, C55, D04
Keywords: Assistive technology, data quality, verification capacity, citizen science

*Department of Economics, National University of Singapore. Email: xdu@nus.edu.sg.

1 Introduction

Citizen science platforms, in which volunteer members of the public submit observational data, aggregate millions of self-reported observations that underpin research and policy worldwide. eBird, the largest biodiversity citizen science database, contains over one billion bird observations and feeds directly into IUCN Red List assessments, national conservation planning, and hundreds of published studies (Sullivan et al., 2014, 2009). In the United States, 52% of state heritage programs use eBird data in regulatory environmental review (Callaghan et al., 2025). The EU's Common Bird Index and Grassland Butterfly Index rely on analogous citizen-collected monitoring data for biodiversity directives. Similar platforms, including iNaturalist for biodiversity, Galaxy Zoo for astronomy (Lintott et al., 2008), and community networks for air and water quality, form a growing infrastructure of crowd-sourced scientific data (Fraisl et al., 2022; Bonney et al., 2014; Callaghan et al., 2021a). The reliability of this infrastructure depends on verification systems that are inherently capacity-constrained. Expert reviewers are volunteers whose numbers have remained roughly constant since 2017 despite a near-quadrupling of annual submissions.

Artificial intelligence is transforming how users interact with these platforms. Beginning around 2021, rapid advances in image recognition, generative AI, and image synthesis tools dramatically lowered the cost of producing convincing species identifications and fabricated photographic evidence.[1] Before this AI boom, submitting a plausible rare species record required genuine ornithological expertise or a real photograph. After 2021, anyone with a phone can generate a species name from an AI identification app, produce synthetic wildlife images using generative AI, and submit plausible-looking records at near-zero effort. Did the AI boom inadvertently increase data manipulation on citizen science platforms?

The question connects to a well-documented pattern in environmental economics. When monitoring is imperfect, agents strategically manipulate the data. Chinese cities game air pollution readings around regulatory thresholds (Chen et al., 2012; Ghanem and Zhang,

[1]The 2021 treatment date is somewhat discretionary: AI identification tools existed before 2021, and generative image models continued improving after. I use 2021 as the primary treatment date because it marks the inflection point in both model capability and mass adoption. Robustness checks using 2022 as the alternative treatment date yield similar coefficients (Table S5), and the event study with 2022 as the reference year (Figure S4) confirms that the break occurs at AI adoption timing.

2014), U.S. pollution sources increase emissions on unmonitored days (Zou, 2021), and Chinese firms strategically dilute wastewater to meet discharge standards (He et al., 2020). In each case, regulated firms exploit gaps in government monitoring to misreport. I show that the same strategic logic operates in a fundamentally different institutional setting. In citizen science platforms, the the reporters are individual volunteers, the monitors are fellow volunteers with fixed capacity, and the tool that enables manipulation was designed to assist rather than deceive. The dual-use problem arises not from deliberate evasion of regulation but as an unintended byproduct of lowering skill barriers. More broadly, AI-mediated content generation is expanding rapidly across domains that depend on the credibility of user-generated reports, from fake product reviews (He et al., 2022) to political misinformation (Goldstein et al., 2023) to synthetic images that fool expert judges (Murabayashi, 2023), and citizen science platforms face an analogous vulnerability wherever verification is costly and capacity-constrained.

Brazil is an ideal setting for studying this problem. The country harbors about 1,900 confirmed bird species, more than any other country outside Colombia, and its remaining rainforest, cerrado, and Atlantic Forest fragments constitute some of the most biodiverse ecosystems on earth. Brazilian biodiversity data are consequential for policy: ICMBio, the federal conservation agency, uses occurrence databases including eBird for Red List assessments of more than 200 threatened bird species, and international organizations such as the IUCN incorporate Brazilian records into global species range maps that determine conservation funding and land-use regulations (Rodrigues et al., 2006). At the same time, enforcement capacity is limited. Even forest management relies heavily on satellite remote sensing because on-the-ground monitoring is infeasible across the Amazon's vast territory (Burgess et al., 2012; Assunção et al., 2015). Verifying whether a reported bird sighting is genuine depends on a fixed corps of roughly 150–200 volunteer reviewers. To date, I am not aware of any systematic check on the integrity of eBird data in Brazil.

I study this question using a triple difference (DDD) design applied to eBird Brazil, which covers 5,572 municipalities over 2017-2025. The unit of observation is a municipality-year-rarity cell. The first difference exploits the 2021 AI boom as the treatment date. The second difference compares rare species, defined as the bottom tercile of pre-2021 observation

counts and computed before the treatment to avoid circularity, to common species. The third difference uses municipality-level forest cover from MapBiomas Collection 9 (Souza et al., 2020), because rare forest-dependent species cannot plausibly occur in municipalities with little remaining forest. Low-forest municipalities serve as natural "falsification zones" where rare species records are intrinsically suspect. The identifying assumption is that, absent AI tools, trends in rare versus common species records would have been parallel across low- and high-forest municipalities. I validate this assumption with event studies showing flat pre-trends and placebo tests at false treatment dates.

The Harpy Eagle (*Harpia harpyja*), one of the world's largest raptors, is strictly dependent on continuous lowland rainforest canopy. In the pre-AI period of 2017-2020, there were zero eBird records of this species in municipalities with less than 5% forest cover. After the AI boom, multiple new "first appearances" emerged in deforested municipalities in Maranhão and Piauí, areas with less than 2% forest cover and no contiguous habitat. These records survived eBird's automated filters and were never flagged for expert review. This is not an isolated case. Novel rare-species and low-forest municipality pairs increased from 262 per year in the pre-AI period, 1,047 over four years from 2017 to 2020, to 315 per year post-AI, 1,574 over five years from 2021 to 2025, a 20 percent annual increase.[2]

The mechanism operates through what I call a *deterrence-erosion multiplier*. AI reduces fabrication costs, the additional fakes dilute the fixed-capacity review system, the resulting drop in detection probability encourages still more fabrication, and the direct effect is amplified by a multiplier greater than one. The main results are as follows. First, rare species records in ecologically implausible areas increase by about 22 percent post-AI, while the share of these records flagged for expert review falls by 9.7 percentage points, nearly double the sample mean review rate of 5.2 percent. Event study plots show no pre-trends before 2021 and a sharp break at the AI boom date. The continuous forest specification confirms that the manipulation signal strengthens monotonically as ecological implausibility increases.

Second, I document patterns consistent with AI-enabled fabrication rather than genuine

[2]The raw totals (1,047 vs. 1,574) are not directly comparable because the pre-AI and post-AI periods differ in length. The per-year figures (262 vs. 315) normalize by period length and provide the appropriate comparison.

range expansion. Night-time submissions of rare species, between 22:00 and 06:00 when real birding is impossible, surge by 23 percent post-AI. The review rate decline concentrates in municipalities with high mobile coverage, where AI tool adoption is feasible. A mirror-image decomposition shows that the reviewed share falls by 9.7 percentage points while the unreviewed share rises by 9.7 percentage points; records escape the review system entirely rather than passing through it. The review rate decline is concentrated among single-rare-species checklists. Marginal fakes evade review, while egregious multi-rare-species checklists are still caught.

Third, mechanism tests distinguish between competing explanations. The effect is not driven by reviewer dilution. Common species review rates are unchanged post-AI, and controlling for common-species volume barely moves the rare-species coefficient. The "better identification" interpretation, that AI tools help novices correctly identify rare species and reduce the need for review, predicts higher review rates, the opposite of what I find. It also cannot explain night-time submissions or the concentration in low-forest areas where rare species cannot occur.

Fourth, I identify who manipulates and quantify the contamination risk. The effect concentrates in municipalities with birding tourism infrastructure, but the weekend pattern rules out organized guide fraud. The weekend premium is large in non-tourism areas but dampened in tourism areas. Charismatic species are not differentially targeted. The profile points to casual AI-app users, weekend hobbyists with mobile coverage. Applying the estimated 22% excess to the 15,646 rare species records in low-forest municipalities post-2021 yields roughly 2,800 suspect records, including threatened species used in IUCN Red List assessments and Brazil's national Red List.

One interpretive concern is whether the results reflect genuine ecological changes rather than fabrication. I address this by showing event study patterns where a sharp break at 2021 is inconsistent with gradual range shifts, forest cover in treatment municipalities continued to decline, and the night-time submission surge cannot reflect real observation.

This paper contributes to four literatures. First, a large body of work studies how AI transforms content production. Brynjolfsson et al. (2025) provide prominent recent evidence that generative AI raises worker productivity, showing large gains for customer service

agents. The same cost reduction that improves legitimate performance in their setting enables fabrication in mine. Other work focuses on AI-generated misinformation in commercial and political settings (Acemoglu, 2024, 2021; Goldstein et al., 2023; Allcott and Gentzkow, 2017). A parallel literature documents strategic misreporting in online platforms, including fake product reviews on Yelp and Amazon, promotional manipulation by hotels, and opinion fraud in online forums (Luca and Zervas, 2016; Mayzlin et al., 2014; He et al., 2022; Dellarocas, 2006). Similar patterns appear in offline settings such as teacher cheating on standardized tests and strategic misreporting of trade values (Jacob and Levitt, 2003; Fisman and Wei, 2004). These studies focus on agents who deliberately set out to deceive, from hotel managers commissioning fake reviews to firms falsifying emissions data. The marginal fabricator in my setting is different. A casual user who points an AI app at a bird, accepts the species name it returns, and submits the record without recognizing the identification as implausible. The fabrication is a byproduct of AI-mediated misidentification, not a deliberate act of fraud. This distinction matters because it implies that the standard policy responses to manipulation, such as increasing penalties and improving detection, may be less effective when the manipulators do not perceive themselves as cheating.

Second, the monitoring and enforcement literature, from Becker (1968) and Stigler (1970) through recent field experiments, has studied how auditing probability, auditor quality, wages, and penalties deter misconduct (Harrington, 1988; Gray and Deily, 1996; Olken, 2007; Duflo et al., 2013, 2018; Di Tella and Schargrodsky, 2003; Polinsky and Shavell, 1979; Mookherjee and Png, 1995). This work models enforcement from the supply side, how monitoring resources shape deterrence. I document a demand-side channel. A technology shock that floods the verification system with harder-to-detect fakes, endogenously eroding deterrence even when enforcement effort is unchanged. The multiplier formalized in Proposition 1 shows that the feedback between fabrication volume and review capacity creates amplification beyond the direct cost reduction. Unlike commercial platforms with scalable detection algorithms, the verification system studied here is a scarce volunteer-staffed public good with fixed capacity, so the capacity constraint generates amplification absent from existing models.

Third, environmental economists have extensively documented that regulated agents

strategically manipulate environmental data when monitoring is imperfect, gaming air pollution readings around regulatory thresholds, increasing emissions on unmonitored days, and diluting wastewater to meet discharge standards (Chen et al., 2012; Ghanem and Zhang, 2014; Zou, 2021; He et al., 2020; Greenstone et al., 2022). These studies focus on regulated firms manipulating their own emissions data under government monitoring. I extend this logic to citizen science platforms, where the reporters are individual users rather than firms, the monitoring is crowd-sourced rather than regulatory, and the AI tool that enables manipulation was designed to assist rather than deceive.

Fourth, the citizen science data quality literature focuses predominantly on unintentional errors arising from variation in observer expertise, developing protocols, statistical guidelines, and proficiency metrics to address observer skill heterogeneity (Sullivan et al., 2009, 2014; Kosmala et al., 2016; Johnston et al., 2021; Kelling et al., 2015). I document a qualitatively different problem, strategic manipulation at scale, enabled by a technology shock that lowered the barrier to producing plausible-but-false records. Conservation policy increasingly depends on occurrence databases for species range estimation and threat assessment (Ferraro and Pattanayak, 2006; Rodrigues et al., 2006; Jetz et al., 2012), and strategic behavior by landowners already undermines conservation outcomes in settings like the Brazilian Amazon (Burgess et al., 2012, 2024). I identify a new threat vector, strategic data generation that corrupts the information base on which conservation decisions rest. On the platform design side, the existing literature on reputation systems and quality assurance in two-sided markets (Tadelis, 2016; Filippas et al., 2022; Einav et al., 2016; Nosko and Tadelis, 2015) does not consider a dual-use dilemma where the same tool that increases participation and lowers barriers to entry simultaneously degrades the platform's core public-good data asset.

From the policy perspective, the findings highlight an underappreciated tension in the deployment of assistive AI for environmental monitoring. Governments and international organizations are increasingly relying on citizen-generated biodiversity data to inform conservation policy. Brazil's ICMBio uses eBird records in Red List assessments for more than 200 threatened bird species, and the IUCN integrates citizen science occurrence data into global species range maps that determine funding allocations and land-use regulations (Rodrigues et al., 2006; Jetz et al., 2012). The European Union's Biodiversity Strategy for 2030

(COM/2020/380) explicitly calls for expanding citizen science monitoring networks. As these policy pipelines deepen their dependence on crowd-sourced data, the integrity of the underlying databases becomes a first-order policy concern. The AI boom generated thousands of suspect records that propagate into conservation assessments without correction.

The same structural vulnerability, valuable self-reports, costly verification, and fixed reviewer capacity, extends to environmental self-monitoring by regulated firms (Duflo et al., 2013), community-based water quality testing, participatory mapping, and crowd-sourced pollution reporting. Deploying assistive AI without simultaneously scaling verification infrastructure is equivalent to lowering the expected penalty for misconduct.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting. Section 3 presents a stylized model of AI-enabled data fabrication with capacity-constrained verification. Section 4 describes the data. Section 5 lays out the empirical strategy. Section 6 presents the main results, heterogeneity analysis, and robustness checks. Section 7 investigates mechanisms. Section 8 concludes.

2 Background

2.1 eBird and the AI Boom

eBird, operated by the Cornell Lab of Ornithology, is the world's largest biodiversity citizen science platform (Sullivan et al., 2009, 2014). Users submit checklists reporting species observed at a given location, date, and time, along with counts and observation effort. The cumulative database exceeds one billion observations contributed by more than 900,000 observers. eBird data are used extensively in ecology, conservation biology, and increasingly in economics. Liang et al. (2020) use eBird data to show that the Clean Air Act saved roughly 1.5 billion birds by improving air quality, and Liang et al. (2025) link economic activity to biodiversity decline across the United States using the same platform. Callaghan et al. (2021b) derive global population estimates for 9,700 bird species from eBird records.

eBird employs a tiered data quality system. At the first stage, automated filters flag records that are unusual for the given location, date, and species. Each species in each

region has a filter threshold, a maximum expected count for a given date range, set by regional experts based on historical observation frequencies. A record triggers a flag if the reported species has never or rarely been observed in that region during that time of year, or if the reported count exceeds the threshold. Crucially, these filters operate on species-location-date combinations, not on the content of the submission itself. They cannot assess whether a photograph is genuine or AI-generated, nor whether the observer actually visited the reported location. A single rare-species record submitted from an established birding hotspot during the appropriate season may pass the filter entirely, even if the observation never occurred. At the second stage, flagged records are routed to regional volunteer expert reviewers, who examine the supporting evidence, including photographs, audio recordings, checklist details, and the observer's track record, and either confirm or invalidate the record. Only records that survive this process enter the public database used by researchers and policymakers. The review system is inherently capacity-constrained. Brazil has roughly 150–200 active regional reviewers, and this number has not scaled proportionally with the post-2021 submission surge.[3] As Kosmala et al. (2016) note, volunteer reviewer recruitment is a persistent bottleneck across citizen science platforms.

