Permit Incentives and Taxonomic Precision: Strategic Underreporting in Deep-Sea Mining Data*

Andrés de Loera

Harvard University

Last updated: May 04, 2026

Abstract

The International Seabed Authority (ISA) requires deep-sea mining contractors operating in international waters beyond national jurisdiction to conduct biological surveys of their prospectively-mined tracts, creating a public good as a byproduct of private exploration. However, contractors face a principal-agent problem: they have incentives to underinvest in species detection if discoveries could create obstacles to mining approval. I test for this "taxonomic deflation" using over 160,000 biological samples from the ISA's DeepData database, comparing taxonomic reporting detail in prospectively-mined tracts versus conservation zones. I show that samples from conservation zones are significantly more likely to report species-level identification, a pattern that persists after controlling for contractor identity, year, and environmental variables. A counterfactual exercise suggests that tens of thousands of additional observations would include species identifications absent the observed reporting gap. I discuss the ISA's exploration requirement as a second-best solution to the deep sea's public goods problem and consider policy options for mitigating the resulting incentive distortions.

1 Introduction

Can the imprecision of deep sea mining data be explained by scientific difficulties alone? Biological exploration of the deep sea is genuinely difficult: depths exceeding four kilometers require specialized remotely operated vehicles, specimens often arrive damaged or chemically altered in ways that conceal diagnostic features (Kürzel et al. 2022), and the taxonomic expertise needed for species-level identification of deep-sea invertebrates is scarce (Frutos et al. 2022). Yet the parties responsible for most of this exploration are not entirely impartial—they are mining contractors seeking permits to extract nodules from the seafloor, whose

*I am grateful to Joe Aldy, Ed Glaeser, Wolfram Schlenker, Jim Stock, and Charles Taylor for invaluable guidance and feedback. I am also grateful for helpful comments and discussion from Karl Aspelund, Dale Squires, Jim Sanchirico, and participants of the Harvard Environmental Economics Workshop. All errors are my own.

commercial interests may not align with scientific thoroughness. The Clarion-Clipperton Zone (CCZ), the most active target for deep-sea mining, is estimated to harbor over 7,000 metazoan species, the vast majority unknown to science (Rabone, Wiethase, et al. 2023), and recent surveys continue to find that most newly collected specimens belong to previously undescribed taxa (Jażdżewska et al. 2025). In such conditions, coarse identification could reflect the limits of the discipline, but could also be driven by strategic incentives. A contractor whose sampling uncovers a species that triggers conservation review faces delays, regulatory redesigns, and potential permit denial. This paper argues that the observed imprecision of deep sea mining data is partially explained by strategic choices rather than scientific necessity alone.

While mining has commercial benefits, environmental impact assessment is critical to determining if, how, and where mining should be permitted. Scientific exploration is a prerequisite for an effective environmental impact assessment, but no entity has sufficient incentive to adequately explore the deep sea: mining contractors do not benefit from scientifically valuable discoveries that do not serve their commercial interests, such as the discovery of new species and traits, while third-party explorers do not bear the opportunity cost of delaying mining if they are slow in their exploration. The ISA's operative solution has been to require mining contractors to conduct scientific exploration as part of the process to get approval to mine. This paper discusses the logic and challenges of this approach, focusing on the perverse incentive it creates for mining companies to underexplore their tracts when possible. Specifically, I ask whether mining contractors engage in "taxonomic deflation" when exploring their prospectively-mined tracts, relative to non-contract areas.

The ISA requires mining contractors to submit biological data from their exploration areas to the DeepData database. I exploit the fact that the same contractors also occasionally submit data from areas outside their exploration contracts, including Areas of Particular Environmental Interest (APEIs) — collectively, non-contract areas — where no mining permit is at stake. APEIs are central to the comparison: they are large rectangular blocks sited under the ISA's explicit "representativity" criterion, designed to span the major environmental gradients of the CCZ rather than to surround unusual habitats. They are, by design, the closest available analogue to contract tracts. This within-contractor, across-zone comparison holds fixed the inherent difficulties of deep-sea taxonomy and the reporting infrastructure, isolating the effect of the permit incentive on reporting precision.

I first develop a theoretical model in which a contractor simultaneously chooses exploration intensity (how many samples to collect) and taxonomic precision (how specifically to identify the organisms recovered). The key insight is that these are distinct instruments with asymmetric regulatory incentives. More sampling unambiguously helps the contractor by demonstrating survey compliance, but taxonomic precision cuts both

ways: sufficient precision signals scientific credibility, while excessive precision risks uncovering species that trigger conservation review and jeopardize the mining permit. Outside contract areas, where no permit is at stake, this tension disappears. The model generates two testable predictions: contractors conduct more physical sampling on prospectively-mined tracts than on non-contract areas, but report those samples with less taxonomic precision.

Using over 160,000 biological specimen records from the Clarion-Clipperton Zone submitted by 12 mining contractors to DeepData, I test these predictions. The dependent variable is whether each specimen record includes a species-level identification. I estimate linear probability models that incrementally add controls: a baseline comparison of zone types, then contractor fixed effects (absorbing differences in scientific capacity across firms), year fixed effects (absorbing improvements in identification methods over time), and environmental controls matched to each sample's location from Bio-ORACLE rasters. Across specifications, samples from non-contract areas are significantly more likely to include a species identification than prospectively-mined tract samples, and the result survives increasingly demanding controls. As a robustness exercise, I progressively condition on taxonomy fixed effects from phylum through genus; the gap persists even when comparing organisms classified to the same genus, ruling out compositional differences in sampled taxa as an explanation. I then exploit a structural feature of the dataset—a small but consistent share of biological samples are collected by one contractor inside *another* contractor's polygon—to run two complementary robustness exercises: a within-sampler comparison of own-tract vs. guest-tract reporting (holding scientific capacity fixed) and a within-polygon comparison of owner-collected vs. guest-collected samples (holding ecology fixed). Both designs deliver coefficients consistent with the headline finding.

A counterfactual exercise simulates species reporting as if all observations were reported with non-contract-area-level precision, holding exploration intensity and all other covariates fixed. This exercise allows me to decompose the missing species-level identifications in the dataset into two components: a portion attributable to strategic incentives (the deflation wedge) and a portion attributable to the genuine limits of deep-sea taxonomy. Even by the most conservative estimate, strategic incentives account for roughly seven percent of the missing species identifications—on the order of 8,000 observations under the most conservative bound and as many as 37,000 under the largest, all of which would carry a species name absent the permit-denial mechanism. I discuss alternative explanations for the pattern, including the inherent difficulty of deep-sea identification, differences in survey objectives across zones, reporting timing, and data quality, and argue that none can fully account for the systematic within-contractor gap. I conclude by discussing policy options, distinguishing between instruments that address the precision wedge (decoupling species discovery from permit denial through safe harbor provisions, sealed reporting, or independent taxonomic review) and those

that address the exploration wedge (raising the scientific return to precise identification through data prizes or genetic resource markets). Section 2 presents the institutional background, Section 3 presents the model, Section 4 describes the empirical methods, Section 5 presents results, Section 6 discusses these results, Section 7 discusses policy responses, and Section 8 concludes.

This paper connects to several strands of the economics literature. The first is the literature on strategic manipulation of self-reported environmental data. A well-documented finding is that firms' self-reported environmental data frequently diverge from independently measured values (de Marchi and Hamilton 2006), and that third-party auditors can be compromised when their compensation depends on the regulated firm (Duflo et al. 2013). Regulated agents also respond dynamically to enforcement, adapting their evasion strategies as they learn the inspection regime (Gonzalez-Lira and Mobarak 2021), and regulatory pressure can paradoxically reduce the quantity and quality of voluntary disclosures (Zakota 2025). In some cases, agents manipulate the monitoring process itself rather than the data it produces.[1] On the other hand, mandatory disclosure requirements can meaningfully improve environmental outcomes, though firms respond heterogeneously (Bennear and Olmstead 2008; Doshi et al. 2013), and firms may selectively disclose favorable information while withholding unfavorable results (Lyon and Maxwell 2011). This paper contributes to this literature by identifying a form of misreporting that operates not through falsification of observed quantities but through underproduction of information itself: contractors do not fabricate coarse identifications; they rationally forgo the precise ones. The second strand concerns the perverse incentives created by species-protection regulation. Brown and Shogren (1998) and Langpap et al. (2018) survey the economics of the Endangered Species Act (ESA), emphasizing the tension between conservation benefits that accrue nationally and costs that fall on private actors harboring species. This tension generates strategic responses: private landowners near known endangered species colonies have been shown to preemptively destroy potential habitat to avoid triggering ESA protections (Lueck and Michael 2003), while ESA policy gives landowners no incentive to cooperate with species information collection on their land (Polasky and Doremus 1998). Crucially, institutional design can mitigate these perverse incentives: safe-harbor agreements have been shown to reduce preemptive habitat destruction (Byl 2019). The mechanism in this paper is structurally similar: an agent with a valuable asset manipulates the conditions that would trigger conservation oversight. The key difference is that landowners destroyed potential habitat (a physical act) while mining contractors suppress taxonomic precision (an informational one). Finally, the paper contributes to the literature on public goods provision under asymmetric incentives (Harrington 1988; Heyes and Rickman 1999): the ISA's

[1] Mu et al. (2021) show that local governments strategically shut down air pollution monitors on days when air quality is expected to deteriorate. Kim and Lyon (2011) find that firms participating in the DOE's Voluntary Greenhouse Gas Registry reported emissions reductions even as their actual emissions increased.

exploration requirement is a clever attempt to use private rents as a carrot for public good provision, and the taxonomic deflation documented here is precisely the kind of shirking on uncontracted margins that such arrangements invite.