AI-powered species identification apps, such as the Merlin Bird ID app developed by the Cornell Lab, have been available since the mid-2010s, using deep convolutional neural networks trained on millions of labeled images and sound recordings.[4] These apps are free, integrate with eBird, and require only a smartphone with mobile connectivity. However, the relevant treatment is not any single app but the broader AI boom beginning around 2021. Advances in generative AI and image synthesis tools, including text-to-image models capable of producing photorealistic wildlife imagery, dramatically expanded the potential for fabrication. Before 2021, producing a convincing fake photograph of a rare species required specialized skills; after the AI boom, synthetic images can be generated from a text

[3]The 150–200 figure is based on the count of unique reviewer identifiers active in the eBird Brazil review system during any calendar year. Table 11 confirms empirically that reviewer workload rises one-for-one with the post-AI submission surge, with no offsetting increase in reviewer headcount.

[4]I use Merlin as an illustrative example of the broader class of AI identification tools, not to single it out as the source of fabrication. The Cornell Lab of Ornithology both operates eBird and develops Merlin; the findings in this paper reflect an unintended consequence of the broader AI boom, not a flaw specific to any one application. Other tools (Google Lens, iNaturalist's CV model, generative image generators) contribute to the same cost reduction.

prompt. The combination of AI identification tools with generative image capabilities created a qualitative shift in the cost of fabrication.[5] The timing of this AI boom was driven by global technology development, not by local conditions in Brazilian municipalities, providing exogeneity for identification.

The key economic feature is that the AI boom dramatically reduces the skill barrier to producing plausible species records. Before the AI boom, submitting a plausible record for a rare species, say a Harpy Eagle (*Harpia harpyja*), required either genuine expertise or a deliberately staged encounter with convincing photographic evidence. Misidentifying a common hawk as a Harpy Eagle would have been immediately transparent to any knowledgeable reviewer. After the AI boom, the same submission requires only a phone. Point the camera at any bird and let an AI app generate a species name, or produce a synthetic image of the target species using generative AI, then submit to eBird. Yet neither AI identification apps nor eBird's automated filters can distinguish between genuine field observations and AI-generated or AI-assisted fabrications. This creates what I call the dual-use dilemma, whereby tools designed to lower identification costs for genuine birders simultaneously lower fabrication costs for strategic reporters.

The phenomenon has attracted media attention. Audubon Magazine reported concerns that "as the popularity of AI-powered Sound ID grows, so do concerns about how imperfect AI impacts a trove of scientific data" (Sun, 2023). Mongabay documented cases where "AI-generated wildlife photos make conservation more difficult" by corrupting species databases (Butler, 2026). PetaPixel estimated that over 90% of wildlife images on some Facebook nature photography pages are now AI-generated (Growcoot, 2025).[6] These anecdotal accounts motivate a systematic empirical investigation.

[5] I cannot directly observe whether a given submitter used an AI identification app, a generative image tool, or neither. The research design identifies the combined effect of the post-2021 AI boom on record quality, not the contribution of any single tool. This is a feature of the identification strategy: the treatment is the aggregate cost reduction from the AI ecosystem, which is what matters for the deterrence-erosion mechanism.

[6] The Growcoot (2025) figure is a journalistic estimate based on a single Facebook group and should be treated as illustrative rather than systematic evidence. I cite it to establish that AI-generated wildlife imagery is a recognized phenomenon, not to claim a precise prevalence rate.

2.2 Forest Cover and Ecological Plausibility

Many rare bird species in Brazil are forest-dependent. Species such as the Black-fronted Piping-Guan (*Pipile jacutinga*) and the Seven-colored Tanager (*Tangara fastuosa*) require large tracts of contiguous forest for foraging, nesting, and territory maintenance. A rare forest-dependent species cannot plausibly occur in a municipality where deforestation has eliminated its habitat. I exploit this ecological constraint as a ground-truth measure of record implausibility.

Forest cover data come from MapBiomas Collection 9 (Souza et al., 2020), which provides annual land-use classification maps at 30-meter resolution for all of Brazil from 1985 to 2023. The dataset has been widely validated and is used extensively in economics research on deforestation and environmental policy (Balboni et al., 2023; Carreira et al., 2024; Assunção et al., 2015). Burgess et al. (2012) use similar satellite-based deforestation data to study the political economy of tropical forest loss.

I aggregate pixel-level forest classifications to the municipality level. Forest is defined as the union of forest formation, class 3, savanna formation, class 4, mangrove, class 5, and floodable forest, class 6. The forest share for municipality i in year t is the ratio of forest pixels to total pixels. I define "low forest" as below-median pre-treatment forest share averaged over 2017-2020. Low-forest municipalities are concentrated in the southern agricultural belt, the northeastern interior including the Caatinga, and the deforestation frontier of eastern Amazonia.

If AI tools enable fabrication, the manipulation signal should be strongest where genuine rare species records are least plausible, namely in low-forest municipalities. The continuous forest specification confirms that this relationship is monotonic. The coefficient on the triple interaction of post-AI, rarity, and ecological implausibility is 0.267, indicating that each percentage point reduction in forest cover is associated with a stronger manipulation signal.

2.3 Birding Tourism in Brazil

Brazil hosts the world's third-largest avifauna, with about 1,900 confirmed species. Bird-watching tourism is a significant economic activity in certain municipalities, particularly

those near biodiversity hotspots such as the Atlantic Forest remnants, the Pantanal, and the Amazon fringes. Local guides and ecotourism operators benefit directly from the diversity and rarity of species observed in their territory. eBird species lists function as de facto marketing tools for local guides: visiting birders routinely consult eBird hotspot pages to choose destinations, so a municipality with confirmed records of rare or endemic species attracts more paying visitors.

This economic structure creates incentives for strategic manipulation. If a local operator can inflate the species list for their municipality by adding rare species that would attract high-paying international birders, the return on a fabricated record is tangible. The model in Section 3 formalizes this through the private value parameter v. I exploit cross-municipality variation in pre-existing birding tourism intensity, measured by pre-2021 eBird checklist density and alternatively by a binary high-birding-tourism indicator based on the above-median cutoff, to test whether the manipulation signal concentrates where economic incentives are strongest.

3 Stylized Model

I use a simple framework extending Becker (1968) to organize the empirical predictions. The framework clarifies why AI's impact on data quality operates through a multiplier, because the direct cost reduction is amplified by the endogenous erosion of the review system. This amplification mechanism distinguishes the setting from standard models of strategic misreporting (Kartik, 2009; Crawford and Sobel, 1982) and connects to the monitoring literature's emphasis on the interaction between misconduct volume and enforcement capacity (Stigler, 1970; Polinsky and Shavell, 1979).

The model treats fabrication as a deliberate cost-benefit decision, but the empirical evidence points to a spectrum of behavior. At one end are strategic fabricators who knowingly submit false records and weigh detection risk against the private value of an accepted record. At the other end are casual misidentifiers who accept an AI app's output without verifying whether the identification is plausible and submit the record in good faith. The fabrication cost c_i can be reinterpreted to span both types. For deliberate cheaters, c_i is the effort cost

of constructing a convincing fake; for casual misidentifiers, c_i is the cognitive cost of verifying whether the AI identification is plausible. AI reduces c_i for both types. It lowers construction costs for cheaters and eliminates the verification step for misidentifiers. The key insight is that both types enter the review system and overwhelm fixed reviewer capacity, so the deterrence-erosion multiplier operates through the volume of implausible records regardless of the submitter's intent.

Consider municipalities indexed by habitat quality $h \in [0,1]$. Genuine rare species arrivals follow $\lambda(h)$ with $\lambda'(h) > 0$ and $\lambda(h)$ small for low h, meaning rare species are unlikely in areas with little habitat. Each municipality has a population of potential reporters. Reporter i faces a fabrication cost $c_i \sim F(c)$ drawn from a smooth distribution with density $f(c) > 0$. This cost captures the effort required to produce a convincing fake record, including obtaining a plausible photo, constructing a credible checklist, and submitting it through eBird's interface. AI, parameterized by $\Delta \geq 0$, uniformly reduces this cost to $\max(c_i - \Delta, 0)$. Before the AI boom when $\Delta = 0$, fabrication required ornithological expertise. After the AI boom when $\Delta > 0$, the skill barrier drops discretely.

An accepted rare species report yields private value $v > 0$ to the reporter, reflecting tourism revenue, community prestige, contribution to a municipality's species list, or conservation influence. The value v is heterogeneous across municipalities but taken as given by individual reporters.[7]

The platform has aggregate review capacity $K > 0$, representing the total number of records that volunteer expert reviewers can examine per period. This capacity is fixed in the short run.[8] The review rate is $\phi = \min(K/N, 1)$, where N is total rare reports across all municipalities. Reviewed fakes are detected with probability $\pi \in (0,1)$ and removed.[9]

[7] In principle, v could be endogenous to the species list: a municipality with many confirmed rare species may be more attractive to birding tourists, raising the marginal value of an additional record. The model abstracts from this feedback by treating v as fixed. Allowing v to increase in the cumulative species list would generate an additional amplification channel beyond the deterrence-erosion multiplier, making the current specification conservative.

[8] The fixed-K assumption reflects institutional constraints specific to volunteer-based review systems. eBird reviewers are unpaid regional experts recruited through personal networks by the Cornell Lab; there is no open application process and no mechanism to scale reviewer supply in response to submission volume. Brazil has roughly 150–200 active regional reviewers, a number that has remained approximately constant since 2017 despite a near-quadrupling of annual submissions. Table 11 confirms empirically that reviewer workload rises one-for-one with the post-AI submission surge, with no offsetting increase in review capacity.

[9] In practice, π likely varies with the quality of the fabrication; a poorly constructed checklist is easier to

Reporter i with post-AI fabrication cost $c_i - \Delta$ fabricates if and only if the expected benefit exceeds the expected cost:

$$v(1-\phi\pi) > c_i - \Delta \tag{1}$$

where $v(1-\phi\pi)$ is the expected value of a fabricated report, accepted with probability $1-\phi\pi$, with ϕ denoting the probability of being reviewed and π the detection probability conditional on review. Rearranging, the fabrication threshold is $c^* = v(1-\phi\pi) + \Delta$.

In equilibrium, the fabrication threshold c^* and the review rate ϕ are jointly determined by two conditions:

$$c^* = v(1-\phi\pi) + \Delta \tag{2}$$

$$\phi = \frac{K}{\bar{\lambda} + F(c^*)} \tag{3}$$

where $\bar{\lambda} = \sum_h \lambda(h)$ is aggregate genuine arrivals. Equation (2) is the marginal fabricator's indifference condition. Reporters with cost below c^* fabricate; those above do not. Equation (3) is the review rate given total submissions. As the fabrication threshold rises, more fakes enter the system, total submissions $N = \bar{\lambda} + F(c^*)$ increase, and the review rate falls.

The key insight is the feedback between these two equations. An increase in Δ, which parameterizes AI availability, directly raises c^* as more reporters find it worthwhile to fabricate. This increases $F(c^*)$, which reduces ϕ through equation (3). The lower review rate further raises c^* through equation (2), inducing additional fabrication. This positive feedback loop generates amplification.

Proposition 1 (AI increases manipulation with amplification)**.** *Under the conditions that the feedback term lies strictly between 0 and 1, which holds for K sufficiently large, the effect*

detect than a single plausible-looking record. Allowing π to decrease in AI availability Δ would introduce an additional channel through which AI erodes deterrence, reinforcing the amplification result. The constant-π assumption is therefore conservative: it understates the heterogeneity in fabrication strategies and the total effect of AI on data quality.

of AI on the fabrication threshold satisfies

$$\frac{dc^*}{d\Delta} = \frac{1}{1 - \underbrace{\frac{vK\pi f(c^*)}{[\bar{\lambda} + F(c^*)]^2}}_{\text{feedback term} \in (0,1)}} > 1 \tag{4}$$

Proof. Substituting (3) into (2) gives $c^* = v\left(1 - \frac{K\pi}{\bar{\lambda}+F(c^*)}\right) + \Delta$. Totally differentiating with respect to Δ: $\frac{dc^*}{d\Delta} = \frac{vK\pi f(c^*)}{[\bar{\lambda}+F(c^*)]^2} \cdot \frac{dc^*}{d\Delta} + 1$. Solving yields the expression above. The feedback term is positive since all components are positive, and less than one when K is sufficiently large relative to $v\pi f(c^*)$, ensuring a unique stable equilibrium. □

The multiplier $dc^*/d\Delta > 1$ means that AI's effect on manipulation exceeds the direct cost reduction. The amplification is larger when the private value v of accepted records is higher, when the detection probability π is higher, and when the density of marginal fabricators $f(c^*)$ is higher.

Proposition 2 (Manipulation stronger in low-habitat areas)**.** *The percentage effect of AI on total reports is decreasing in habitat quality:*

$$\frac{\partial}{\partial \Delta} \log N(h) = \frac{F'(c^*) \cdot dc^*/d\Delta}{\lambda(h) + F(c^*)} \quad \textit{is decreasing in } h \tag{5}$$

Proof. Total reports in municipality h are $N(h) = \lambda(h) + F(c^*)$. The numerator $F'(c^*) \cdot dc^*/d\Delta$ is independent of h, since AI reduces costs uniformly. The denominator $\lambda(h) + F(c^*)$ is increasing in h because $\lambda'(h) > 0$. Hence the ratio is decreasing in h. □

The intuition is that AI raises $F(c^*)$, the number of fakes, uniformly across municipalities, but in low-h areas where $\lambda(h)$ is small, fakes constitute a larger share of total reports. In high-h areas, genuine arrivals dilute the fake signal. This motivates the triple DiD: the coefficient on the triple interaction of post-AI treatment, species rarity, and low forest cover captures the excess manipulation signal in ecologically implausible areas.

The model generates additional predictions that I test in the empirical analysis.

Corollary 1 (Tourism value amplifies manipulation)**.** *The equilibrium fabrication level* $F(c^*)$

is increasing in v. Municipalities with higher private value from rare species records, proxied by birding tourism intensity, should exhibit larger manipulation signals.

Corollary 2 (AI access mediates the effect)**.** *The manipulation signal is increasing in Δ. Municipalities with higher AI access, proxied by mobile coverage required for AI identification apps, should show larger effects.*

Empirical Prediction 1 (Temporal signatures of fabrication)**.** *Because AI eliminates the time-of-day constraint on identification, these tools work any time, but genuine birding requires daylight, fabricated records should be disproportionately submitted at night.*

Empirical Prediction 2 (Indiscriminate AI predicts null charismatic and seasonal effects)**.** *If manipulation is driven by indiscriminate AI misidentification rather than targeted tourism fraud, there should be no differential effect for charismatic species and no seasonal concentration during the tourist season.*

4 Data

4.1 eBird Brazil

I obtain the complete eBird dataset for Brazil from eBird's data request system, covering all checklist submissions from January 2017 through March 2025.[10] The raw data contain individual checklist-level observations with species identifications, geographic coordinates (mapped to municipalities using IBGE shapefiles), date, time of day, observer effort (duration, distance, number of observers), review status, and a unique observer identifier. Brazil's eBird database has grown rapidly, from about 1.2 million records in 2017 to over 4.5 million in 2024.

I aggregate the raw data to the municipality-year-rarity level. Each observation in the raw data is a species-level record within a checklist; a single checklist may contain multiple species records. The outcome log(records) counts total species-level observations in a cell, and the review rate is the share of species-level records in a cell that were flagged for expert

[10]The 2025 data represent a partial year of three months. Results are robust to dropping 2025 entirely.

review. The review rate measures scrutiny, not validation. A flagged record may ultimately be confirmed as genuine.[11] Rare species are defined as the bottom tercile of pre-2021 national observation counts over 2017-2020. This definition is computed entirely in the pre-treatment period to avoid circularity, so species that become more commonly reported after the AI boom do not retroactively change classification. The bottom tercile corresponds to roughly 630 species out of Brazil's ∼1,900 confirmed species.[12]

The key outcome variables are as follows. First, log of total observation records in a municipality-year-rarity cell, where the plus-one transformation accommodates zeros and results are robust to inverse hyperbolic sine and Poisson specifications[13]; and second, the review rate, defined as the share of records flagged for expert review in a cell, the primary data quality indicator reflecting eBird's internal assessment of record plausibility.

The sample covers 5,572 municipalities over 2017-2025. The baseline DiD sample, which includes all municipalities with at least one record, contains 32,657 municipality-rarity-year cells. Pre-AI average review rates over 2017-2020 are 4.8% overall and substantially higher for rare species, reflecting their inherently greater implausibility. Post-AI average review rates over 2021-2025 are 5.2%, but this aggregate masks substantial heterogeneity across rarity and forest cover categories.

4.2 MapBiomas Forest Cover

Municipality-level forest cover comes from MapBiomas Collection 9 (Souza et al., 2020), the most complete land-use dataset available for Brazil. MapBiomas uses Landsat satellite imagery at 30-meter resolution, processed through Google Earth Engine, with random forest classifiers trained on over 80,000 labeled samples per biome. Classification accuracy for forest

[11] This distinction matters because a high review rate does not imply poor data quality. The review rate captures the intensity of the verification process, which is the channel through which the deterrence-erosion multiplier operates.

[12] Results are robust to alternative rarity thresholds. Table S8 reports estimates using the bottom 20th and 40th percentiles. The triple interaction coefficient increases monotonically with stricter rarity definitions, from 0.151 at P40 to 0.280 at P20, consistent with the model prediction that rarer species exhibit stronger manipulation signals.

[13] The inverse hyperbolic sine transformation $\sinh^{-1}(Y)$ avoids the arbitrary plus-one and produces nearly identical coefficients. A Poisson pseudo-maximum-likelihood specification accommodates the count nature of the data and yields qualitatively similar results. I report the log specification throughout for interpretability.

categories exceeds 89% in independent validation samples.[14]

I define forest pixels as the union of four MapBiomas categories: forest formation, class 3; savanna formation, class 4; mangrove, class 5; and floodable forest, class 6. These categories collectively capture the habitat types required by forest-dependent bird species. The forest share for municipality i in year t is the ratio of forest pixels to total municipality area. I define the binary indicator Low Forest as below-median average forest share over 2017-2020. The median pre-treatment forest share is about 35%, with substantial geographic variation: Amazon municipalities average above 80%, while northeastern interior and southern agricultural municipalities often fall below 10%.

For the triple DiD sample, I restrict to municipalities that can be matched to MapBiomas forest data and have at least one eBird record in the sample period, yielding $N = 19{,}214$ municipality-year-rarity cells.

4.3 Other Data

I construct several additional variables for the heterogeneity and mechanism analyses. Night submissions are defined as records submitted between 22:00 and 06:00 local time, hours when genuine birdwatching is impossible for virtually all species.[15] About 3% of all records fall in this window pre-AI. Urban/mobile coverage is proxied by an above-median urbanization rate indicator, which correlates strongly with 4G/5G mobile coverage required for AI identification apps to function. Birding tourism intensity is measured by log pre-2021 checklist density, defined as total checklists per municipality over 2017-2020, and alternatively by a binary high-birding-tourism indicator based on the municipality's ranking in the national distribution. Protected areas are defined using ICMBio, the Chico Mendes Institute for Biodiversity Conservation, shapefiles, with the municipality-level variable measuring the share of area within federal or state conservation units. Charismatic species are coded as iconic, large-bodied species visible in ecotourism marketing: raptors such as eagles and hawks,

[14] If MapBiomas misclassifies some forest pixels as non-forest or vice versa, the resulting measurement error in the forest share variable would attenuate the triple DiD estimates toward zero, making the current specification conservative.

[15] A small number of legitimate nocturnal species exist in Brazil, including owls (Strigidae, Tytonidae), nightjars (Caprimulgidae), and potoos (Nyctibiidae). The night-time results are robust to excluding known nocturnal families from the rare species sample.

large parrots such as macaws and Amazons, toucans, and large antbirds. Endemic species are flagged based on range restriction to Brazil or to specific biomes, such as Atlantic Forest endemics. Weekend is an indicator for Saturday and Sunday submissions. Dry season covers May through September, the peak period for birding tourism in most of Brazil.

Table 1 presents summary statistics for the main analysis sample. Figure 1 shows the cross-sectional distribution of the key variables. Panel (a) maps eBird record density by municipality, showing that birding activity is concentrated in southeastern Brazil, along the Atlantic coast, and near major urban centers, with sparse coverage in the interior northeast and western Amazon. Panel (b) maps pre-treatment forest cover from MapBiomas, with the Amazon basin, Atlantic Forest remnants, and Pantanal region showing the highest values, while the southern agricultural belt and northeastern Caatinga are largely deforested. The spatial mismatch between areas of high birding activity and areas of high forest cover is an important feature of the identification strategy: low-forest municipalities with birding activity form the treatment group where rare species records are ecologically implausible.

5 Empirical Strategy

The baseline specification is a two-way fixed effects DiD comparing rare to common species before and after the AI boom:

$$Y_{irt} = \alpha_{ir} + \alpha_t + \beta\,(\text{Post}_t \times \text{Rare}_r) + \varepsilon_{irt} \tag{6}$$

where Y_{irt} is the outcome, log records or review rate, for municipality i, rarity group $r \in \{0, 1\}$, and year t. $\text{Post}_t = \mathbf{1}[t \geq 2021]$ is the AI boom indicator. α_{ir} are municipality $\times$ rarity fixed effects and α_t are year fixed effects. Standard errors are clustered at the municipality level.[16]

The main specification adds the forest cover dimension to form a triple DiD:

$$Y_{irt} = \alpha_{ir} + \alpha_t + \beta_1(\text{Post}_t \times \text{Rare}_r) + \beta_2(\text{Post}_t \times \text{LF}_i) + \beta_3(\text{Post}_t \times \text{Rare}_r \times \text{LF}_i) + \varepsilon_{irt} \tag{7}$$

[16]Results are robust to alternative clustering at the state or biome level, which allows for broader spatial correlation in unobservables. State-level clustering produces slightly wider confidence intervals but all main coefficients remain significant at conventional levels.

where LF_i is the predetermined low-forest indicator, fixed at the pre-treatment average over 2017–2020. The coefficient of interest is β_3, the differential change in outcomes for rare species in low-forest municipalities after the AI treatment. The model in Section 3 predicts $\beta_3 > 0$ for log records and $\beta_3 < 0$ for the review rate.

To validate the parallel trends assumption, I estimate an event study specification:

$$Y_{irt} = \alpha_{ir} + \alpha_t + \sum_{k \neq 2020} \gamma_k \left(\mathbf{1}[t = k] \times \text{Rare}_r \times \text{LF}_i \right) + \mathbf{X}'_{irt}\delta + \varepsilon_{irt} \tag{8}$$

where 2020 is the omitted reference year, the last pre-treatment year. Under the identifying assumption, the pre-treatment coefficients γ_k for $k < 2021$ should be zero.

The identifying assumption for β_3 in equation (7) is that, absent the AI boom, trends in rare versus common species records would have evolved similarly in low- versus high-forest municipalities. This is a parallel trends assumption in triple-difference form: even if rare species records follow different trajectories from common species records, captured by Post × Rare, and even if low-forest municipalities trend differently from high-forest ones, captured by Post × Low Forest, the interaction between these two differential trends would have been zero absent the AI treatment. Three pieces of evidence support this assumption. First, the event study in Figure 2 shows flat pre-trends: the coefficients on Rare × Low Forest × Year are statistically indistinguishable from zero for all pre-treatment years, 2017-2020. This rules out gradual divergence in rare-species reporting across forest cover categories before the AI boom. Second, the timing of the AI boom was driven by global technology development, not by local conditions in Brazilian municipalities. There is no plausible channel through which municipality-level birding patterns in Brazil could have influenced the global trajectory of AI development. Third, I conduct placebo tests using 2018 and 2019 as false treatment dates on the pre-treatment sample and find null effects, as shown in Figure S3, confirming the absence of pre-existing differential trends.

The main threat to identification is COVID-19, which began affecting Brazil in March 2020. COVID-19 could bias the estimates if mobility restrictions differentially affected birding activity for rare versus common species across low- and high-forest municipalities. I address this concern through three complementary tests. First, I shift the treatment date

to 2022, after the main wave of COVID restrictions and vaccine rollout, and find similar coefficients. If COVID were driving the results, the effect should weaken or disappear when treatment is moved past the pandemic period. Second, I control for Google Community Mobility indices and their interactions with rarity and forest cover. The triple interaction coefficient is essentially unchanged at 0.208 versus 0.200 for record counts. Third, the event study using 2022 as the reference year shows that 2021 is flat, with the sharp break occurring at AI adoption timing rather than COVID timing, as shown in Figure S4.

A second concern is that deforestation itself could be changing during the sample period, confounding the forest cover variable. I address this by defining low forest using pre-treatment forest share averaged over 2017-2020, which is predetermined relative to the 2021 treatment. Time-varying forest cover changes within the sample period are absorbed by the municipality $\times$ rarity fixed effects. In addition, MapBiomas data show that forest cover in low-forest treatment municipalities continued to decline during 2021-2023, which would work against finding increased rare species records, since habitat is deteriorating, not improving.

A third concern is that AI identification tools might genuinely improve species identification accuracy, leading to more correct rare species records rather than fabricated ones. I address this "better identification" alternative in detail in Section 7.

Finally, I note that the triple DiD design has a built-in placebo dimension. If the results were driven by a confound affecting rare species generally but not differently across forest cover levels, or a confound affecting low-forest municipalities generally but not differently across rarity categories, these confounds would be captured by the two-way interactions β_1 and β_2 rather than the triple interaction β_3. The fact that β_3 is precisely estimated and significant while β_1, the Post $\times$ Rare coefficient for log records, is near zero provides reassurance that the identifying variation is specific to the intersection of rarity and ecological implausibility.

6 Results

6.1 Baseline and Main Results

Table 2 reports the baseline DiD estimates from equation (6). In Column (1), the outcome is the expert review rate, defined as the share of species observation records flagged for volunteer review in each municipality-rarity-year cell. The coefficient on Post × Rare is −0.061, with a dependent variable mean of 0.048. The point estimate implies that rare species observation records experience a 6.1 percentage point decline in expert review rates after the AI boom, corresponding to 42 percent of the sample standard deviation of 14.5 percentage points. The overall eBird Brazil database contains about 1,900 confirmed species, of which roughly 630 qualify as rare under the bottom-tercile definition; the review rate decline therefore affects the quality assurance process for one-third of Brazil's avifauna. Figure 4 plots the corresponding event study: the pre-2021 coefficients are flat, and a sharp decline in review rates begins precisely at 2021, growing in magnitude through 2024.

Table 3 reports the main triple DiD estimates from equation (7). In Column (1), the outcome is the log of species observation record counts in each municipality-rarity-year cell. The coefficient on Post × Rare is 0.041, small and imprecisely estimated, indicating that the volume increase for rare species is not a general post-AI trend. The coefficient on Post × Low Forest is −0.082, marginally significant, suggesting a slight contraction of overall birding activity in low-forest areas. This contraction works against the triple interaction finding: if legitimate birders avoid low-forest areas post-AI, the denominator of total records shrinks, making the 22 percent estimate conservative. The coefficient of interest, Post × Rare × Low Forest, is 0.200, with a dependent variable mean of 2.806. The point estimate implies that rare species observation records in low-forest municipalities, where such species are ecologically implausible, increase by about 22 percent after the AI boom, corresponding to 9.5 percent of the sample standard deviation. In absolute terms, the implied excess corresponds to roughly 2,800 suspect rare species records in low-forest municipalities across the post-2021 period, out of a total of roughly 15,600 rare species records in these areas.[17]

[17] The 2,800 figure is computed as $(\exp(0.200) - 1) \times \bar{Y}_{\text{pre}}$, where $\bar{Y}_{\text{pre}}$ is the average pre-AI rare species record count in low-forest municipalities, summed over the five post-treatment years. The calculation accounts for the log transformation and the pre-treatment base, not simply 22% of 15,600.

By comparison, eBird Brazil as a whole contains about 4.5 million annual records; the excess suspect records thus represent a small share of total submissions but a large share of the specific cells, rare species in low-forest areas, that are most consequential for conservation assessments.

In Column (2), the outcome is the expert review rate. The coefficient on Post $\times$ Rare $\times$ Low Forest is -0.097. The share of rare species records flagged for expert review falls by 9.7 percentage points in low-forest municipalities, corresponding to 67 percent of the sample standard deviation of 14.5 percentage points. The review system is not merely weakened but effectively collapsed for rare species in these areas. The two-way interaction Post $\times$ Low Forest is positive and significant for the review rate at 0.006, indicating marginally increased scrutiny of common species in low-forest areas, which makes the rare-species review rate decline even more striking. Reviewers are paying more attention to common species in these areas but less attention to rare ones.

In Column (3), the outcome is the novel location rate, defined as the share of records at locations where the species was not previously recorded. The triple interaction is -0.026, negative but imprecisely estimated, indicating that the surge in rare species records is not concentrated at novel locations within municipalities. This is consistent with, but not confirmatory of, fabricated records being submitted at established birding hotspots where they are less conspicuous rather than at novel locations that would draw automatic scrutiny.

Figure 2 plots the event study coefficients from equation (8). For both outcomes, the pre-2021 coefficients are flat and statistically indistinguishable from zero. A sharp break occurs at 2021, the year of the AI boom in Brazil. The post-2021 coefficients are positive and growing for record counts (Panel a) and negative and growing in magnitude for the review rate (Panel b). For record counts, the effect appears immediately in 2021 and grows through 2024, consistent with gradual diffusion of AI adoption. For review rates, the decline is more immediate and sustained, suggesting rapid overwhelm of the review system. The monotonically increasing pattern post-2021 is consistent with the deterrence-erosion multiplier. As more fakes enter the system, the review rate falls further, inducing additional fabrication in subsequent years.