2 Background

2.1 The Deep Sea as a Commons

The resources of the ocean seabed beyond national jurisdiction are considered "the common heritage of mankind" under the United Nations Convention on the Law of the Sea (UNCLOS). These resources include critical minerals, such as cobalt and lithium, but also biodiversity. While deep-sea minerals have immediate economic value and could prove useful for decarbonization, there is significant concern that deep-seabed mining could destroy valuable biodiversity (Niner et al. 2018; Levin et al. 2020). These concerns are not speculative: Jones et al. (2025) revisited experimental mining tracks in the CCZ 44 years after disturbance and found that large sessile fauna showed very little recovery, and the first industrial-scale mining trial found a 37% reduction in macrofaunal abundance and a 32% reduction in species richness in directly impacted areas (Stewart et al. 2026). At present, less than 0.001% of the deep seabed has been visually observed (Bell et al. 2025). Rabone, Wiethase, et al. (2023) estimate that there are likely over 7,000 species in the Clarion-Clipperton Zone (CCZ), the most promising area for deep-seabed mining, the vast majority of which are currently unknown to science. Bonifácio et al. (2021) find significant morphological and genetic diversity within Polychaeta, the most prevalent class of deep-sea worms, and recent research cataloging Amphipoda, an order of small crustaceans, in the CCZ found that a majority of species discovered had been previously unknown (Jażdżewska et al. 2025). In a setting with such immense scientific uncertainty, exploration is critical for establishing the consequences of exploitation through mining.

2.2 The ISA's Exploration Requirement

The ISA is responsible for managing the deep seabed and has given out exploration contracts that entitle certain country-backed contractors to explore a specific tract (Blanchard et al. 2023). These are officially termed "exploration areas" by the ISA, but I will refer to them as prospectively-mined tracts throughout to contrast them with conserved areas. While the eventual goal of an exploration contract is to convert to an exploitation contract, these contracts currently require that the company intending to mine a certain area must collect environmental and biological data about it. The straightforward justification for this requirement is that it provides information necessary for the ISA to determine whether mining should in

fact be allowed in a particular location. More indirectly, this requirement forces prospective mining firms to provide a public good (non-rival, non-excludable information about the marine genetic resources of the deep sea) which is critical for decision makers like the ISA to balance the benefits of mining and the benefits of conservation (Jones et al. 2019).

One of the key deliberations over whether deep-sea mining should be permitted at all has been the role of undiscovered species and the precautionary principle (Jaeckel 2017). If firms are required to fund exploration prior to exploitation, this can prevent the possibility that mining destroys a valuable ecosystem or genetic resource before scientists can properly characterize whatever value it may have or deem it necessary for conservation. In fact, deep-sea organisms are increasingly prominent in marine bioprospecting, with growing numbers of patent filings tied to deep-sea genetic material (Zhivkoplias et al. 2024; Krusberg et al. 2024). If the deep seabed might have valuable biological resources, there is quasi-option value to delaying mining in favor of exploration (Arrow and Fisher 1974; Henry 1974); particularly given evidence that mining impacts persist for decades and recovery is at best partial (Jones et al. 2025). However, there may be insufficient incentives for exploring the deep sea, at least on the timeline required for eventual mining, because the entities that would most benefit from the discoveries may not have incentives to move quickly. Furthermore, even if the major benefit of undiscovered deep-sea species is the eventual discovery of pharmaceutical applications, pharmaceutical companies will still have insufficient incentive to fund bioprospecting today since they are unlikely to secure intellectual property rights over their discoveries. The ISA's exploration requirement is therefore a clever solution: it both incentivizes exploration and makes the body responsible for exploration (the contractor) internalize the cost of delays on mining. Alternative solutions, such as giving the same enterprise exclusive jurisdiction over both mineral and biological resources in an area, are legally infeasible and may still insufficiently incentivize bioprospecting because its benefits can be hard to monetize. On the other hand, pausing mining indefinitely until a scientific body deems the deep sea "sufficiently explored" is likely also legally infeasible and creates its own tension between exploration and conservation.

2.3 Taxonomic Deflation and the Comparison Strategy

However, this clever solution comes with one significant potential downside: it can incentivize the mining enterprise to explore poorly. Since the mining enterprise's returns come only from eventually mining the tract, it has strong disincentives against discovering anything that might lead to them not being allowed to mine later, such as rare or valuable species. This incentive could be addressed somewhat by mandating relatively high standards for the quality and thoroughness of reporting, the use of third-party explorers, or other procedures, but it could still operate through some channels. The company responsible for explo-

ration can try to shirk on uncontracted margins, such as hiring different teams of scientists, using different equipment, or altering categorization decisions in cases of uncertainty. In particular, the exploring body can attempt "taxonomic deflation," the flip side of "taxonomic inflation" where subpopulations of a species are erroneously categorized as independent species (Isaac et al. 2004). Taxonomic inflation has been used in service of conservationist goals, such as by making subpopulations eligible for endangered species protection (Ghezelayagh et al. 2025). So too can taxonomic deflation undermine conservation, by failing to identify distinct species that should be considered independently for protection. In both cases, the fundamental difficulties of defining a species give exploitable leeway to parties with their own goals and incentives (George and Mayden 2005). Nevertheless, the species is the fundamental unit of biodiversity, and so "deflation" to the phylum or class will necessarily understate a significant amount of biodiversity (Claridge et al. 1997).

All biological and environmental samples are supposed to be reported to the ISA in a standardized format and are ultimately made available to the public through the DeepData database. However, there is significant heterogeneity in the quality of data reported on DeepData. Rabone, Horton, et al. (2023) document several issues with the database, including duplication of entries, inconsistent or absent identifiers, and significant data quality issues for taxonomy variables. Even absent strategic underreporting, the biological data from the deep sea are often taxonomically coarse (Costello and Chaudhary 2017; Kennedy and Rotjan 2023), and the ISA has been significantly less transparent than other international resource managers (Ardron 2018). Therefore, it is not enough to point out the relative lack of taxonomic information in DeepData: the data from prospectively-mined tracts must be compared against a suitable comparison group.

The primary comparison in this paper is between contractor-allocated exploration tracts and all other areas in the CCZ: Areas of Particular Environmental Interest and areas outside any formal zone designation, which I refer to collectively as non-contract areas. All samples in the DeepData database, including those from APEIs and outside areas, are submitted by mining contractors as part of their exploration programs. This means the comparison is within the same set of reporting entities, holding fixed the inherent difficulties of deep-sea exploration, the reporting process through DeepData, and each contractor's scientific methodologies and institutional partnerships. If the same contractors' data show systematically less taxonomic detail in their contract areas than in non-contract areas, this is consistent with taxonomic deflation. The APEI comparison is conceptually the cleanest: these zones are explicitly designated to be free of mining and biodiversity reporting there carries no permit consequences. Regression results restricted to the APEI-only comparison are presented in the Appendix and are consistent with the main analysis. Figure 1 shows the administrative boundaries of the CCZ, with contractor tracts (colored) and APEIs (gray) labeled. The spatial distribution of biological samples in the DeepData database is shown in Appendix Figure 9.

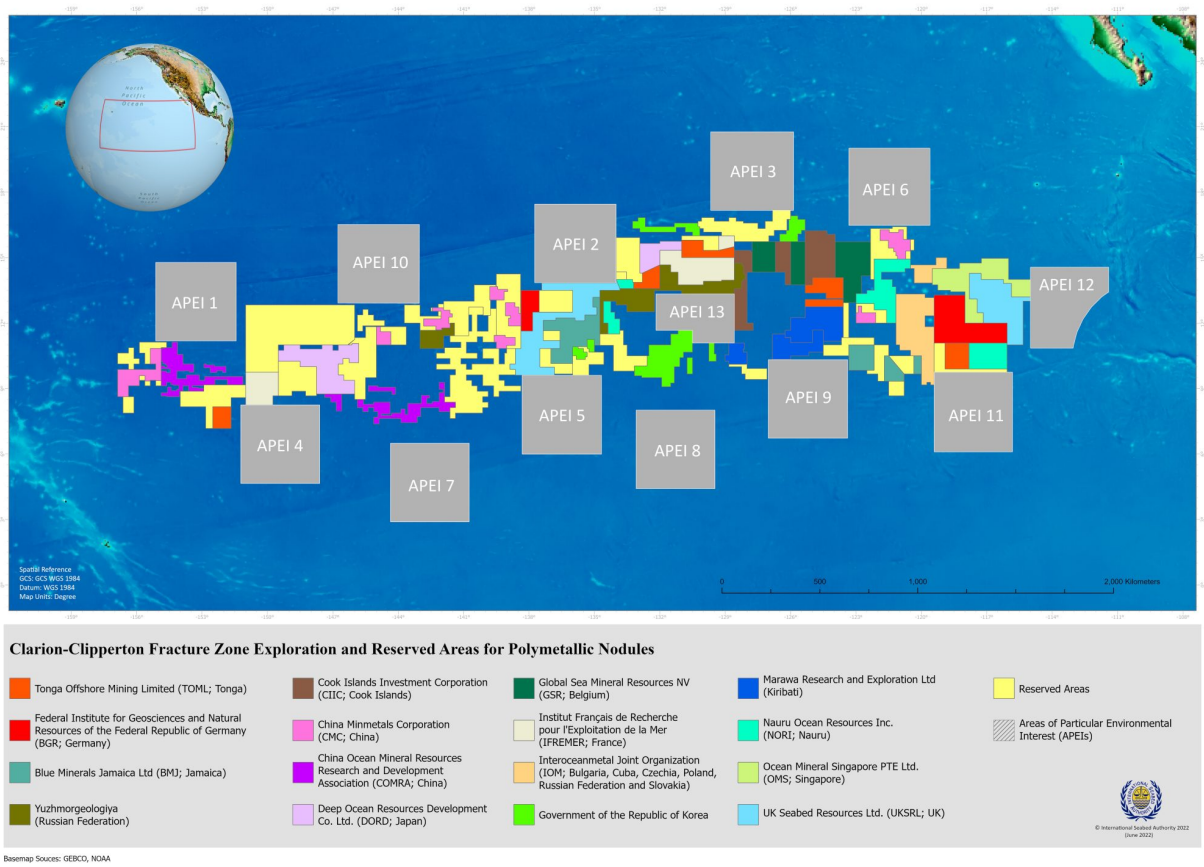

Figure 1: Contractor exploration areas (colored) and Areas of Particular Environmental Interest (gray) in the Clarion-Clipperton Zone. Source: International Seabed Authority.

3 Model

Before turning to the empirical analysis, I develop a theoretical model to fix ideas and clarify the mechanism underlying the empirical strategy. The model demonstrates that the observed pattern of taxonomic deflation follows naturally from rational contractor behavior under the ISA's regulatory structure. The model generates two predictions that correspond directly to the empirical findings: contractors conduct more physical sampling on prospectively-mined tracts than in non-contract areas, but report those samples with less taxonomic precision. This decoupling of effort and precision is the signature of strategic taxonomic deflation.

The key insight is that exploration intensity and taxonomic precision are distinct instruments, and the regulatory environment creates asymmetric incentives for each. More sampling unambiguously improves a contractor's chances of regulatory approval by demonstrating survey compliance. Taxonomic precision, by contrast, cuts both ways: some specificity signals scientific credibility, but excessive precision risks uncovering

species that trigger conservation review and jeopardize the mining permit. Outside contract areas, where no permit is at stake, this tension disappears and contractors have no reason to withhold precise identifications. The gap in taxonomic reporting between the two zones therefore reflects incentives, not capacity.

3.1 Setup

Consider a mining contractor that simultaneously chooses two instruments: exploration intensity $e \geq 0$, capturing the volume of biological sampling it conducts, and taxonomic precision $\tau \in [0, 1]$, capturing how specifically it identifies the organisms it recovers. These are conceptually and empirically distinct: a contractor can collect many samples while reporting all observations only at the phylum level, or conduct a small but meticulous survey. The DeepData database encodes exactly this distinction: it records both the number of biological samples submitted and the taxonomic level at which each is identified.

The regulator observes both instruments and makes an approval decision. Exploration intensity increases approval probability by demonstrating survey compliance with ISA requirements. Taxonomic precision has a non-monotone effect: the regulator rewards sufficient precision as evidence of scientific diligence, but high precision risks identifying species that trigger conservation review. I model the probability of permit approval as:

$$\pi(e, \tau) = [1 - \exp(-\alpha e)] \left[1 - \phi(\tau - \tau^*)^2\right] \tag{1}$$

The first term captures the compliance benefit of exploration: approval probability rises with exploration intensity at rate $\alpha > 0$, approaching 1. The second term captures the precision sweet spot: the regulator rewards precision up to $\tau^* \in (0, 1)$—enough specificity to be credible—but penalizes deviations in either direction at rate $\phi > 0$, where ϕ captures the regulatory bite of a species discovery. The quadratic penalty is symmetric around τ^* for tractability, but the economically important direction is the penalty for *exceeding* τ^*: a contractor that identifies too precisely risks uncovering species that trigger conservation review. The penalty for falling *below* τ^* captures the regulator's ability to reject submissions that are insufficiently detailed to be scientifically credible. The two instruments enter multiplicatively: the payoff from precision is increasing in sampling volume, since precise identifications only matter if there are samples to identify.

Costs are separable: $c(e, \tau) = \kappa e^2 + \lambda \tau^2$. The value of the mining permit is $V > 0$. The contractor's expected payoff on a prospectively-mined tract is:

$$\Pi_M(e,\tau) = [1-\exp(-\alpha e)]\left[1-\phi(\tau-\tau^*)^2\right]V - \kappa e^2 - \lambda\tau^2 \tag{2}$$

3.2 Optimal Choices on Prospectively-Mined Tracts

The contractor chooses exploration intensity e_M and taxonomic precision τ_M jointly, where the subscript M denotes the prospectively-mined tract regime (and A will denote the non-contract area regime below). The first-order conditions are:

$$\frac{\partial \Pi_M}{\partial e}: \quad \alpha\exp(-\alpha e_M)\left[1-\phi(\tau_M-\tau^*)^2\right]V = 2\kappa e_M \tag{3}$$

$$\frac{\partial \Pi_M}{\partial \tau}: \quad -2\phi(\tau_M-\tau^*)\left[1-\exp(-\alpha e_M)\right]V = 2\lambda\tau_M \tag{4}$$

The first condition equates the marginal compliance benefit of additional sampling with its marginal cost. The second pins down optimal precision. Let $A_M \equiv 1-\exp(-\alpha e_M) \in (0,1)$ denote the compliance value of the contractor's chosen exploration level. Solving the second condition explicitly:

$$\tau_M = \frac{\phi\tau^* A_M V}{\lambda + \phi A_M V} \tag{5}$$

The contractor always chooses $\tau_M < \tau^*$: it strategically undershoots the regulator's ideal precision level. The undershoot $\tau^* - \tau_M = \frac{\lambda\tau^*}{\lambda+\phi A_M V}$ is strictly positive for any finite V and ϕ. The undershoot is decreasing in the compliance value A_M, but only slowly: even contractors that have invested heavily in exploration still deflate their taxonomic reporting substantially, because the very magnitude of their investment raises the stakes of a species discovery.

3.3 Optimal Choices in Non-Contract Areas

In non-contract areas, there is no permit at stake, and the contractor faces no regulatory penalty for precise identification. The contractor surveys these areas to generate scientific value, which I model as proportional to the product of effort and precision: what matters is the number of precisely-identified observations, not mere sampling volume:

$$\Pi_A(e,\tau) = b\cdot e\cdot\tau - \kappa e^2 - \lambda\tau^2 \tag{6}$$

where $b > 0$ is the marginal scientific value of a precise observation. The first-order conditions are $b\tau_A = 2\kappa e_A$ and $be_A = 2\lambda\tau_A$, which together imply that exploration intensity and taxonomic precision are proportional and jointly determined by the scientific return b. The contractor has no reason to drive a wedge between sampling volume and taxonomic precision.

3.4 The Central Predictions

Comparing the two regimes yields the paper's twin predictions.

Prediction 1: More Exploration on Prospectively-Mined Tracts. When V is large relative to b, the contractor explores prospectively-mined tracts more intensively because the permit value V far exceeds the scientific return b that motivates surveys in non-contract areas. From the first-order condition on prospectively-mined tracts, exploration is driven by αV: the marginal gain from demonstrating compliance to the regulator. In a setting where mining rights represent large commercial values, $\alpha V \gg b$, so contractors deploy substantially more physical sampling on their prospectively-mined tracts.

Prediction 2: Less Precision on Prospectively-Mined Tracts. Despite greater exploration intensity, contractors choose strictly lower taxonomic precision on prospectively-mined tracts: $\tau_M < \tau^*$ whereas τ_A is unconstrained by regulatory risk. On prospectively-mined tracts, precision above τ^* threatens permit denial, creating a strategic motive to stop short of full identification. Outside contract areas, no such threat exists and the contractor identifies as precisely as the science warrants.

Predictions 1 and 2 are complementary. They imply that the raw correlation between sampling volume and taxonomic precision is negative across tract types: the zone with more samples has fewer species identifications per sample. This is the signature of strategic taxonomic deflation, and is difficult to reconcile with a resource-constraints story, under which effort and precision would naturally move together. Define the precision-per-unit-effort ratio $\rho \equiv \tau/e$. Then:

$$\rho_M = \tau_M/e_M \quad < \quad \tau_A/e_A = \rho_A \tag{7}$$

This ratio is directly observable in the DeepData database as the share of submitted samples reporting a species-level identification, and is precisely what the regressions below estimate.

3.5 Further Implications

The model generates additional comparative statics that, while not directly tested, sharpen intuition about how the deflation mechanism operates.

Asset value and deflation intensity. The model implies that $\partial e_M/\partial V > 0$ and $\partial \tau_M/\partial V < 0$: higher tract value simultaneously increases exploration intensity and decreases taxonomic precision. More commercially valuable deposits (those with greater nodule density or surveyed during periods of high commodity prices) should produce larger sample counts but coarser identifications within the same contractor. This is a potentially testable implication.

Regulatory bite. The undershoot $\tau^* - \tau_M$ is increasing in ϕ: as the regulatory consequences of species discovery become more severe, the contractor has greater reason to depress precision below the credible reporting threshold. Testing this directly would require variation in regulatory stringency that is not available in the current setting.