Figure 3 shows the geographic pattern. Panel (a) maps the percentage change in rare

species records by municipality, comparing post-AI to pre-AI periods. Panel (b) overlays the low-forest municipality boundaries. Rare record growth is widespread but visually concentrated in low-forest municipalities, the "falsification zones" identified by the triple DiD.

6.2 Heterogeneity

I examine six dimensions of heterogeneity, each motivated by testable predictions from the model. Table 4 reports results from all six dimensions together; full regression details appear in Appendix Tables S1–S4.

The first heterogeneity dimension is night-time submissions, reported in Table 4, Panel A, Columns (1)–(2). AI identification tools are available 24 hours a day, while genuine birdwatching is overwhelmingly a dawn and daytime activity. The triple interaction Post × Rare × Night is 0.204 for the log of record counts. Night-time rare species submissions surge by about 23 percent post-AI. Importantly, the overall level of night submissions declines sharply: the coefficient on Post × Night is −0.389, indicating a general 32 percent reduction in nocturnal reporting. The rare-specific night surge occurs against this declining backdrop, making it a particularly strong signature of fabrication. The review rate interaction in Column (2) is positive but imprecisely estimated at 0.008, suggesting that the review system does not differentially flag night-time rare species records.

The second dimension is mobile coverage, reported in Table 4, Panel A, Columns (3)–(4). AI identification apps require 4G/5G mobile connectivity to download species identification packs and to function in the field. The triple interaction Post × Rare × Urban is −0.053 for the review rate. High-coverage municipalities see a 5.3 percentage point steeper decline in review rates for rare species, which corresponds to 102 percent of the sample mean. The volume effect in Column (3) is directionally positive but imprecise at 0.028, likely because urban areas have both higher AI adoption and denser reviewer presence that partially offsets the volume increase.

The third dimension is endemic species, reported in Table 4, Panel B, Column (1). The interaction Post × Endemic-rich family is −0.004 for record counts. The coefficient is economically trivial and imprecisely estimated. There is no evidence of endemic-specific targeting, consistent with indiscriminate AI misidentification, since AI tools do not distinguish

endemic from non-endemic species.

The fourth dimension is charismatic species, reported in Table 4, Panel B, Column (2). The interaction Post × Charismatic is −0.013 for record counts, imprecisely estimated. Charismatic species, large iconic birds like raptors, macaws, and toucans, are not differentially targeted. Tourism-motivated fraud should favor charismatic species; indiscriminate AI misidentification should not.

The fifth dimension is novel municipality records, reported in Table 4, Panel B, Column (3). The coefficient on Post × Rare for the novel-municipality rate is −0.031. The dependent variable mean is 0.611, corresponding to 5.1 percent of the mean. The novel-municipality rate actually falls post-AI, consistent with fabricated records appearing in municipalities where the species was already recorded at least once, making them less conspicuous to eBird's automated novelty filters.

The sixth dimension is protected areas, reported in Table 4, Panel B, Column (4). The interaction Post × Rare × Protected is −0.001 for record counts, with neither the volume nor the review rate coefficient significant. The null result rules out conservation-area-specific fraud as the driver.

For the six heterogeneity tests, I compute Benjamini-Hochberg adjusted p-values to account for multiple comparisons. Night and Urban remain significant at $p_{\text{adj}} < 0.007$. The novel municipality coefficient is also statistically significant after adjustment, though its sign is negative, consistent with fabricated records appearing in municipalities where the species was already recorded. Endemic, Charismatic, and Protected remain imprecise at $p_{\text{adj}} > 0.4$. I apply the BH correction to these six pre-specified heterogeneity dimensions rather than to the full set of robustness and mechanism tests, which address different questions.

6.3 Robustness

The main threat to identification is COVID-19. Brazil experienced major pandemic waves in early 2021, precisely when the AI boom began. If mobility restrictions differentially affected birding activity for rare versus common species across low- and high-forest municipalities, the triple DiD could capture pandemic disruption rather than AI-enabled manipulation. I take this concern seriously and address it with three tests, leading with the strongest.

The most direct test shifts the treatment date to 2022, past Brazil's main COVID restrictions and vaccine rollout. If the pandemic drives the results, moving treatment beyond the recovery period should eliminate or substantially weaken the effect. It does not. The triple interaction moves from 0.200 to 0.158 for record counts and from −0.097 to −0.087 for the review rate, both remaining statistically significant (Table S5). The event study using 2022 as the reference year (Figure S4) is equally informative. The year 2021 is flat and indistinguishable from 2020, with the sharp break occurring at AI adoption timing rather than COVID timing. Second, I add Google Community Mobility indices and their interactions with rarity and forest cover. The triple interaction is essentially unchanged at 0.208 for record counts and −0.100 for the review rate (Table S6). Mobility itself is insignificant ($p = 0.91$), confirming that COVID behavioral changes are orthogonal to the manipulation signal.

A second concern is that the binary low-forest indicator may introduce an artificial discontinuity at the median. I address this by replacing the binary indicator with continuous ecological implausibility, defined as one minus forest share, so that higher values indicate less forest. As reported in Table S7, the triple interaction is 0.267 for record counts and −0.120 for the review rate, both statistically significant. The continuous specification confirms that the manipulation signal is not an artifact of the median split and strengthens monotonically as ecological implausibility increases. Figure S2 shows a binscatter of the Post × Rare coefficient by forest cover decile, confirming the monotonic gradient.

A third concern is sensitivity to the rarity definition. I vary the rarity cutoff from P20, the strictest definition covering the bottom 20 percent of species, to P40, the most lenient definition covering the bottom 40 percent. As reported in Table S8, the triple interaction for record counts moves from 0.280 at P20 to 0.195 at the baseline P33 to 0.151 at P40. For the review rate, the gradient is from −0.198 at P20 to −0.092 at P33 to −0.043 at P40. The monotonic gradient is consistent with the model prediction that rarer species show stronger manipulation signals. At the strictest threshold, the volume effect is 32 percent and the review rate decline is 19.8 percentage points, more than four times the sample mean review rate of 0.045.

A fourth concern is that 2025 data may understate true totals due to submission lags.

Dropping 2025 from the sample produces virtually identical coefficients, as shown in Table S9. The triple interaction is 0.200 for record counts, identical to the baseline, confirming that the 2025 data neither drive nor distort the main finding.

A fifth concern is pre-existing differential trends. I conduct placebo tests using 2018 and 2019 as false treatment dates on the pre-treatment sample of 2017-2020. As shown in Figure S3, the placebo triple interaction for record counts is −0.001 at the 2018 cutoff and 0.020 at the 2019 cutoff, both economically trivial and imprecisely estimated, confirming the absence of pre-existing differential trends.

A sixth concern is that the excess rare species records might trigger eBird's automated detection systems, suggesting they are not truly evasive. Table S10 tests this directly. The Post × Rare coefficient on the novel location rate is −0.000 and on the private location rate is −0.003, both imprecisely estimated. The coefficient on the suspicious checklist rate, checklists with two or more rare species, is −0.027. The excess records do not appear at novel or private locations and do not cluster into suspicious multi-species checklists. If anything, the suspicious checklist rate *declines*, consistent with marginal fakes being individually plausible single-species submissions that fly under automated filters. Figure S1 additionally shows that forest cover in low-forest municipalities continued to decline through 2023, ruling out habitat recovery as an alternative explanation.

7 Mechanisms

The main result establishes that rare species records in ecologically implausible areas increase, and their review rates decline, after the introduction of AI identification tools. This section investigates what drives these patterns and who is responsible.

7.1 Ruling Out Reviewer Dilution

I first test whether the review rate decline reflects genuine reviewer dilution, a mechanical overwhelm of reviewers by the total volume of post-AI submissions. Table 5 presents results from two tests that reject this interpretation. In Column (1), I estimate the effect of Post-AI on review rates for common species only. The coefficient on Post-AI is −0.0013, imprecisely

estimated at conventional levels. Common species review rates are entirely unaffected by the post-AI submission surge, ruling out a general dilution story. In Columns (2) and (3), I add controls for common species volume and its interaction with Post to the baseline rare-species review rate regression. The Post × Rare coefficient moves from −0.0506 in Column (2) to −0.0484 in Column (3) with the full set of volume controls, less than 6 percent shrinkage. The review rate decline for rare species is volume-robust and not mechanically driven by reviewers processing more records overall. Table S11 provides a complementary test. Review rates actually *increase* by 2.0 percentage points in low-reviewer-density municipalities post-AI, confirming that the rare-species review decline is not a general capacity shortage but is specific to the rare-species channel. These tests rule out dilution at the national level; dilution within specific high-submission municipalities remains possible, and Table 11 provides the geographic disaggregation.

7.2 Ruling Out Better Identification

A natural alternative interpretation is that AI tools genuinely improve species identification accuracy, so that novice birders correctly identify more rare species post-AI. Under this "better ID" hypothesis, the increase in rare species records reflects real ecological observations, and review rates should *rise*, not fall, because better-identified records are more likely to pass review. Three pieces of evidence reject this interpretation.

First, the review rate decomposition. Table 6 decomposes the review outcome into its components. In Column (1), the outcome is the reviewed share and the triple interaction Post × Rare × Low Forest is −0.0970. In Column (2), the outcome is the unreviewed share and the coefficient is 0.0970, a perfect mirror image.[18] Records in treated cells escape the review system entirely rather than passing through it and being confirmed, the opposite of the "better ID" prediction. In Column (3), the outcome is the too-good-to-be-true rate, hereafter TGTB, the share of checklists with two or more rare species. The triple interaction is 0.005, imprecisely estimated. In Column (4), the outcome is the novel location rate, with

[18]The symmetry between the reviewed and unreviewed share coefficients is mechanical: since reviewed + unreviewed = 1, the coefficients must sum to zero. The substantive content is in the levels and the interaction pattern, the fact that the decline concentrates in rare species in low-forest municipalities, not the mirror image itself.

a triple interaction of −0.026, also imprecisely estimated. The decomposition confirms that records are evading review, not passing through it.

Table 7 provides further evidence by separating TGTB from non-TGTB checklists. In Column (1), the outcome is the count of TGTB records and the triple interaction is 0.155, marginally significant. In Column (2), the outcome is the review rate conditional on TGTB, and the coefficient is −0.013, imprecisely estimated. The TGTB review rate mean is 0.092, nearly double the overall review rate of 0.052, confirming that egregious checklists still attract scrutiny. In Column (3), the outcome is the review rate for non-TGTB records, and the coefficient is −0.105, corresponding to 204 percent of the non-TGTB sample mean of 0.051. The review rate decline concentrates entirely among single-rare-species checklists. Marginal fakes evade review, while egregious multi-species fakes are still caught.

Second, the "better ID" hypothesis cannot explain the concentration of effects at night. Night-time submissions, between 22:00 and 06:00 when genuine birding is impossible for virtually all species, surge by 23 percent post-AI. Genuine improvement in identification accuracy would operate during daylight hours when birds are actually observed. Third, the effect concentrates in low-forest municipalities where rare forest-dependent species cannot plausibly occur regardless of identification accuracy. Better identification tools cannot generate correct identifications of species that are ecologically absent. The mirror-image decomposition, the night-time surge, and the low-forest concentration are each individually inconsistent with better identification and collectively decisive against it.

Having ruled out better identification, I turn to the economic incentives behind the manipulation signal and the identity of the manipulators.

7.3 Tourism Incentives and Who Manipulates

I next investigate the role of birding tourism incentives. Table 8 presents results using two measures of tourism intensity. In Columns (1)–(2), the heterogeneity variable is a binary high-birding-tourism indicator. The triple interaction Post × Rare × High-birding is 0.455 for record counts in Column (1). The point estimate implies that high-tourism municipalities see a 58 percent larger increase in rare species records relative to low-tourism municipalities, corresponding to 21.5 percent of the sample standard deviation. Checklist density proxies

for birding tourism activity rather than directly measuring the private value v in the model, so these results are consistent with Corollary 1 rather than a direct test of it. In Column (2), the review rate interaction is -0.047, directionally negative but imprecise. In Columns (3)–(4), I replace the binary indicator with continuous log pre-2021 checklist density. The triple interaction Post $\times$ Rare $\times$ log(pre-checklists) is 0.236 for record counts in Column (3) and -0.024 for the review rate in Column (4). Municipalities with one standard deviation higher pre-existing birding activity show about 24 percent more excess rare records and a 2.4 percentage point steeper review rate decline post-AI.

To distinguish organized tourism guides from casual AI users, I examine weekend patterns. Table 9 presents results. In Columns (1)–(2), the triple interaction Post $\times$ Rare $\times$ Weekend is 0.094 for record counts. Rare records surge by about 10 percent on weekends post-AI, corresponding to 3.6 percent of the sample mean. The review rate interaction in Column (2) is 0.010, imprecisely estimated. In Columns (3)–(4), I add the four-way interaction with birding tourism intensity. The three-way term Post $\times$ Rare $\times$ Weekend, which now captures the weekend effect in non-tourism areas, is 0.310. The four-way interaction Post $\times$ Rare $\times$ Weekend $\times$ High Birding is -0.218. The weekend effect is substantially smaller in tourism areas: the four-way interaction offsets roughly two-thirds of the weekend premium ($0.218/0.310 \approx 0.70$). This pattern inverts the natural prior. If professional guides drove the manipulation, they would concentrate activity on weekends when paying tourists visit, and the weekend premium should be *larger* where tourism infrastructure exists. Instead, the weekend surge is concentrated in non-tourism municipalities, precisely where organized guides have no economic incentive to operate. The result points to a different type of manipulator: casual weekend hobbyists with smartphones and AI apps, birding recreationally in areas without professional tourism infrastructure. In tourism areas, guides work throughout the week, smoothing any weekend signal.

7.4 Null Results and Reviewer Capacity

Two null results sharpen the interpretation. Table 10 reports charismatic species interactions. In Columns (1)–(2), the interaction Post $\times$ Rare $\times$ Charismatic is -0.021 for record counts and -0.022 for the review rate. Neither is statistically significant. Tourism-motivated fraud

should favor charismatic species; indiscriminate AI misidentification should not. In Columns (3)–(4), I add the four-way interaction with birding tourism. The four-way term Post × Rare × Charismatic × High Birding is −0.077, imprecisely estimated. Even in high-tourism areas, charismatic species are not differentially targeted.

Table 10 also reports seasonality tests. The interaction Post × Rare × Dry Season is −0.020 for record counts and 0.005 for the review rate. Neither coefficient is significant. Manipulation is year-round, inconsistent with tourism-driven seasonal fraud and consistent with AI misidentification that occurs whenever a user submits a record. Figure S6 confirms the absence of seasonal concentration: monthly Post × Rare coefficients are flat across the calendar year.

Reviewer capacity is directly affected. Table 11 shows that the triple interaction Post × Rare × Low Forest is 0.200 for log reviewer workload in Column (1) and 0.097 for the unreviewed rate in Column (2). Reviewer workload rises in treated cells while the unreviewed share increases by the same magnitude as the reviewed share declines, confirming that the review system is overwhelmed rather than scaled. Figure 5 plots the event study of reviewer workload. Pre-2021 coefficients are flat, and a sharp increase begins at 2021, growing through 2024.