These behavioral implications have direct consequences for welfare. To see this, consider a social planner who maximizes total surplus, adding the social value $s > 0$ of a precisely-identified species (capturing option value for conservation, ecosystem services, and future research):

$$W(e, \tau) = \pi(e, \tau) \cdot V - c(e, \tau) + s \cdot e \cdot \tau \tag{8}$$

Comparing the social optimum to the contractor's private optimum on prospectively-mined tracts reveals two distinct wedges. The first is an *exploration wedge*: the contractor ignores the social value s of each precisely-identified sample, so there is underexploration relative to the social optimum even on prospectively-mined tracts where exploration is commercially motivated. This does not contradict Prediction 1—contractors explore prospectively-mined tracts more than non-contract areas because the permit value V is large, but they still explore *less than the social planner would* because they do not internalize the scientific externality. The second is a *precision wedge*: the contractor is deterred from full taxonomic precision by the permit risk ϕ, driving τ_M below the socially optimal level independent of how much sampling it conducts. The contractor could identify more precisely at no additional sampling cost, but rationally chooses not to.

The model's parameters map onto distinct classes of policy intervention. Reducing ϕ (through safe harbor provisions, sealed reporting, or independent taxonomic review) addresses the precision wedge by decoupling discovery from denial. Raising s (through data prizes, genetic resource markets, or publication incentives) addresses the exploration wedge by internalizing the scientific externality. Importantly, the two instru-

ments are complements: reducing ϕ amplifies the contractor's precision response to any increase in scientific incentives. I return to these policy directions in greater detail after the empirical analysis.

3.6 Mapping to the Empirical Design

The model maps directly onto the regression design below. The contract area indicator partitions observations into those where the permit-denial mechanism is operative ($\phi > 0$, prospectively-mined tracts) and those where it is not ($\phi = 0$, non-contract areas). Contractor fixed effects absorb firm-level variation in the scientific return b, cost parameters κ and λ, and exploration technology α, isolating the effect of the regulatory penalty on taxonomic precision. Year fixed effects absorb improvements in sampling methodology over time.

Under this mapping, the estimated coefficient $\hat{\beta}$ on the contract area indicator recovers the precision gap $\tau_M - \tau_A$: the difference in species reporting rates induced by the regulatory mechanism, holding contractor and year fixed. This gap can be expressed as $\hat{\beta} = \tau_M - \tau_A = \frac{\phi \tau^* A_M V}{\lambda + \phi A_M V} - \tau_A$, where τ_A is pinned down by b and the cost parameters. Intuitively, $|\hat{\beta}|$ is increasing in ϕ and V: the larger the stakes, the wider the gap. The counterfactual analysis—which simulates non-contract-area reporting precision for all observations—corresponds to shutting down the $\phi > 0$ mechanism while holding all else fixed. Crucially, the model also clarifies why sampling volume cannot explain the finding: Prediction 1 implies that prospectively-mined tracts have more samples, so any purely ecological or resource-constraints explanation would predict *better* species coverage there, not worse. The observed reversal is the signature of strategic taxonomic deflation.

4 Methods

I test for taxonomic deflation by examining the taxonomic resolution of data reported on DeepData. Each observation in the dataset is a biological specimen record submitted by a contractor to the ISA. When reporting biological specimens to DeepData, some records include as many as eight taxonomic levels: Kingdom, Phylum, Class, Order, Family, Subfamily, Genus, and Species. Far more commonly, records only report as far as the phylum or class. I focus on whether each record includes a species-level identification, as the species is the fundamental unit of biodiversity assessment (Claridge et al. 1997). The dependent variable is therefore binary: it equals one if the specimen record includes a species name, and zero otherwise. I compare species reporting in samples from prospectively-mined tracts to those from non-contract areas, accounting for various possible confounding variables.

The dataset contains over 160,000 biological specimen records from the CCZ, submitted by 12 mining

contractors to DeepData. These records span prospectively-mined tracts—allocated to specific contractors for mineral exploration—as well as APEIs and areas outside both zone types. All records in the dataset, including those from non-contract areas, were submitted by mining contractors; the dataset does not include samples from independent academic expeditions or other non-contractor sources. This is central to the empirical strategy: the comparison is between samples from different zone types reported by the same set of contractors, not between contractor data and independently-collected scientific data. I match each record to environmental data from Bio-ORACLE rasters (Assis et al. 2024) based on its reported longitude and latitude, yielding an analysis dataset of 160,967 observations with complete environmental covariates. Table 3 in the Appendix reports summary statistics by zone type, including the number of records, year range, and species reporting rates.

I begin by showing the likelihood of species reporting in the raw data, then proceed to regression models that estimate the probability that a specimen record includes a species identification. The main explanatory variable is whether the sample comes from a contractor-allocated exploration tract (a contract area). I present four specifications that incrementally add controls, showing how the estimated effect changes (and survives) as potential confounds are absorbed. I estimate a linear probability model (LPM), which is appropriate in this instance because the objective is to identify the average marginal effect of contract area status on a binary outcome, and the linear model provides the most straightforward coefficient interpretation (Angrist and Pischke 2008). All standard errors are robust to heteroskedasticity. The first specification regresses the species-reporting indicator on an indicator for whether a sample comes from a contract area:

$$\text{Species Reported}_{ict} = \text{Intercept} + \beta \cdot \text{ContractArea}_i + \varepsilon_{ict} \tag{9}$$

The coefficient of interest, β, captures the difference in the mean probability of species-level identification between contract area samples and non-contract area samples (APEIs and outside areas) without accounting for possible confounding factors. Each observation i is a biological specimen record from contractor c in year t, and ε_{ict} is the error term. A negative β indicates that specimens from contract areas are less likely to include a species-level identification than specimens from non-contract areas.

The base specification omits potential confounding variables, so the estimated β could reflect factors other than strategic behavior. First, differences in taxonomic resolution could be driven by differences in scientific methodology rather than strategic incentives. To account for this, I add fixed effects for the mining contractor and for the year. Contractor fixed effects absorb average differences in reporting rates across the 12 contractors in the data. Each contractor is a country-backed entity with an ISA exploration contract for

one or more tracts (e.g., Nauru Ocean Resources Incorporated). This is important because contractors that generally report more taxonomic detail might also be the ones submitting more non-contract area samples, which could generate a spurious relationship in the raw data. Year fixed effects absorb average differences in reporting rates by year, accounting for improvements in scientific methods and taxonomic knowledge over time. These regression equations are the following:

$$\text{Species Reported}_{ict} = \beta \cdot \text{ContractArea}_i + \text{Contractor FE}_c + \varepsilon_{ict} \tag{10}$$

$$\text{Species Reported}_{ict} = \beta \cdot \text{ContractArea}_i + \text{Contractor FE}_c + \text{Year FE}_t + \varepsilon_{ict} \tag{11}$$

A second possible source of confounding comes from ecological differences between non-contract areas and prospectively-mined tracts which could affect species identification rates. The threat would be most acute if APEIs were placed where mining is unattractive *because* the biology there is unusual—producing a comparison group that is systematically not like the contract tracts. The ISA's siting principles are designed to do the opposite. APEIs are large rectangular blocks (most are 160{,}000~km^2) sited under an explicit "representativity" criterion: they are intended to span the major environmental gradients of the CCZ so that whatever fauna are eventually documented there resemble what would be displaced by mining elsewhere (Wedding et al. 2013; Smith et al. 2021). They were carved out of the CCZ grid rather than chosen to surround unusual habitats. APEIs are, by design, the closest analogue to contract tracts that the regulatory architecture allows. The available evidence is consistent with this: Washburn et al. (2021) find that environmental conditions inside APEIs largely overlap those inside contract areas, with some heterogeneity at the margins. To absorb that residual heterogeneity, I add four environmental variables that capture the major axes of variation across the CCZ: depth (which influences pressure, habitat type, and benthic community composition), sea surface temperature (capturing the east-west and north-south productivity gradients), surface chlorophyll concentration (a direct measure of primary productivity and thus the organic carbon flux that sustains abyssal communities), and dissolved bottom oxygen (which constrains benthic metabolic capacity). All environmental data are from Bio-ORACLE rasters matched to individual sample locations. Variance inflation factors for the environmental controls are all below conventional thresholds for concern. Nodule abundance, another potentially relevant variable, is not included because spatially resolved nodule data are not publicly available at the sample level; however, the contractor fixed effects partially absorb nodule-related variation, since prospectively-mined tracts were allocated in part based on nodule density. The fourth regression equation is the following:

$$\text{Species Reported}_{ict} = \beta \cdot \text{ContractArea}_i + \text{Environmental Controls}_i + \text{Contractor FE}_c + \text{Year FE}_t + \varepsilon_{ict} \quad (12)$$

where environmental controls include sea surface temperature, surface chlorophyll, dissolved bottom oxygen, and depth matched to each sample's location.

A remaining concern is that differences in species reporting could reflect differences in the *composition* of organisms sampled rather than strategic behavior. If contract areas happen to contain taxa that are inherently harder to identify to the species level, the gap could arise even without deflation. Importantly, taxonomy is itself endogenous to the exploration process: contractors who deflate will report coarser taxonomic classifications, so conditioning on taxonomy absorbs variation that may itself be generated by the deflation mechanism. Nonetheless, as a robustness exercise, I present results from specifications that each condition on a single taxonomy fixed effect (phylum, class, order, family, or genus, omitting subfamily which is sparsely reported) on top of the full set of exogenous controls. Each specification restricts the sample to observations with a valid identification at the relevant taxonomic level. This is a conservative test: if the gap survives even after conditioning on endogenous taxonomy at the finest level (genus), it cannot be attributed to compositional differences in the organisms sampled at any taxonomic level.