The mechanism tests point to a consistent profile. Casual users with smartphones in areas with mobile coverage, submitting AI-assisted rare species records at unusual times such as nights and weekends, without targeting charismatic species or timing submissions to the tourist season. The channel is indiscriminate AI misidentification by non-expert users, amplified by the capacity-constrained review system. I cannot determine from the data whether the excess records are spread across many casual misidentifiers or concentrated among a small number of prolific fabricators, a distinction that matters for the optimal policy response.

The mechanism tests align with every prediction of the model in Section 3. The manipulation signal concentrates where habitat is absent, where tourism incentives are strong, where AI access is feasible, and at times when real birding is impossible, while the predicted nulls for charismatic species, seasonality, and protected areas hold as well.

8 Conclusion

This paper documents a deterrence-erosion multiplier in citizen science data. AI lowers fabrication costs, the resulting flood of marginal fakes overwhelms the fixed-capacity review system, and the reduced detection probability induces further fabrication. The evidence from eBird Brazil lines up with the model on every dimension I can test, including the predicted null results for charismatic species and seasonality, which rule out the organized-fraud alternative.

I estimate that roughly 2,800 excess suspect records appeared in low-forest municipalities post-2021, including roughly 200 involving threatened species across 60 unique species. Novel rare-species and low-forest municipality pairs increased from 262 per year pre-AI to 315 per year post-AI, a 20 percent annual increase. These records enter species distribution models, IUCN Red List assessments, and Brazil's national conservation planning by ICMBio (Rodrigues et al., 2006). The consequences propagate through multiple channels. Inflated species ranges can bias threat assessments in unpredictable directions (Jetz et al., 2012), potentially misallocating conservation resources away from actual habitat corridors (Ferraro and Pattanayak, 2006). More broadly, eBird data inform hundreds of ecological studies annually (Sullivan et al., 2014; Callaghan et al., 2021a), and contaminated records propagate through biodiversity indices and meta-analyses without systematic correction.

The results motivate three concrete platform design interventions. First, geographic plausibility filters. eBird's existing filters check species-location-date combinations against historical frequencies, but do not incorporate habitat suitability. Cross-referencing species habitat requirements with satellite-derived land cover data would flag records of forest-dependent species in low-forest municipalities for mandatory expert review. Second, AI-provenance tagging. Require disclosure when a species identification was generated by an AI tool and weight AI-assisted rare-species records differently in downstream databases. Since voluntary disclosure may be unreliable, this intervention is most effective if implemented at the app-developer level, with AI identification tools automatically flagging their outputs when integrated with eBird. Third, targeted reviewer allocation. Dynamically direct review resources to high-risk cells rather than distributing them uniformly, raising the effective

detection rate without requiring an increase in total reviewer supply. All three interventions involve a tradeoff. Geographic plausibility filters would also flag genuine vagrant records by expert birders in low-forest areas, increasing false positives. The optimal policy depends on the relative costs of false positives for conservation assessments versus false negatives from undetected fabrication.

These findings have implications for the broader trajectory of AI deployment. AI identification tools are becoming more powerful and more widely adopted across citizen science platforms, including iNaturalist and community environmental monitoring networks. As AI capabilities improve, the cost of generating plausible fabricated records will continue to fall, while the incentives to manipulate biodiversity data remain strong, because species occurrence records directly influence conservation funding allocations, land-use regulations, and environmental impact assessments. In Brazil alone, the ICMBio Red List for more than 200 threatened bird species relies on occurrence databases including eBird; similar dependence exists in countries across the tropics where biodiversity conservation needs are most acute and institutional monitoring capacity is most limited. The deterrence-erosion multiplier identified in this paper implies that the data quality problem will worsen nonlinearly as AI adoption accelerates, unless verification systems are proactively scaled.

The dual-use dilemma extends beyond citizen science. The underlying structure, costly verification, valuable self-reports, and AI-driven reduction in the cost of producing plausible signals, characterizes many settings, from fake product reviews (Luca and Zervas, 2016; Mayzlin et al., 2014; He et al., 2022), academic fraud, financial misreporting, and environmental self-monitoring (Duflo et al., 2013, 2018). The deterrence-erosion multiplier applies specifically to settings where verification relies on volunteer or fixed-capacity reviewers, self-reports carry material value, and there is no scalable automated detection. Not all citizen science or self-reporting settings share these features, but those that do face the same nonlinear vulnerability. The policy implication is that deploying assistive AI without simultaneously scaling verification capacity is equivalent to lowering the expected cost of misconduct.

I cannot empirically distinguish deliberate fabricators from casual misidentifiers, and this limitation has policy implications. Raising penalties deters the former but not the latter, while platform-side interventions such as geographic plausibility filters and AI-provenance

tagging address both. The contamination estimates are cell-level, and I cannot flag individual fabricated records, which limits what platforms can do directly with these results. The analysis also focuses on Brazil, where birding tourism incentives and low reviewer density provide clean identification. Brazil's unusually low reviewer-to-submission ratio makes it a near-worst-case setting for the deterrence-erosion multiplier; in the United States and Europe, where eBird has substantially more reviewers per submission, the amplification would be attenuated. The mechanism should still operate wherever reviewer capacity is fixed and submissions grow, but the magnitude may be smaller.

References

Acemoglu, Daron, "Harms of AI," Working Paper 29247, National Bureau of Economic Research 2021.

_ , "The Simple Macroeconomics of AI," Working Paper 32487, National Bureau of Economic Research 2024.

Allcott, Hunt and Matthew Gentzkow, "Social Media and Fake News in the 2016 Election," *Journal of Economic Perspectives*, 2017, *31* (2), 211–236.

ao, Juliana Assunç Clarissa Gandour, and Romero Rocha, "Deforestation Slowdown in the Brazilian Amazon: Prices or Policies?," *Environment and Development Economics*, 2015, *20* (6), 697–722.

Balboni, Clare, Aaron Berman, Robin Burgess, and Benjamin A. Olken, "The Economics of Tropical Deforestation," *Annual Review of Economics*, 2023, *15*, 723–754.

Becker, Gary S., "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 1968, *76* (2), 169–217.

Bonney, Rick, Jennifer L. Shirk, Tina B. Phillips, Andrea Wiggins, Heidi L. Ballard, Abraham J. Miller-Rushing, and Julia K. Parrish, "Next Steps for Citizen Science," *Science*, 2014, *343* (6178), 1436–1437.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work," *Quarterly Journal of Economics*, 2025, *140* (2), 889–942.

Burgess, Robin, Francisco Costa, and Benjamin A. Olken, "The Brazilian Amazon's Double Reversal of Fortune," Technical Report, Working paper 2024.

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

Butler, Rhett Ayers, "AI-Generated Wildlife Photos Make Conservation More Difficult," *Mongabay* 2026. https://news.mongabay.com/short-article/2026/02/ai-generated-wildlife-photos-make-conservation-more-difficult/.

Callaghan, Corey T., Alistair G. B. Poore, Thomas Mesaglio, Angela T. Moles, Shinichi Nakagawa, Clint Roberts, Jodi J. L. Rowley, Adriana Vergeś, John H. Wilshire, and William K. Cornwell, "Three Frontiers for the Future of Biodiversity Research Using Citizen Science Data," *BioScience*, 2021, *71* (1), 55–63.

_ **, Carly Winnebald, Blaze Smith, Brittany M. Mason, and Laura López-Hoffman**, "Citizen Science as a Valuable Tool for Environmental Review," *Frontiers in Ecology and the Environment*, 2025, *23* (1), e2808.

_ , **Shinichi Nakagawa, and William K. Cornwell**, "Global Abundance Estimates for 9,700 Bird Species," *Proceedings of the National Academy of Sciences*, 2021, *118* (21), e2023170118.

Carreira, Igor, Francisco Costa, and João Paulo Pessoa, "The Deforestation Effects of Trade and Agricultural Productivity in Brazil," *Journal of Development Economics*, 2024, *167.*

Chen, Yuyu, Guo Zhong Jin, Naresh Kumar, and Guang Shi, "Gaming in Air Pollution Data? Lessons from China," *B.E. Journal of Economic Analysis & Policy*, 2012, *13* (3), 1–43.

Crawford, Vincent P. and Joel Sobel, "Strategic Information Transmission," *Econometrica*, 1982, *50* (6), 1431–1451.

Dellarocas, Chrysanthos, "Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms," *Management Science*, 2006, *52* (10), 1577–1593.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan, "Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India," *Quarterly Journal of Economics*, 2013, *128* (4), 1499–1545.

_ , _ , _ , **and** _ , "The Value of Regulatory Discretion: Estimates from Environmental Inspections in India," *Econometrica*, 2018, *86* (6), 2123–2160.

Einav, Liran, Chiara Farronato, and Jonathan Levin, "Peer-to-Peer Markets," *Annual Review of Economics*, 2016, *8*, 615–635.

Ferraro, Paul J. and Subhrendu K. Pattanayak, "Money for Nothing? A Call for Empirical Evaluation of Biodiversity Conservation Investments," *PLOS Biology*, 2006, *4* (4), e105.

Filippas, Apostolos, John J. Horton, and Joseph Golden, "Reputation Inflation," *Marketing Science*, 2022, *41* (4), 733–745.

Fisman, Raymond and Shang-Jin Wei, "Tax Rates and Tax Evasion: Evidence from "Missing Imports" in China," *Journal of Political Economy*, 2004, *112* (2), 471–500.

Fraisl, Dilek, Gerid Hager, Baptiste Bedessem, Margaret Gold, Pen-Yuan Hsing, Finn Danielsen, Colleen Hitchcock et al., "Citizen Science in Environmental and Ecological Sciences," *Nature Reviews Methods Primers*, 2022, *2*, 64.

Ghanem, Dalia and Junjie Zhang, "'Effortless Perfection:' Do Chinese Cities Manipulate Air Pollution Data?," *Journal of Environmental Economics and Management*, 2014, *68* (2), 203–225.

Goldstein, Josh A., Girish Sastry, Micah Musser, Renee DiResta, Matthew Gentzel, and Katerina Sedova, "Generative Language Models and Automated Influence Operations: Emerging Threats and Potential Mitigations," *arXiv preprint arXiv:2301.04246*, 2023.

Gray, Wayne B. and Mary E. Deily, "Compliance and Enforcement: Air Pollution Regulation in the U.S. Steel Industry," *Journal of Environmental Economics and Management*, 1996, *31* (1), 96–111.

Greenstone, Michael, Guojun He, Ruixue Jia, and Tong Liu, "Can Technology Solve the Principal-Agent Problem? Evidence from China's War on Air Pollution," *American Economic Review: Insights*, 2022, *4* (1), 54–70.

Growcoot, Matt, "AI is Ruining Nature Photography on Facebook," *PetaPixel* 2025. https://petapixel.com/2025/08/07/ai-is-ruining-nature-photography-on-facebook/.

Harrington, Winston, "Enforcement Leverage When Penalties Are Restricted," *Journal of Public Economics*, 1988, *37* (1), 29–53.

He, Guojun, Shaoda Wang, and Bing Zhang, "Watering Down Environmental Regulation in China," *Quarterly Journal of Economics*, 2020, *135* (4), 2135–2185.

He, Sherry, Brett Hollenbeck, and Davide Proserpio, "The Market for Fake Reviews," *Marketing Science*, 2022, *41* (5), 896–921.

Jacob, Brian A. and Steven D. Levitt, "Rotten Apples: An Investigation of the Prevalence and Predictors of Teacher Cheating," *Quarterly Journal of Economics*, 2003, *118* (3), 843–877.

Jetz, Walter, Jana M. McPherson, and Robert P. Guralnick, "Integrating Biodiversity Distribution Knowledge: Toward a Global Map of Life," *Trends in Ecology and Evolution*, 2012, *27* (3), 151–159.

Johnston, Alison, Wesley M. Hochachka, Matthew E. Strimas-Mackey, Viviana Ruiz Gutierrez, Orin J. Robinson, Erica T. Miller, Tom Auer, Steve T. Kelling, and Daniel Fink, "Analytical Guidelines to Increase the Value of Community Science Data: An Example Using eBird Data to Estimate Species Distributions," *Diversity and Distributions*, 2021, *27* (7), 1265–1277.

Kartik, Navin, "Strategic Communication with Lying Costs," *Review of Economic Studies*, 2009, *76* (4), 1359–1395.

Kelling, Steve, Alison Johnston, Wesley M. Hochachka, Marshall Iliff, Daniel Fink, Jeff Gerbracht, Carl Lagoze, Frank A. La Sorte, Travis Moore, Andrea Wiggins, Weng-Keen Wong, Chris Wood, and Jun Yu, "Can Observation Skills of Citizen Scientists Be Estimated Using Species Accumulation Curves?," *PLOS ONE*, 2015, *10* (10), e0139600.

Kosmala, Margaret, Andrea Wiggins, Alexandra Swanson, and Brooke Simmons, "Assessing Data Quality in Citizen Science," *Frontiers in Ecology and the Environment*, 2016, *14* (10), 551–560.

Liang, Yuanning, Ivan Rudik, and Eric Zou, "Economic Activity and Biodiversity in the United States," *Journal of Public Economics*, 2025, *250*, 105482.

_ , _ , Eric Yongchen Zou, Alison Johnston, Amanda D. Rodewald, and Catherine L. Kling, "Conservation Cobenefits from Air Pollution Regulation: Evidence from Birds," *Proceedings of the National Academy of Sciences*, 2020, *117* (49), 30900–30906.

Lintott, Chris J., Kevin Schawinski, Anže Slosar, Kate Land, Steven Bamford, Daniel Thomas, M. Jordan Raddick, Robert C. Nichol, Alex Szalay, Dan Andreescu et al., "Galaxy Zoo: Morphologies Derived from Visual Inspection of Galaxies from the Sloan Digital Sky Survey," *Monthly Notices of the Royal Astronomical Society*, 2008, *389* (3), 1179–1189.

Luca, Michael and Georgios Zervas, "Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud," *Management Science*, 2016, *62* (12), 3412–3427.

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier, "Promotional Reviews: An Empirical Investigation of Online Review Manipulation," *American Economic Review*, 2014, *104* (8), 2421–2455.

Mookherjee, Dilip and Ivan P. L. Png, "Corruptible Law Enforcers: How Should They Be Compensated?," *Economic Journal*, 1995, *105* (428), 145–159.

Murabayashi, Allen, "What Does Generative AI Mean for Bird and Nature Photography?," *Audubon Magazine*, Summer 2023 2023. https://www.audubon.org/magazine/what-does-generative-ai-mean-bird-and-nature-photography.

Nosko, Chris and Steven Tadelis, "The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment," *NBER Working Paper 20830*, 2015.

Olken, Benjamin A., "Monitoring Corruption: Evidence from a Field Experiment in Indonesia," *Journal of Political Economy*, 2007, *115* (2), 200–249.

Polinsky, A. Mitchell and Steven Shavell, "The Optimal Tradeoff between the Probability and Magnitude of Fines," *American Economic Review*, 1979, *69* (5), 880–891.