Two further robustness exercises exploit a structural feature of the dataset uncovered by overlaying official ISA contract polygons on each observation's coordinates: a small but consistent share of biological samples are collected by one contractor inside *another* contractor's polygon, typically during joint expeditions or transit segments. These "guest" samples are processed by the same scientific institutions as own-tract samples, but carry no permit-related incentive for the sampler. This enables (i) a within-sampler design comparing each contractor's own-tract reporting to its guest-tract reporting (holding scientific capacity fixed) and (ii) a within-polygon design comparing owner-collected to guest-collected samples inside the same physical tract (holding ecology fixed). Together these designs sandwich the headline effect from above and below using identifying variation that the contractor-and-environmental controls in the main regressions cannot use.

Across all specifications, if contract area samples are reported with lower taxonomic resolution after accounting for confounding variables, then β should be negative and statistically significant.

5 Results

5.1 Exploration Intensity across Zones

The model's first prediction, that contractors conduct more physical sampling on prospectively-mined tracts than on non-contract areas, can be evaluated directly from the DeepData records before any regression analysis. Figure 2 plots the total number of biological specimens submitted from prospectively-mined tracts versus non-contract areas.

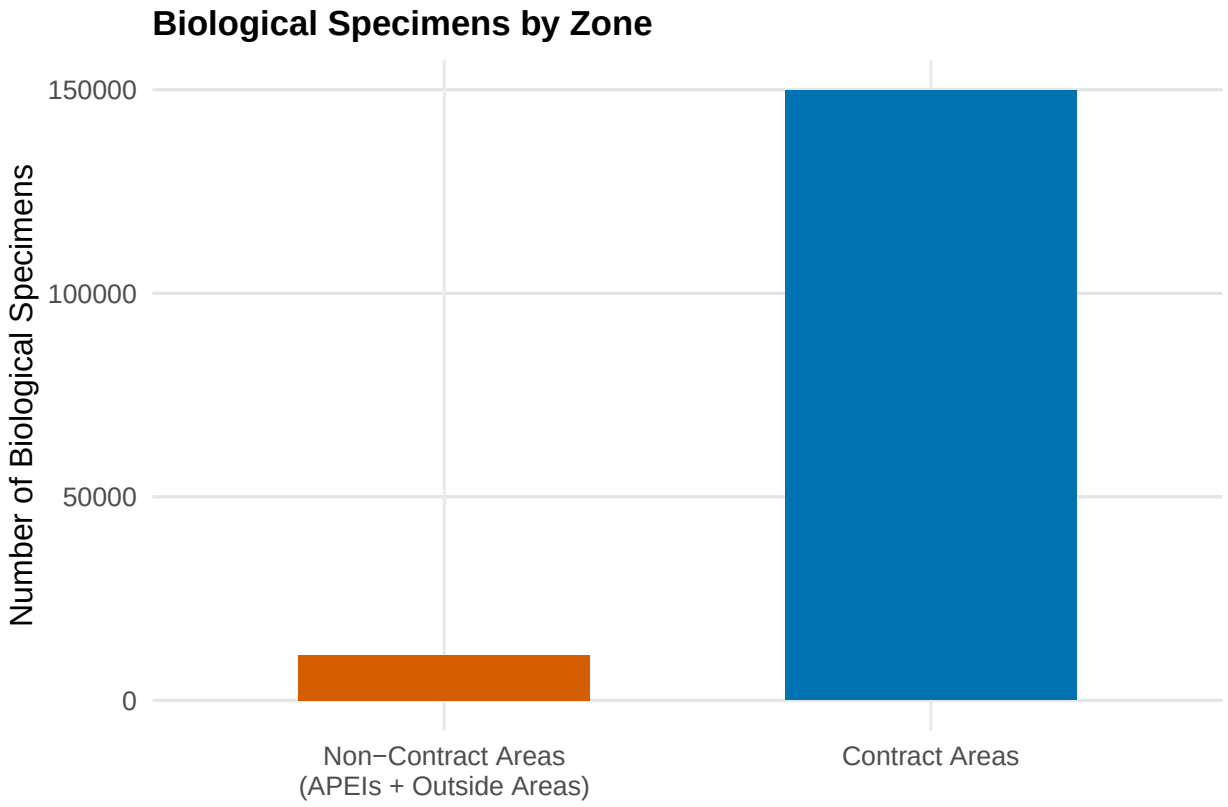

Figure 2: Consistent with Prediction 1, contractors conduct far more physical sampling on tracts they intend to mine.

Consistent with Prediction 1, the overwhelming majority of biological samples in DeepData originate from contract areas rather than non-contract areas. This is unsurprising given that the commercial value of a mining permit dwarfs the scientific return from surveys outside the contract area ($\alpha V \gg b$ in the model's terms). The critical question is whether this greater exploration effort translates into correspondingly greater taxonomic detail. Under a pure resource-constraints explanation, the zone with more samples should also

exhibit better species coverage. The model's second prediction is that the opposite holds.

5.2 Taxonomic Reporting by Zone

The raw data confirm Prediction 2. Figure 3 shows that most taxonomic levels are more likely to be reported in non-contract areas. For example, 44.3% of observations from non-contract areas report the species of the sample, compared to just 16.4% of observations from contract areas. While this raw difference could reflect many things, the pattern persists after adding controls.

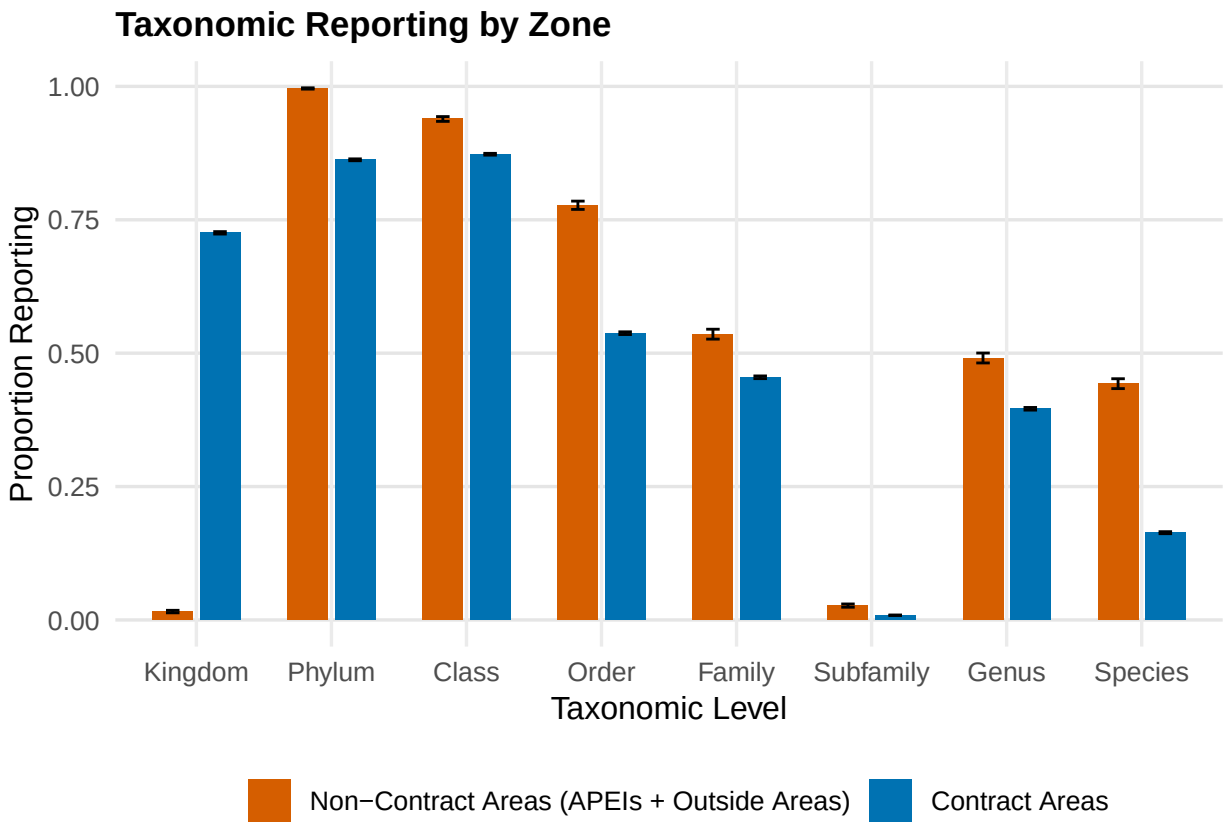

Figure 3: Proportion of specimen records with a non-missing value at each taxonomic level. Each level is evaluated independently from the others: a record counts as reported at a given level if and only if its corresponding taxonomic field (Kingdom, Phylum, ..., Species) contains a non-missing value, regardless of whether coarser or finer levels are also populated. The gap widens at finer taxonomic levels, consistent with strategic deflation at the margins where species discovery poses the greatest regulatory risk. 95% confidence intervals shown.

5.3 Regression Results

Regression results provide a formal test of Prediction 2. Figure 4 plots the estimated $\hat{\beta}$ (the precision gap $\tau_M - \tau_A$ from the model) and 95% confidence intervals for four specifications that incrementally add exogenous controls. The coefficient on contract area status is negative and statistically significant in four of the four designs, implying anywhere from 28 to 6 percentage points lower probability of species reporting in contract areas relative to non-contract areas. As the model predicts, contractor fixed effects, which absorb variation in scientific capacity across firms, reduce but do not eliminate the gap (Appendix Figures 10 and 11 summarize sampling volume and species reporting by contractor). Adding year fixed effects reduces the coefficient further ($p = 0.000$) but it remains negative and significant at conventional thresholds; adding environmental controls sharpens precision further. Full regression results are reported in Table 4 in the Appendix.