Rodrigues, Ana S. L., John D. Pilgrim, John F. Lamoreux, Michael Hoffmann, and Thomas M. Brooks, "The Value of the IUCN Red List for Conservation," *Trends in Ecology and Evolution*, 2006, *21* (2), 71–76.

Souza, Carlos M., Julia Z. Shimbo, Marcos R. Rosa, Leandro L. Parente, Ane A. Alencar, Bernardo F. T. Rudorff, Heinrich Hasenack, Marcelo Matsumoto, Laerte G. Ferreira, Pedro W. M. Souza-Filho et al., "Reconstructing Three Decades of Land Use and Land Cover Changes in Brazilian Biomes with Landsat Archive and Earth Engine," *Remote Sensing*, 2020, *12* (17), 2735.

Stigler, George J., "The Optimum Enforcement of Laws," *Journal of Political Economy*, 1970, *78* (3), 526–536.

Sullivan, Brian L., Christopher L. Wood, Marshall J. Iliff, Rick E. Bonney, Daniel Fink, and Steve Kelling, "eBird: A Citizen-Based Bird Observation Network in the Biological Sciences," *Biological Conservation*, 2009, *142* (10), 2282–2292.

_ , Jocelyn L. Aycrigg, Jessie H. Barry, Rick E. Bonney, Nicholas Bruns, Caren B. Cooper, Theo Damoulas, André A. Dhondt, Tom Dietterich, Andrew Farnsworth et al., "The eBird Enterprise: An Integrated Approach to Development and Application of Citizen Science," *Biological Conservation*, 2014, *169*, 31–40.

Sun, Alice, "Merlin Is Magical, but It Still Makes Mistakes," *Audubon Magazine* 2023. https://www.audubon.org/magazine/merlin-magical-it-still-makes-mistakes.

Tadelis, Steven, "Reputation and Feedback Systems in Online Platform Markets," *Annual Review of Economics*, 2016, *8*, 321–340.

Tella, Rafael Di and Ernesto Schargrodsky, "The Role of Wages and Auditing during a Crackdown on Corruption in the City of Buenos Aires," *Journal of Law and Economics*, 2003, *46* (1), 269–292.

Zou, Eric Yongchen, "Unwatched Pollution: The Effect of Intermittent Monitoring on Air Quality," *American Economic Review*, 2021, *111* (7), 2101–2126.

Table 1: Summary Statistics

	Mean	SD	Min	Max	N
Panel A: Triple DiD sample					
log(records + 1)	2.806	2.114	0	10.23	19,214
Review rate	0.052	0.145	0	1	19,214
Low forest (binary)	0.500	0.500	0	1	19,214
Rare (binary)	0.500	0.500	0	1	19,214
Forest share (continuous)	0.382	0.283	0	0.98	19,214
Panel B: Baseline DiD sample					
Review rate	0.048	0.185	0	1	32,657
Panel C: Enriched sample (manipulation signals)					
Novel location rate	0.823	0.278	0	1	22,616
Private location rate	0.673	0.394	0	1	22,616
Suspicious checklist rate	0.068	0.206	0	1	22,616

Notes: Panel A: unit is municipality × rarity × year, restricted to municipalities matched to MapBiomas. Panel B: all municipalities with ≥1 eBird record. Panel C: enriched sample with computed manipulation indicators; the larger sample ($N = 22{,}616$) includes all municipalities with ≥1 eBird record that have sufficient data to compute location novelty and checklist flags, regardless of MapBiomas match status. Sample period: 2017-2025.

Table 2: Baseline Difference-in-Differences: Review Rate

	Review rate
	(1)
Post-2021 × Rare	-0.061^{***}
	(0.010)
Observations	32,657
R^2	0.671
Y-mean	0.048
Y-sd	0.185
Fixed effects	Mun × rarity, Year

Notes: Unit of observation is municipality × rarity × year. Sample includes all municipalities with at least one eBird record. Rare species defined as the bottom tercile of pre-2021 national observation counts. Standard errors clustered at the municipality level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Triple Difference: Post × Rare × Low Forest

	log(records) (1)	Review rate (2)	Novel loc. rate (3)
Post × Rare	0.041	−0.012	0.010
	(0.043)	(0.010)	(0.013)
Post × Low Forest	−0.082*	0.006***	−0.002
	(0.049)	(0.002)	(0.008)
Post × Rare × Low Forest	0.200***	−0.097***	−0.026
	(0.071)	(0.024)	(0.024)
Observations	19,214	19,214	19,214
R^2	0.762	0.726	0.648
Y-mean	2.806	0.052	0.823
Fixed effects		Mun × rarity, Year	

Notes: Unit of observation is municipality × rarity × year. Sample restricted to municipalities matched to MapBiomas forest data. Low Forest defined as below-median average forest share over 2017-2020. Standard errors clustered at the municipality level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: Heterogeneity Analysis

	H1: Night		H5: Mobile/Urban			
	log(n)	Rev. rate	log(n)	Rev. rate		
	(1)	(2)	(3)	(4)		
Panel A: Confirmed predictions						
Post × Rare	0.181***	−0.054***	0.144***	−0.021**		
	(0.035)	(0.011)	(0.048)	(0.010)		
Post × Dimension	−0.389***	0.006***	0.119**	0.005***		
	(0.031)	(0.002)	(0.049)	(0.002)		
Post × Rare × Dimension	0.204***	0.008	0.028	−0.053***		
	(0.044)	(0.014)	(0.065)	(0.019)		
Obs	31,868	31,868	22,601	22,601		
R^2	0.757	0.698	0.757	0.699		
Y-mean	2.570	0.051	2.786	0.052		
	H2: Endemic		H3: Charismatic		H4: Novel mun	H6: Protected
	log(n)	Rev. rate	log(n)	Rev. rate	Novel rate	log(n)
Panel B: Null predictions confirmed						
Post × Rare × Dimension	−0.004	−0.000	−0.013	0.002	−0.031***	−0.001
	(0.022)	(0.002)	(0.017)	(0.002)	(0.011)	(0.067)
Obs	27,471	27,471	30,078	30,078	22,616	22,616
R^2	0.756	0.510	0.758	0.456	0.720	0.756
Y-mean	2.665	0.016	2.611	0.015	0.611	2.786

Notes: Each column is a separate regression. Panel A reports the full triple interaction for confirmed model predictions. Panel B reports the interaction coefficient for null predictions. Sample sizes vary across dimensions because each heterogeneity variable restricts the sample differently: endemic and charismatic species indicators apply only to species with those traits, while novel municipality and protected area indicators apply to all species. Night = 22:00-06:00. Urban = above-median 4G/5G coverage. Endemic = endemic-rich family. Charismatic = iconic large-bodied species. Novel = first municipality record. Protected = above-median ICMBio conservation area. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5: Mechanism: Reviewer Dilution Tests

	Common spp. only	Rare spp. + volume controls	
	Review rate	Baseline	+ Volume × Post
	(1)	(2)	(3)
Post-AI	−0.001		
	(0.001)		
Post × Rare		−0.051***	−0.048***
		(0.010)	(0.010)
log(common volume)			0.000
			(0.001)
Post × log(common volume)			−0.001
			(0.001)
Observations	17,752	22,565	22,565
R^2	0.385	0.698	0.698
Y-mean	0.013	0.051	0.051
Fixed effects	Mun, Year	Mun × rarity, Year	

Notes: Column (1) restricts to common species only (top two terciles). Columns (2)–(3) use rare species review rate as the outcome and add controls for common species volume. Less than 6% coefficient shrinkage from baseline to Column (3). Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6: Review Rate Decomposition: Post × Rare × Low Forest

	Reviewed	Unreviewed	TGTB rate	Novel loc. rate
	(1)	(2)	(3)	(4)
Post × Rare	−0.012	0.012	−0.021*	0.010
	(0.010)	(0.010)	(0.011)	(0.013)
Post × Low Forest	0.006***	−0.006***	0.011***	−0.002
	(0.002)	(0.002)	(0.004)	(0.008)
Post × Rare × Low Forest	−0.097***	0.097***	0.005	−0.026
	(0.024)	(0.024)	(0.020)	(0.024)
Observations	19,214	19,214	19,214	19,214
R^2	0.726	0.726	0.712	0.648
Y-mean	0.052	0.948	0.070	0.823
Fixed effects	Mun × rarity, Year			

Notes: Each column is a separate regression. Reviewed = share flagged for expert review. Unreviewed = 1 − Reviewed; the −0.097 coefficient in Column (1) matches Table 3 Column (2) by construction, as both estimate the same specification. TGTB = share of checklists with ≥2 rare species. Novel location = share at previously unrecorded locations. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 7: Review Evasion: TGTB vs. Non-TGTB Checklists

	log(TGTB records)	Review | TGTB	Review | non-TGTB
	(1)	(2)	(3)
Post × Rare × Low Forest	0.155*	−0.013	−0.105***
	(0.087)	(0.026)	(0.026)
Observations	2,796	2,796	18,923
R^2	0.832	0.828	0.713
Y-mean	2.877	0.092	0.051
Fixed effects		Mun × rarity, Year	

Notes: TGTB (too-good-to-be-true) = checklists with ≥2 rare species. Column (2): review rate among TGTB checklists, which continue to attract scrutiny. Column (3): review rate among non-TGTB checklists, where the entire review rate decline concentrates. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 8: Mechanism: Tourism Incentive

	Binary: High birding		Continuous: log(pre-checklists)	
	log(records)	Review rate	log(records)	Review rate
	(1)	(2)	(3)	(4)
Post × Rare	−0.151*	−0.007	0.089**	−0.029***
	(0.078)	(0.027)	(0.035)	(0.010)
Post × Tourism	−0.486***	0.002	−0.156***	0.001
	(0.047)	(0.003)	(0.018)	(0.001)
Post × Rare × Tourism	0.455***	−0.047	0.236***	−0.024***
	(0.086)	(0.029)	(0.024)	(0.008)
Observations	19,848	19,848	19,848	19,848
R^2	0.755	0.691	0.754	0.692
Y-mean	2.897	0.055	2.897	0.055
Fixed effects		Mun × rarity, Year		

Notes: Columns (1)–(2): High-birding = above-median pre-2021 checklist density. Columns (3)–(4): log(pre-checklists) is continuous log of total eBird checklists in the municipality over 2017-2020. The exact percentage effect in Column (1) is $e^{0.455} - 1 = 57.6\%$; I report 58% in the text. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 9: Mechanism: Weekend Patterns and Organized Guides

	Weekend effect		Weekend × Tourism	
	log(n)	Rev. rate	log(n)	Rev. rate
	(1)	(2)	(3)	(4)
Post × Rare × Weekend	0.094***	0.010	0.310***	−0.013
	(0.033)	(0.012)	(0.064)	(0.023)
Post × Rare × High birding			0.134*	−0.046
			(0.075)	(0.028)
Post × Rare × Wknd × HiBird			−0.218***	0.025
			(0.059)	(0.022)
Observations	34,901	34,901	31,261	31,261
R^2	0.715	0.647	0.713	0.641
Y-mean	2.584	0.050	2.667	0.053
Fixed effects	Mun × rarity × weekend, Year			

Notes: Weekend = Saturday–Sunday. HiBird = above-median pre-2021 birding tourism. Columns (1)–(2) use the full weekend/weekday sample; Columns (3)–(4) restrict to municipalities with non-missing tourism data, reducing the sample. The weekend effect concentrates in non-tourism areas (Column (3)), inconsistent with organized guide fraud. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: Mechanism: Charismatic Species and Seasonality

	Charismatic species		Dry season	
	log(n)	Rev. rate	log(n)	Rev. rate
	(1)	(2)	(3)	(4)
Post × Rare	0.135***	−0.042***	0.160***	−0.049***
	(0.034)	(0.010)	(0.035)	(0.011)
Post × Dimension			0.044*	−0.002
			(0.025)	(0.002)
Post × Rare × Charismatic	−0.021	−0.022		
	(0.050)	(0.019)		
Post × Rare × Dry season			−0.020	0.005
			(0.039)	(0.014)
Observations	36,059	36,059	33,767	33,767
R^2	0.772	0.716	0.702	0.644
Y-mean	2.430	0.041	2.622	0.049
Fixed effects	Mun × rarity × char, Year		Mun × rarity × season, Year	

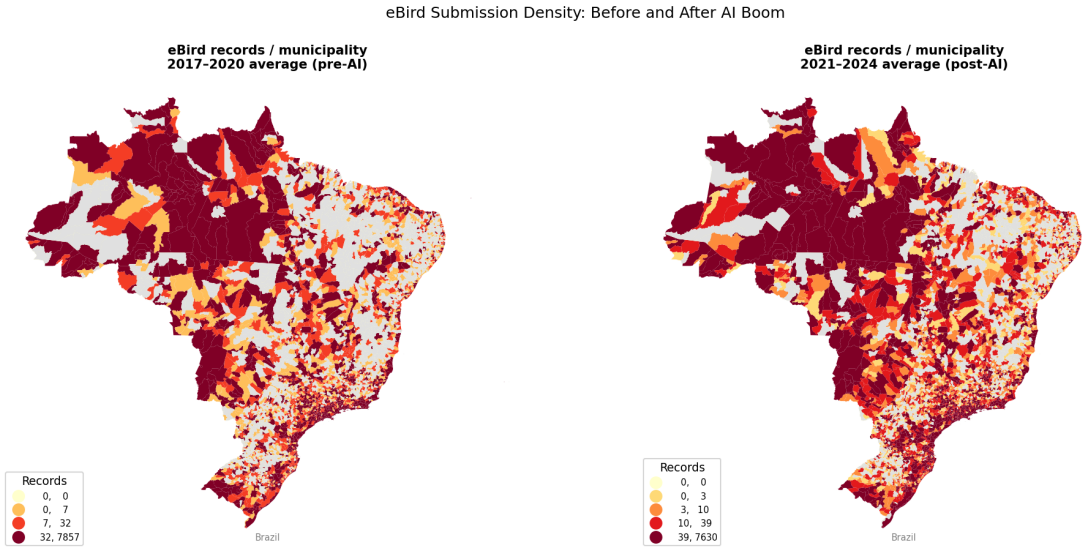
Notes: Charismatic = iconic large-bodied species (raptors, macaws, toucans). Dry season = May–September, the peak birding tourism season. Both null results are consistent with indiscriminate AI misidentification rather than targeted tourism fraud. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 11: Reviewer Capacity and Workload