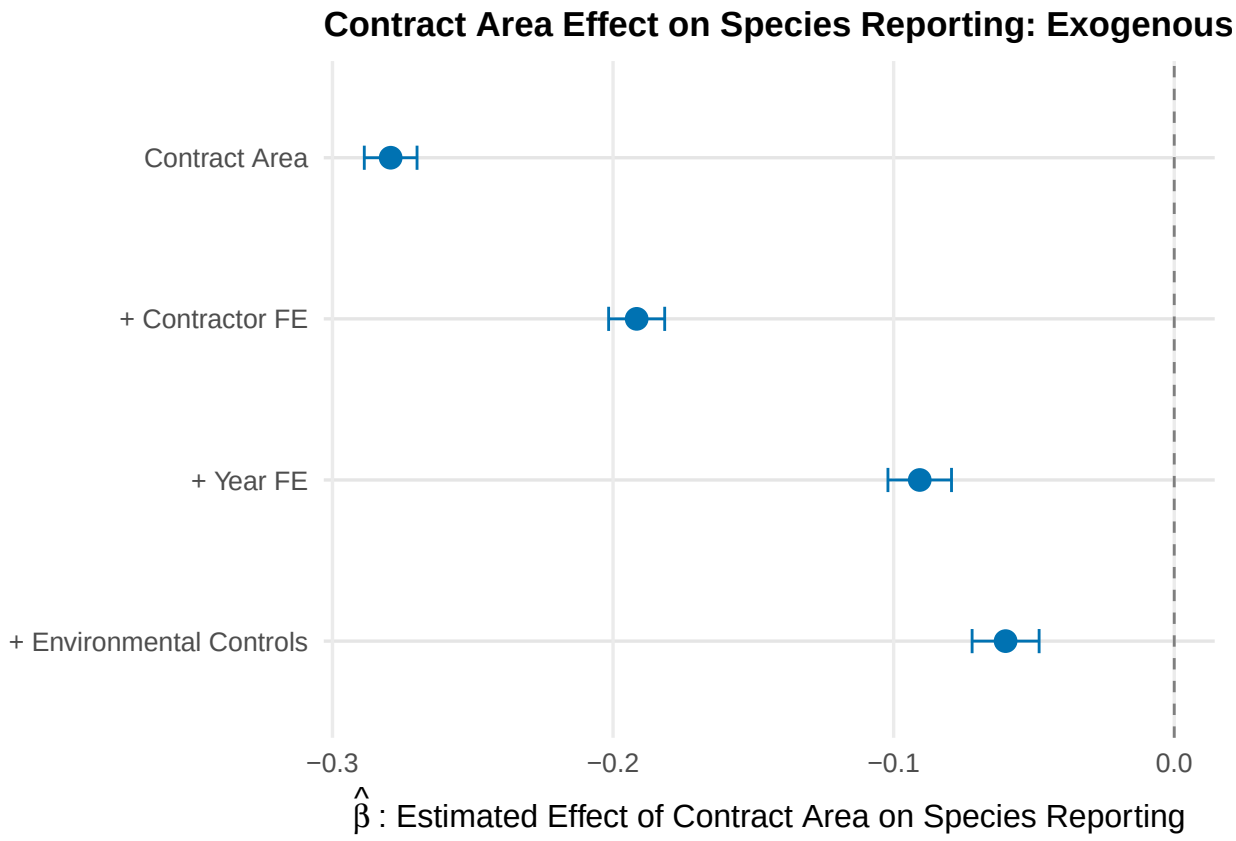

Figure 4: Estimated contract area coefficient ($\hat{\beta}$) from four LPM specifications, each adding incrementally more confounding variable controls. Each point represents the estimated difference in the probability of species-level identification for contract area samples relative to non-contract area samples (APEIs and outside areas). Negative coefficients indicate that species are less likely to be reported from a contract area than from a non-contract area. 95% confidence intervals computed with robust standard errors.

A natural concern is that the gap reflects compositional differences in the organisms sampled across zones rather than strategic behavior. If contract areas happen to contain taxa that are inherently harder to identify, the reporting gap could arise without deflation. To address this, Figure 5 presents a taxonomy robustness exercise that separately conditions on each taxonomic level (from phylum through genus) on top of the full set of exogenous controls (contractor, year, and environmental variables). Importantly, taxonomy is itself endogenous to the exploration process: contractors who deflate will report coarser taxonomic classifications, so conditioning on taxonomy is a conservative test that absorbs variation that may itself be generated by the deflation mechanism. Despite this, the contract area coefficient remains negative and statistically significant even when comparing organisms within the same genus, ruling out compositional differences at

every taxonomic level as an explanation for the gap (Table 5).

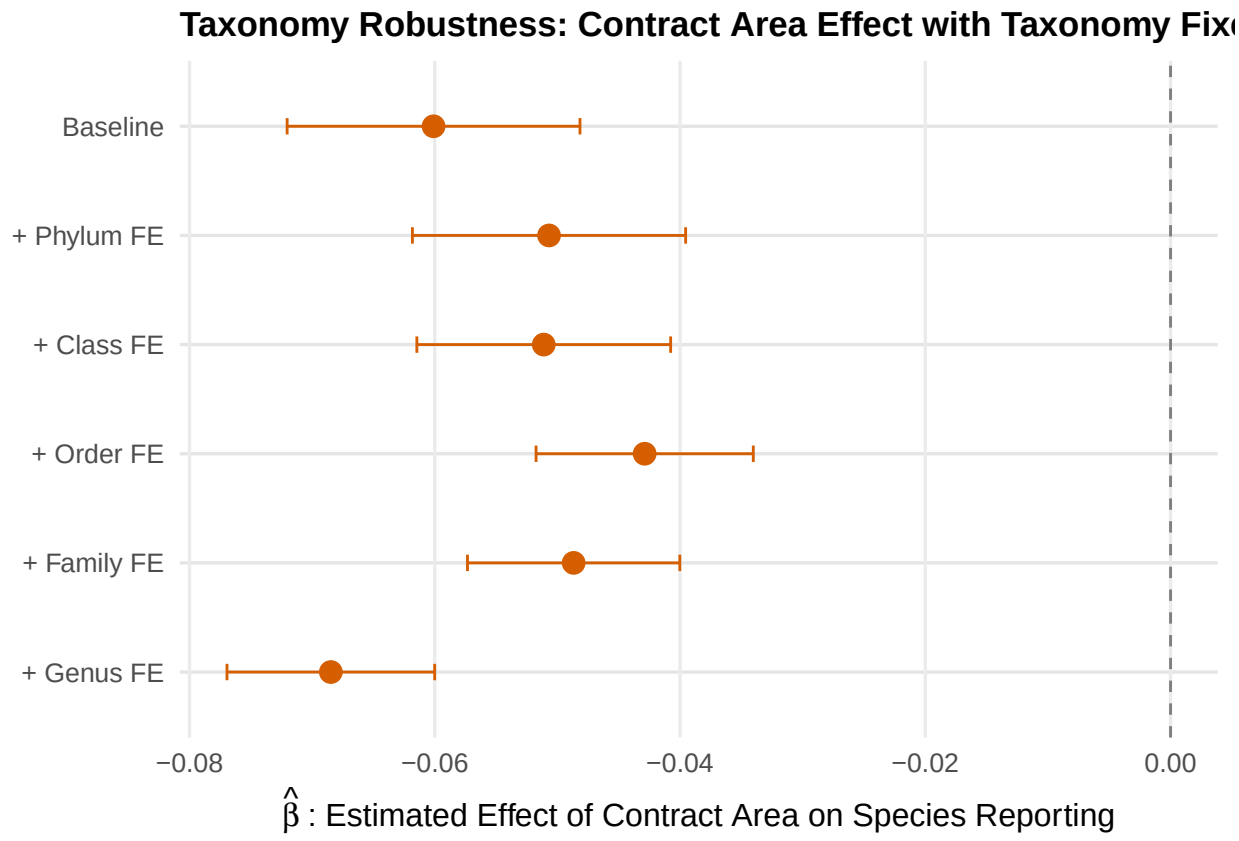

Figure 5: All specifications include contractor, year, and environmental controls. Each specification restricts the sample to observations with a valid identification at the relevant taxonomic level. 95% confidence intervals shown.

The persistence of $\hat{\beta} < 0$ after controlling for contractor identity, year, environmental conditions, and taxonomy indicates that the gap in species-level reporting between contract areas and non-contract areas is robust to a wide range of potential confounding variables. I discuss the interpretation of this pattern, including alternative explanations, in the following sections.

An important caveat is that species-level identification in the deep sea is genuinely difficult. An estimated 88–92% of metazoan species in the CCZ have not yet been formally described (Rabone, Wiethase, et al. 2023), and even expert teams can identify fewer than 20% of observed species (Kennedy et al. 2019). I do not claim that every instance of coarse reporting reflects strategic behavior. However, these difficulties cannot by themselves explain the pattern documented here: all samples come from the same region and face the same identification challenges, so inherent difficulty should depress species reporting symmetrically across

zones. The systematic gap between contract areas and non-contract areas, which persists after controlling for contractor identity, year, environmental conditions, and taxonomy (Figure 5), is what the inherent-difficulty explanation cannot account for. I return to this and other alternative explanations in the Discussion.

5.4 Within-Sampler and Within-Polygon Robustness

The contractor and taxonomy fixed effects above absorb cross-firm differences in scientific capacity and cross-taxon differences in identifiability, but they do not directly hold the *physical location* of each sample fixed. Geological and ecological conditions vary in space, and if contractor tracts are systematically located in regions where species are harder to identify for non-strategic reasons, the regression results above could be confounded. To address this, I exploit a structural feature of the dataset uncovered by overlaying the official ISA contract polygons onto each observation's coordinates: a small but consistent share of biological samples are collected by one contractor inside *another* contractor's polygon, typically during joint expeditions or transit segments. These "guest" samples are processed by the same scientific institution, in the same laboratory pipeline, and often during the same cruise as the sampler's own-tract samples—but they carry no permit-related incentive for the sampling contractor, since the sampler does not hold mining rights at the location where the specimen was collected. Appendix Table 6 documents the full source decomposition.

This cross-contractor sampling enables two distinct robustness designs, each holding a different confounder fixed:

1. *Same contractor, different areas*: a within-sampler comparison that asks whether a given contractor reports samples differently when collecting in their own tract versus when collecting in another contractor's tract. Sampler fixed effects absorb cross-firm differences in scientific capacity, so the surviving variation isolates the regulatory regime change.

2. *Same area, different contractor*: a within-polygon comparison that asks whether, inside a single physical tract, samples collected by the polygon owner differ from samples collected by guest contractors. Polygon fixed effects absorb local ecology, so the surviving variation isolates the sampler-by-regime contrast at that location.

The two designs are complementary: the first holds reporting infrastructure fixed at the cost of leaving ecology to vary, while the second holds ecology fixed at the cost of leaving reporting infrastructure to vary. Both deliver estimates that support the headline finding, and I report them side-by-side in Table **??**. Both designs use the same sample: all observations inside some PMN polygon, across all 12 sampling contractors. Samplers without Owner_i variation contribute nothing to the within-sampler $\hat{\beta}$, and polygons

without Owner_i variation contribute nothing to the within-polygon $\hat{\beta}$, but those observations remain in the regression to inform the year and environmental coefficients. The estimating equations share a common structure:

$$\text{Species Reported}_{ipct} = \beta \cdot \text{Owner}_{ipct} + \mu_p + \gamma_c + \delta_t + X'_{ipct}\eta + \varepsilon_{ipct} \tag{13}$$

where $\text{Owner}_i = 1$ if observation i was collected by the polygon owner and 0 if by a guest. Different combinations of fixed effects identify β from different variation: with sampler fixed effects γ_c but not polygon fixed effects μ_p, β captures the within-sampler difference in species reporting between own-tract and guest-tract samples; with polygon fixed effects but not sampler fixed effects, β captures the within-polygon difference between owner-collected and guest-collected samples; including both is uninformative because polygon and ownership are nearly collinear (each polygon has exactly one owner).

Table 1: Within-sampler design (same contractor, different areas): species reporting on Owner indicator, holding scientific capacity fixed via sampling-contractor fixed effects. Sample: all observations inside some PMN polygon, across all 12 sampling contractors. Samplers without Owner_i variation drop from $\hat\beta$ identification. Robust (HC1) standard errors in parentheses.

	Owner	+ Sampler FE	+ Year FE	+ Env. Controls
Intercept	0.467***	0.284***	1.495***	−47.519***
	(0.005)	(0.006)	(0.009)	(2.371)
Owner sampler	−0.292***	−0.231***	−0.104***	−0.104***
	(0.005)	(0.006)	(0.006)	(0.006)
Depth (m)				0.000***
				(0.000)
Sampler FE	No	Yes	Yes	Yes
Year FE	No	No	Yes	Yes
Environmental Controls	No	No	No	Yes
Sea Surface Temp.				0.628***
				(0.030)
Surface Chlorophyll				−4.758***
				(0.512)
Dissolved Oxygen				6.855***
				(0.341)
Num.Obs.	149 300	149 300	149 300	149 300
R2	0.031	0.110	0.211	0.217

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

Within-sampler estimates are reported in Table 1. Column (1) reports the raw owner-vs-guest gap, $\hat{\beta} = -0.292$: owner-collected specimens are roughly 29 percentage points less likely to carry a species-level identification than guest-collected specimens. Adding sampler fixed effects shrinks the estimate to −0.231, year fixed effects further to −0.104, and environmental controls leave it at −0.104 (SE 0.006). This is somewhat

Table 2: Within-polygon design (same area, different contractor): species reporting on Owner indicator, holding ecology fixed via polygon fixed effects. Sample: all observations inside some PMN polygon. Polygons without Owner_i variation drop from $\hat\beta$ identification. Robust (HC1) standard errors in parentheses.

	Owner	+ Polygon FE	+ Year FE	+ Env. Controls
Intercept	0.467***	0.262***	1.695***	−102.612***
	(0.005)	(0.008)	(0.012)	(1.943)
Owner sampler	−0.292***	−0.203***	−0.400***	−0.398***
	(0.005)	(0.008)	(0.007)	(0.007)
Depth (m)				0.000***
				(0.000)
Polygon FE	No	Yes	Yes	Yes
Year FE	No	No	Yes	Yes
Environmental Controls	No	No	No	Yes
Sea Surface Temp.				0.949***
				(0.025)
Surface Chlorophyll				−2.836***
				(0.498)
Dissolved Oxygen				16.692***
				(0.321)
Num.Obs.	149 300	149 300	149 300	149 300
R2	0.031	0.094	0.209	0.219

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

larger in magnitude than the main paper's preferred specification ($\hat{\beta} = -0.060$ in Figure 4 and column (4) of Table 4), but recovered from a substantially more demanding identification: only within-sampler, within-year variation between own-tract and guest-tract samples is used.

Within-polygon estimates are reported in Table 2. Holding the physical tract fixed by polygon fixed effects alone yields $\hat{\beta} = -0.203$, roughly twice the within-sampler estimate; adding year fixed effects roughly doubles this further to −0.400, and environmental controls leave it at −0.398 (SE 0.007). The fully-controlled within-polygon estimate is roughly four times the within-sampler estimate. The expansion when adding year fixed effects reflects time variation in cross-sampler reporting practices: guest-tract samples are concentrated in particular years and are collected by contractors with above-average reporting rates, so within-polygon, within-year contrasts magnify the gap between owner and guest. This larger magnitude also reflects the fact that the within-polygon comparison picks up cross-sampler differences in scientific capacity at the same location alongside the regulatory regime change, so the fully-controlled within-polygon estimate should be read as an upper bound on the deflation effect under the assumption that residual within-polygon contractor-level heterogeneity reflects the regulatory regime rather than independent capacity differences. The two designs are best read as separate, mutually-reinforcing exercises rather than as a single specification ladder—combining polygon and sampler fixed effects is uninformative because polygon and ownership are nearly

collinear, since each polygon has exactly one owner. The headline finding survives whichever confounder one is most worried about: holding scientific capacity fixed gives a -0.10 effect, holding ecology fixed gives a -0.40 effect. The within-sampler design plausibly understates the true deflation effect (guest-collected samples may receive less intensive post-processing than own-tract samples regardless of regulatory regime), while the within-polygon design plausibly overstates it (cross-firm differences in scientific capacity contribute alongside the regulatory mechanism).

The identifying variation behind the within-sampler estimate is visible in Figure 6, which plots, for each of the five contractors who appear as both polygon owners and as guests, the species-reporting rate of their own-tract samples against the species-reporting rate of their guest-tract samples. The pattern is consistent across firms with substantial guest-tract sample sizes: BGR reports 6% of its own-tract samples to species but 27% of its guest-tract samples (n=2,378); UKSRL reports 30% versus 54% (n=6,553). The other three contractors with both roles (GSR, IFREMER, Korea) have very thin guest-tract samples (n=21 to 52) and should not be over-interpreted; Korea's apparent reversal is identified off 52 guest observations.