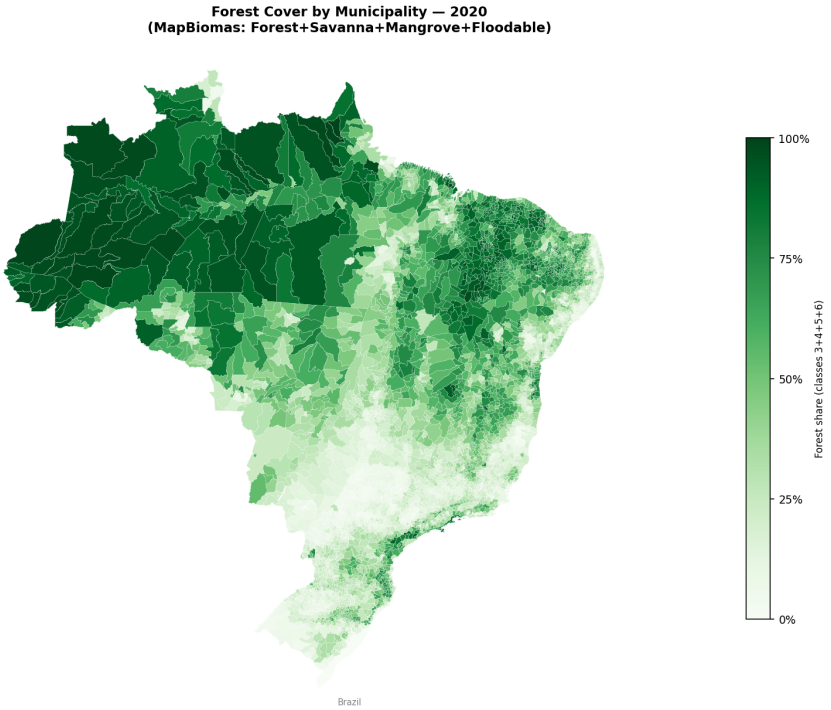
	log(workload)	Unreviewed
	(1)	(2)
Post × Rare	0.002	0.012
	(0.047)	(0.010)
Post × Low Forest	−0.122**	−0.006***
	(0.051)	(0.002)
Post × Rare × Low Forest	0.200***	0.097***
	(0.073)	(0.024)
Observations	19,214	19,214
R^2	0.727	0.726
Y-mean	2.376	0.948
Fixed effects	Mun × rarity, Year	

Notes: log(workload) = log total records per reviewer in the municipality's review region. Unreviewed rate = 1 − review rate. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 1: Cross-Sectional Distribution of eBird Activity and Forest Cover



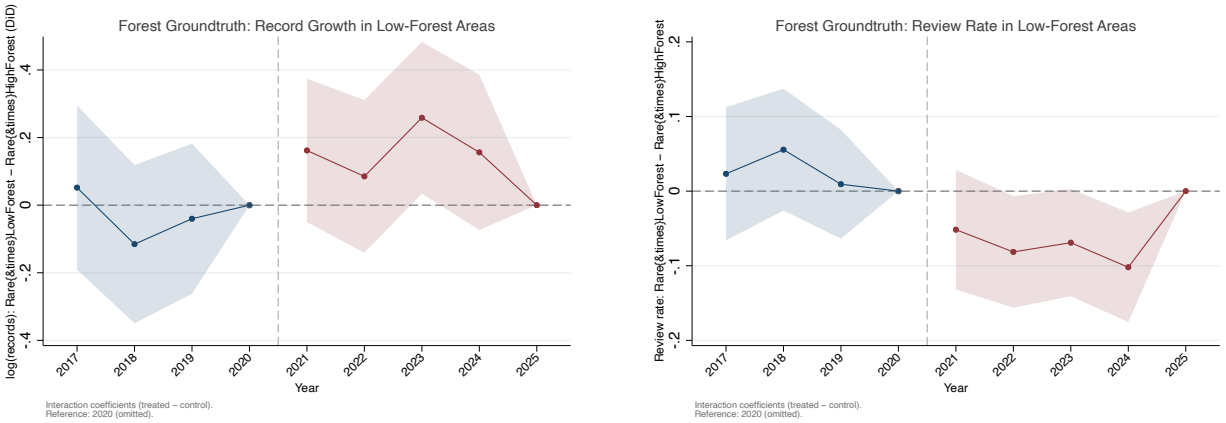
(a) eBird record density by municipality



(b) Pre-treatment forest cover

Notes: Panel (a): log of total eBird observation records by municipality over 2017-2020. Darker shading indicates higher record density. Panel (b): municipality-level forest share from MapBiomas Collection 9, averaged over 2017-2020. Green shading indicates higher forest cover.

Figure 2: Event Study: Triple DiD Coefficients

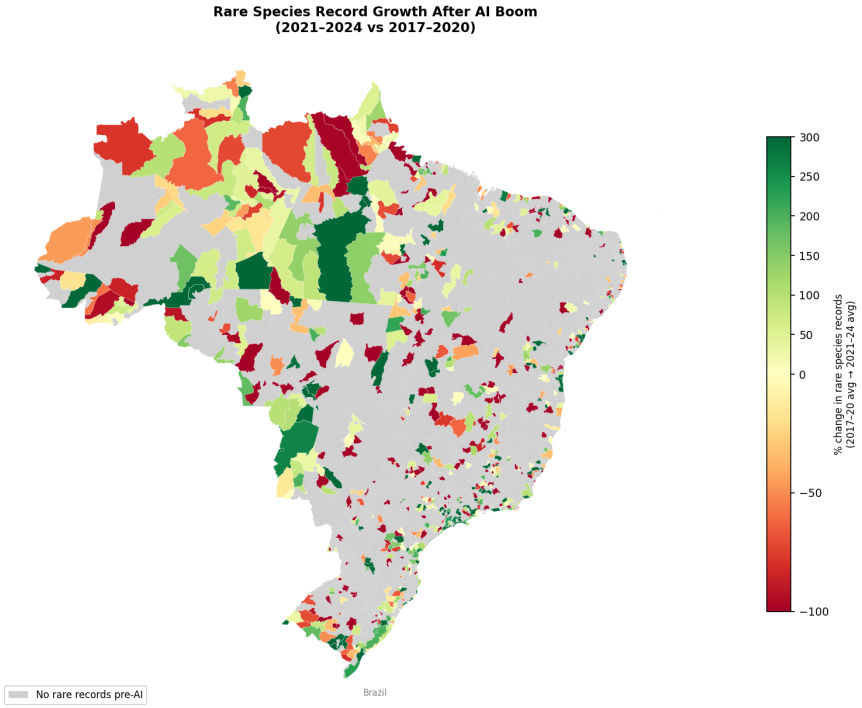


(a) log(records)

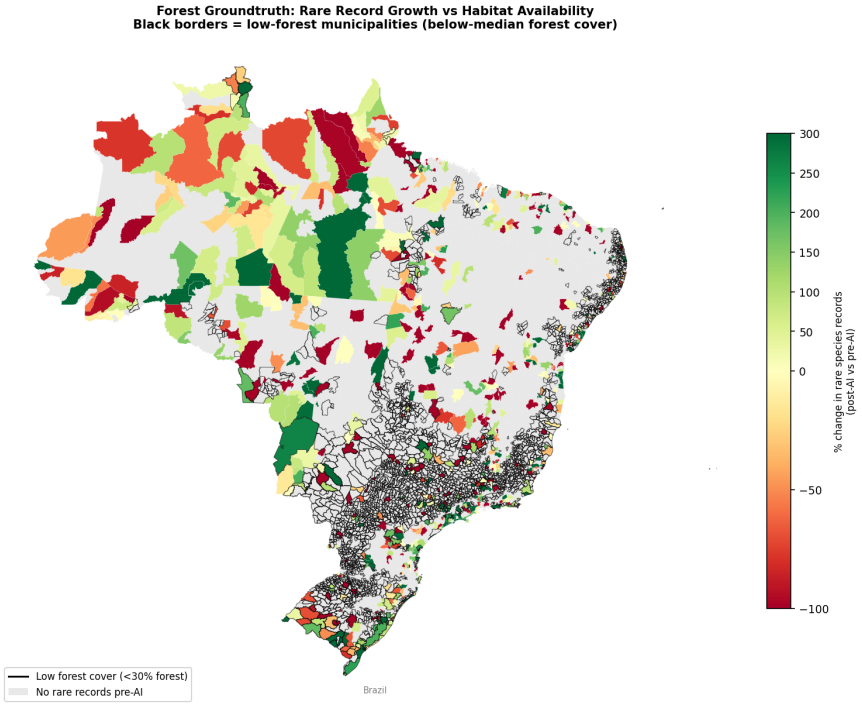
(b) Review rate

Notes: Coefficients from equation (8) with 2020 as the omitted reference year. Shaded areas show 95% confidence intervals based on standard errors clustered at the municipality level. Panel (a): outcome is log(records + 1). Panel (b): outcome is the share of records flagged for expert review.

Figure 3: Geographic Pattern of Rare Record Growth, Post-AI versus Pre-AI



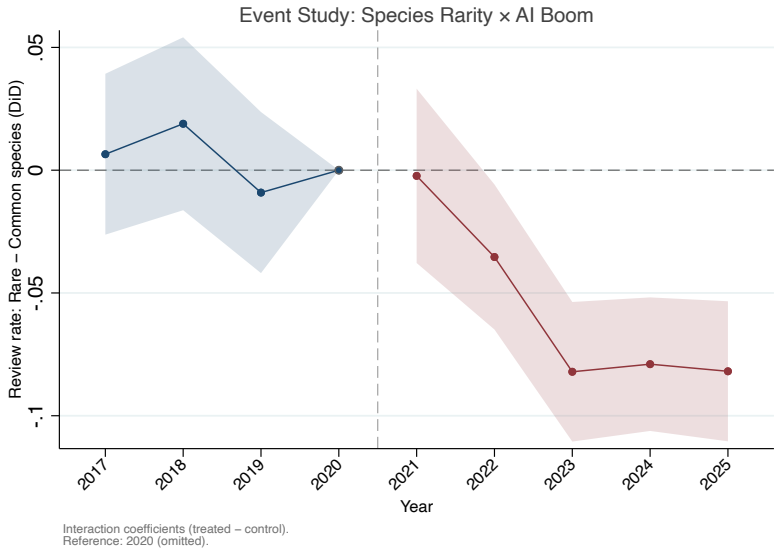
(a) Rare species record growth, percent change



(b) Low-forest municipalities highlighted

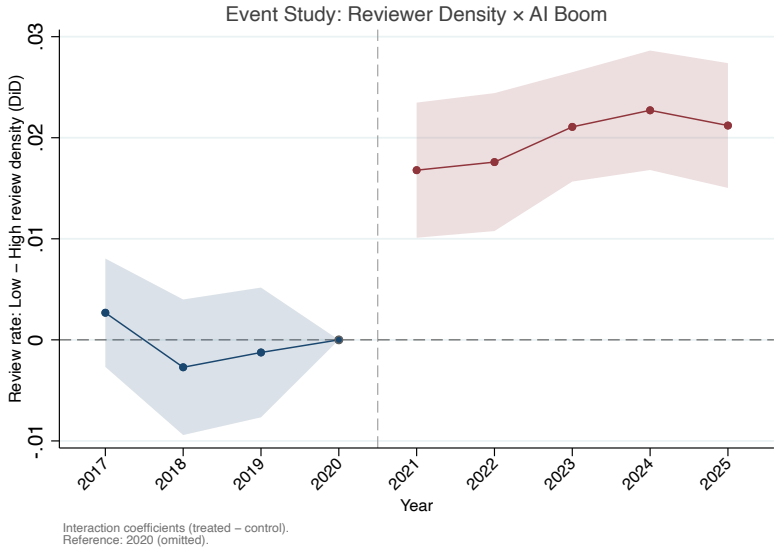
Notes: Panel (a): percentage change in rare species records by municipality, comparing post-AI, 2021-2025, to pre-AI, 2017-2020. Darker shading indicates larger increases. Panel (b): same map with low-forest municipalities, defined as below-median pre-treatment forest share, outlined in black.

Figure 4: Event Study: Post × Rare Coefficients on Review Rate



Notes: Coefficients from the baseline DiD event study with 2020 as the omitted reference year. The outcome is the expert review rate. Pre-2021 coefficients are flat and statistically indistinguishable from zero. A sharp decline begins at 2021, the year of the AI boom, and grows in magnitude through 2024.

Figure 5: Event Study: Reviewer Workload



Notes: Coefficients from an event study of reviewer workload. Pre-2021 coefficients are flat. Post-2021, reviewer workload increases sharply as the volume of submissions rises, confirming that the review system is overwhelmed rather than scaled.

Appendix: Supplementary Tables and Figures

Table S1: Heterogeneity: Night Submissions (H1)

	log(records) (1)	Review rate (2)
Post × Rare	0.181^{***}	-0.054^{***}
	(0.035)	(0.011)
Post × Night	-0.389^{***}	0.006^{***}
	(0.031)	(0.002)
Post × Rare × Night	0.204^{***}	0.008
	(0.044)	(0.014)
Observations	31,868	31,868
R^2	0.757	0.698
Y-mean	2.570	0.051
Fixed effects	Mun × rarity × night, Year	

Notes: Night defined as submissions between 22:00 and 06:00. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table S2: Heterogeneity: Mobile Coverage / Urban (H5)

	log(records) (1)	Review rate (2)
Post × Rare	0.144***	−0.021**
	(0.048)	(0.010)
Post × Urban	0.119**	0.005***
	(0.049)	(0.002)
Post × Rare × Urban	0.028	−0.053***
	(0.065)	(0.019)
Observations	22,601	22,601
R^2	0.757	0.699
Y-mean	2.786	0.052
Fixed effects	Mun × rarity, Year	

Notes: Urban defined as municipality with above-median urbanization rate, proxying for 4G/5G mobile coverage required by AI identification apps. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table S3: Heterogeneity: Protected Areas (H6)

	log(records) (1)	Review rate (2)
Post × Rare	0.122**	−0.029*
	(0.052)	(0.015)
Post × Protected	0.118**	−0.002
	(0.047)	(0.002)
Post × Rare × Protected	−0.001	−0.031
	(0.067)	(0.020)
Observations	22,616	22,616
R^2	0.756	0.699
Y-mean	2.786	0.052
Fixed effects	Mun × rarity, Year	

Notes: Protected defined as municipality with above-median share of ICMBio conservation unit area. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table S4: Heterogeneity: Night × Tourism Four-Way Interaction

	log(records) (1)	Review rate (2)
Post × Rare × Night	0.349**	0.036
	(0.153)	(0.050)
Post × Rare × HiBird	0.452***	−0.046
	(0.093)	(0.043)
Post × Rare × Night × HiBird	−0.247	−0.026
	(0.160)	(0.052)
Observations	28,467	28,467
Y-mean	2.654	0.054
Fixed effects	Mun × rarity × night, Year	

Notes: Four-way interaction testing whether night-time fabrication is stronger in high-tourism areas. The four-way term is negative but imprecise, consistent with the night effect being driven by casual users broadly rather than tourism areas specifically. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table S5: Robustness: COVID-19 Controls and Alternative Treatment Timing

	log(records)		Review rate	
	Treat=2021 (1)	Treat=2022 (2)	Treat=2021 (3)	Treat=2022 (4)
Post × Rare × Low Forest	0.200***	0.158**	−0.097***	−0.087***
	(0.071)	(0.070)	(0.024)	(0.022)
Observations	19,214	19,214	19,214	19,214
R^2	0.762	0.762	0.726	0.726
Y-mean	2.806	2.806	0.052	0.052
Fixed effects	Mun × rarity, Year			

Notes: Columns (1) and (3) use 2021 as the treatment year (baseline). Columns (2) and (4) shift treatment to 2022, past the main COVID recovery period. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table S6: Robustness: Google Mobility Controls

	log(records)		Review rate	
	Baseline	+ Mobility	Baseline	+ Mobility
	(1)	(2)	(3)	(4)
Post × Rare × Low Forest	0.200***	0.208***	−0.097***	−0.100***
	(0.071)	(0.071)	(0.024)	(0.024)
Observations	19,214	19,214	19,214	19,214
R^2	0.762	0.762	0.726	0.726
Y-mean	2.806	2.806	0.052	0.052
Fixed effects	Mun × rarity, Year			

Notes: Columns (2) and (4) add Google Community Mobility "parks and recreation" index and its interactions with rarity and forest cover. Coefficients essentially unchanged. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table S7: Robustness: Continuous Forest Cover Specification

	log(records)		Review rate	
	Binary	Continuous	Binary	Continuous
	(1)	(2)	(3)	(4)
Post × Rare × Low Forest	0.200***		−0.097***	
	(0.071)		(0.024)	
Post × Rare × (1 − Forest %)		0.267**		−0.120***
		(0.129)		(0.035)
Observations	19,214	19,214	19,214	19,214
R^2	0.762	0.762	0.726	0.725
Y-mean	2.806	2.806	0.052	0.052
Fixed effects	Mun × rarity, Year			

Notes: Columns (1) and (3) use the binary Low Forest indicator (baseline). Columns (2) and (4) replace it with continuous (1 − forest share), so that higher values indicate less forest (more ecological implausibility). Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table S8: Robustness: Rarity Threshold Sensitivity

	P20 (1)	P25 (2)	P33 (baseline) (3)	P40 (4)
Panel A: log(records)				
Post × Rare × Low Forest	0.280***	0.171*	0.195***	0.151**
	(0.105)	(0.089)	(0.071)	(0.059)
Panel B: Review rate				
Post × Rare × Low Forest	−0.198***	−0.121***	−0.092***	−0.043**
	(0.051)	(0.036)	(0.024)	(0.017)
Observations	16,926	17,715	19,150	20,878
Y-mean (log)	2.949	2.893	2.809	2.735
Y-mean (rev)	0.045	0.052	0.052	0.046
Fixed effects		Mun × rarity, Year		

Notes: Each column uses a different rarity percentile cutoff. P20 = bottom 20% of pre-2021 species observation counts (strictest). P40 = bottom 40% (most lenient). Stricter thresholds yield larger coefficients, consistent with the model prediction. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table S9: Robustness: Dropping 2025

	Baseline	Triple DiD		Manip. signals		
	Rev. rate (1)	log(n) (2)	Rev. rate (3)	Novel (4)	Private (5)	TGTB (6)
Post × Rare	−0.046***	0.041	−0.012	0.002	−0.013	−0.021**
	(0.010)	(0.043)	(0.010)	(0.011)	(0.013)	(0.009)
Post × Low forest		−0.082*	0.006***			
		(0.049)	(0.002)			
Post × R × LF		0.200***	−0.097***			
		(0.071)	(0.024)			
Observations	19,214	19,214	19,214	19,214	19,214	19,214
Y-mean	0.052	2.806	0.052	0.823	0.664	0.070
Fixed effects			Mun × rarity, Year			

Notes: Sample restricted to 2017-2024 (2025 dropped). Coefficients virtually identical to baseline, confirming 2025 data do not drive results. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table S10: Fabrication Signal Tests: Post × Rare

	Novel location (1)	Private location (2)	Suspicious checklist (3)
Post × Rare	−0.000	−0.003	−0.027***
	(0.011)	(0.012)	(0.009)
Observations	22,616	22,616	22,616
Y-mean	0.823	0.673	0.068
Fixed effects		Mun × rarity, Year	

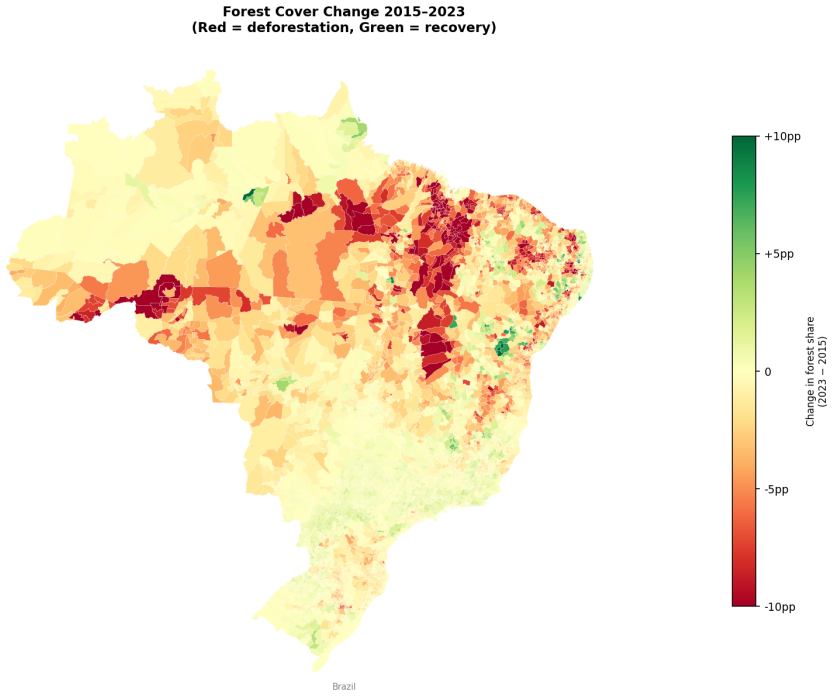
Notes: Column (1): share of records at novel locations. Column (2): share at private locations. Column (3): share of checklists with two or more rare species, a suspicious-checklist indicator. The suspicious checklist rate declines significantly post-AI, indicating that excess records do not trigger automated red flags. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table S11: Review Rate by Reviewer Density

	Review rate (1)
Post × Low density	0.020***
	(0.002)
Observations	26,063
Y-mean	0.017
Fixed effects	Mun × density, Year

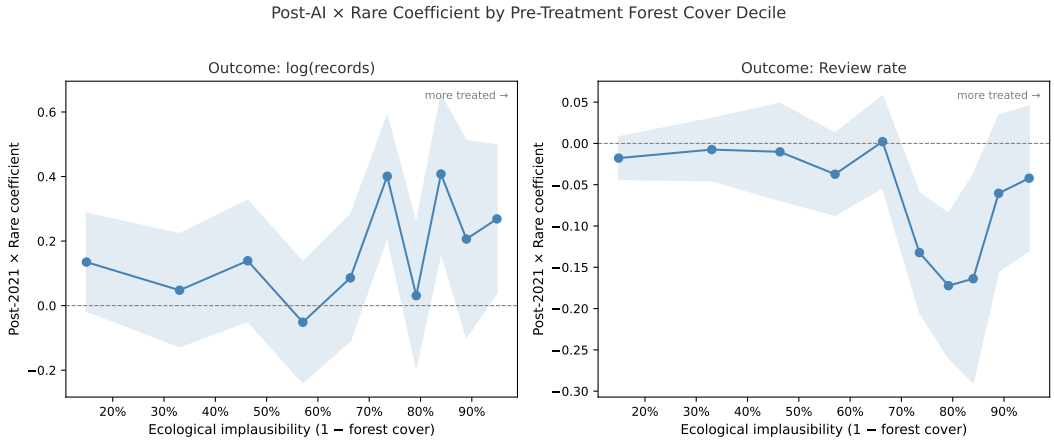
Notes: Low density = below-median pre-2021 reviewer count per municipality. Review rates *increase* in low-reviewer-density areas post-AI, confirming that the rare-species review rate decline documented in the main results is not driven by a general shortage of reviewers. Standard errors clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure S1: Forest Cover Change, 2015-2023



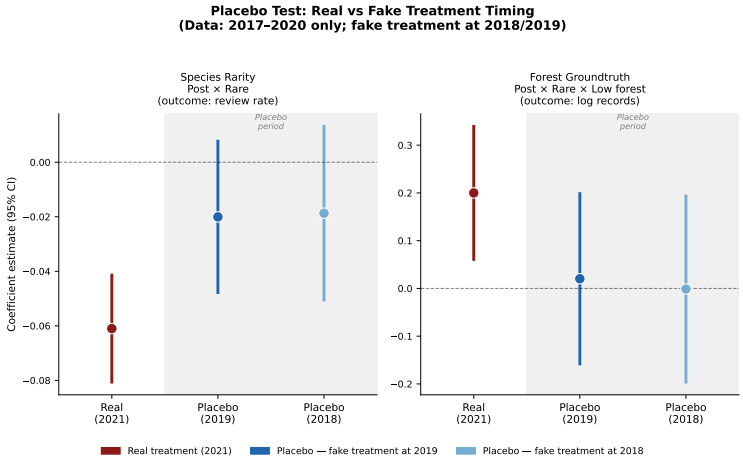
Notes: Municipality-level change in forest share from MapBiomas Collection 9, 2015-2023. Red shading indicates net deforestation; green indicates net forest recovery. Low-forest treatment municipalities continued to lose forest during the post-2021 period, ruling out habitat recovery as an explanation for increased rare species records.

Figure S2: Binscatter: Post × Rare Coefficient by Forest Cover Decile



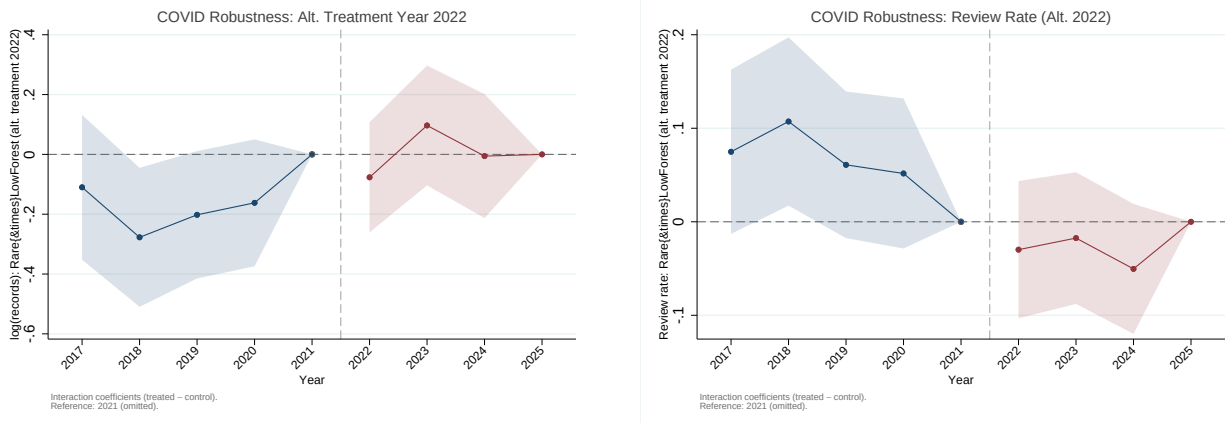
Notes: Each point represents the estimated Post × Rare coefficient within a forest cover decile. The x-axis is ecological implausibility (1 − forest share). The manipulation signal strengthens monotonically as ecological implausibility increases. Shaded area shows 95% confidence intervals.

Figure S3: Placebo Tests: False Treatment Dates, 2018 and 2019



Notes: Triple DiD estimated using 2018 and 2019 as false treatment dates on the pre-treatment, 2017-2020, sample. Both placebo coefficients are small and imprecisely estimated, confirming the absence of pre-existing differential trends.

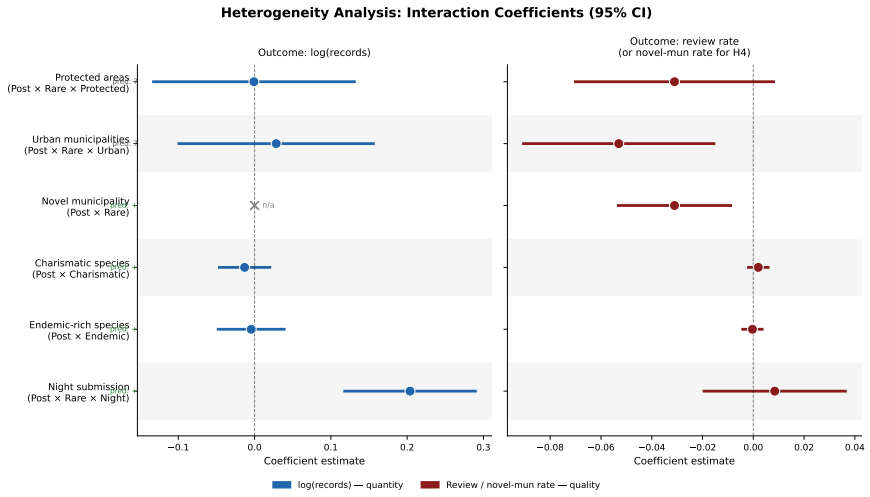
Figure S4: Event Study with 2022 as Treatment Year



(a) log(records) (b) Review rate

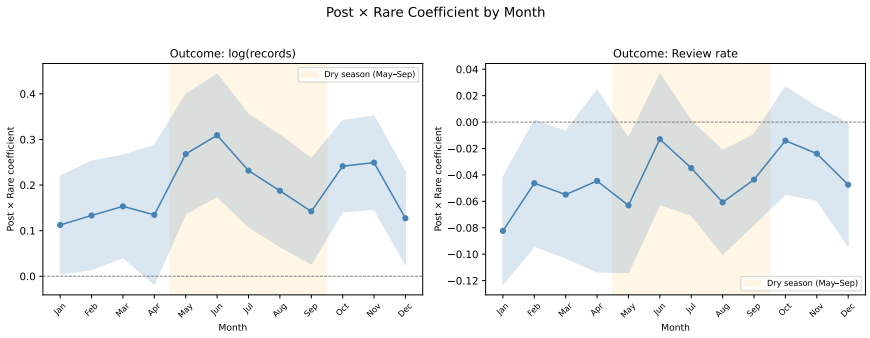
Notes: Event study coefficients with 2021 as the reference year. Pre-2022 coefficients, including 2021, are flat, confirming that the break occurs at AI adoption timing rather than COVID timing.

Figure S5: Heterogeneity: Coefficient Plot Across All Dimensions



Notes: Each point is the triple interaction coefficient from a separate regression. Horizontal bars show 95% confidence intervals.

Figure S6: Seasonal Pattern: Monthly Post × Rare Coefficients



Notes: Monthly coefficients for the Post × Rare interaction. No seasonal concentration during Brazil's dry and birding season, May through September. Manipulation is year-round, consistent with AI misidentification rather than tourism-driven seasonal fraud.

Data Availability

The eBird dataset used in this paper is available through a formal data request to the Cornell Lab of Ornithology (https://ebird.org/data/request). Access requires institutional affiliation, a project description, and agreement to eBird's terms of use; processing typically takes several weeks. MapBiomas forest cover data are publicly available at https://mapbiomas.org. Google Community Mobility data are publicly available through Google. ICMBio protected area shapefiles are publicly available from ICMBio's data portal (https://www.gov.br/icmbio). All replication code, including the API calls used to obtain the eBird data, and all derived datasets will be made available upon publication.