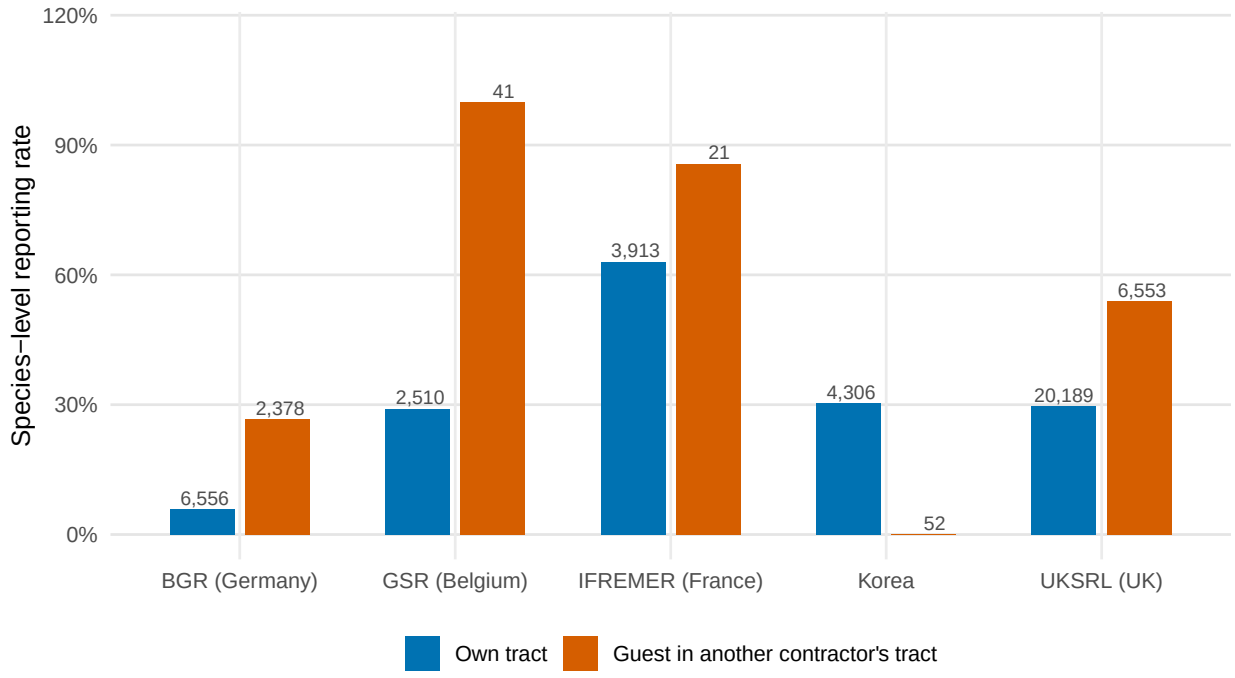

Figure 6: Species-level reporting rate, by sampling contractor and own-vs-guest status. Restricted to the five contractors who appear as both polygon owners and as guests in another contractor's polygon. Numbers above each bar are the underlying sample sizes; bars with very small n (Korea, GSR-guest, IFREMER-guest) should be interpreted with caution. The within-sampler estimate in Table reftab:reg-table-within-sampler averages these contrasts.

5.5 Counterfactual: Reporting without Taxonomic Deflation

The model's welfare analysis implies that taxonomic deflation creates a precision wedge: contractors could identify organisms more precisely at no additional sampling cost but rationally choose not to. To quantify the magnitude of this wedge, I conduct a counterfactual exercise corresponding to shutting down the $\phi > 0$ mechanism while holding exploration effort and all other covariates fixed. Figure 7 plots the predicted number of observations that would include a species entry in the absence of taxonomic deflation, out of the 150,104 observations with complete environmental covariates (the sample on which all four specifications can generate predictions). I generate fitted values from the regression results setting the contract area indicator to zero for every observation: that is, simulating the probability of species reporting as if every observation were reported with the same taxonomic precision as non-contract area samples. I use the largest and smallest estimated $\hat{\beta}$ in absolute value to generate bounds. As Figure 7 shows, this would lead to significantly higher predicted reporting: between 37,372 and 66,467 observations would report the Species of the sampled organism, compared to just 29,050 in actuality.

This counterfactual also decomposes the paper's motivating question. Of the 121,054 observations in the sample that lack a species identification, the share attributable to strategic incentives ranges from 7% (8,322 observations) under the most conservative coefficient to 31% (37,417 observations) under the largest. The complementary shares—93% and 69% respectively—would still lack a species identification even absent any deflation, representing the genuine limits of deep-sea taxonomy: undescribed fauna, specimen damage, and scarce expertise that no policy can immediately overcome. So even by the most conservative estimate, strategic incentives account for thousands of missing species identifications in DeepData; under the largest estimate, tens of thousands. This is a partial-equilibrium exercise: it holds exploration intensity and all covariates fixed, varying only the reporting precision. The model suggests that precision and effort are jointly determined, so eliminating the deflation incentive could also affect sampling volume; the counterfactual should therefore be interpreted as a lower bound on the full welfare gain from removing the incentive distortion.

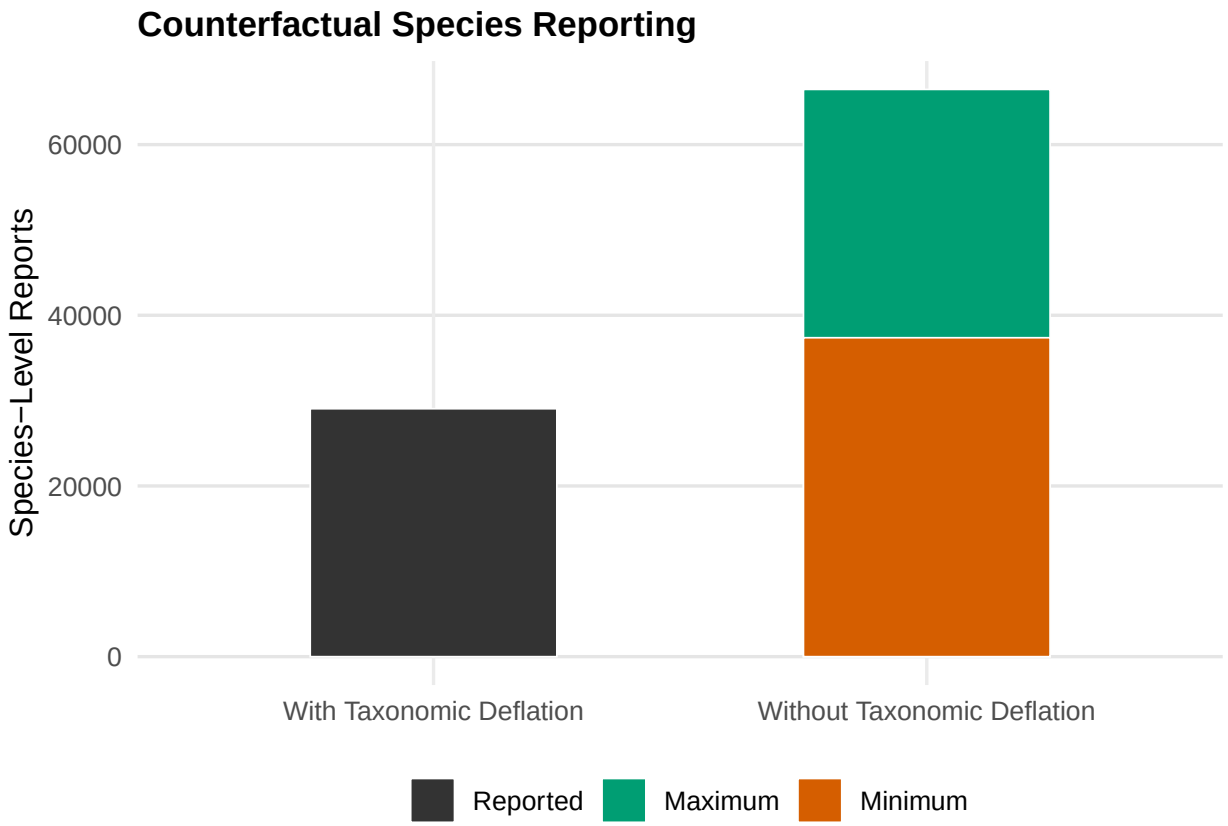

Figure 7: The dark bar plots observed species reports. The vermillion bar plots the counterfactual using the most conservative estimated contract area effect; the green bar extends above it to show the counterfactual using the largest estimated effect.

6 Discussion

The pattern I document fits the principal-agent problem the model describes: the ISA delegates biodiversity assessment to firms whose returns depend on subsequently being allowed to mine, and effort along the precision dimension is hard to monitor. Two features of the empirical evidence are central to the strategic interpretation. First, the gap survives the most demanding controls available—contractor identity, year, environmental conditions, taxonomy down to genus, and (in the within-sampler exercise) institution-level scientific capacity. Second, both spatial robustness designs point in the same direction: when the *same* contractor samples in another firm's tract, it reports more precisely than in its own tract; when the *same* tract receives samples from both its owner and a guest contractor, the owner's samples are coarser. No non-strategic mechanism predicts both patterns. My empirical design documents an association rather than

directly observing intent, but the alternatives I consider below cannot account for the joint within-sampler-and-within-polygon evidence.

6.1 Alternative Explanations

The most prominent objection is that the reporting gap reflects the inherent difficulty of deep-sea taxonomy rather than strategic behavior. Deep-sea identification is genuinely hard: less than 0.001% of the seafloor has been visually surveyed (Bell et al. 2025), expertise is scarce (Frutos et al. 2022), and an estimated 88–92% of CCZ species remain formally undescribed (Rabone, Wiethase, et al. 2023). These constraints almost certainly explain why species-level reporting is low in *absolute* terms across the dataset. They cannot explain why reporting rates differ *systematically* between zones, because every sample in the analysis was collected from the abyssal CCZ and faces the same taxonomic knowledge gaps. If inherent difficulty were the sole driver, it would depress species reporting symmetrically across contract areas and non-contract areas. The gap between zones is precisely the component that science alone cannot explain.

A second alternative is that the gap reflects differences in *survey objectives* between zones. Non-contract area surveys may be designed with biodiversity documentation as a primary goal, receiving greater post-cruise taxonomic follow-up, while surveys of contract areas may prioritize environmental baseline characterization. If non-contract area samples are more likely to be sent to specialized taxonomic laboratories or to receive extended analysis by academic collaborators, species-level identification rates could be higher even without any difference in field effort. This concern is partially addressed by contractor fixed effects, which absorb average differences in each contractor's scientific methodology and institutional partnerships. However, contractor fixed effects cannot absorb *within-contractor* differences in how samples from different zones are processed after collection. If a contractor systematically invests more in the taxonomic analysis of their non-contract area samples than their contract area samples, this would generate the observed pattern. While I cannot directly test for differential post-collection processing, the direction of this explanation is itself consistent with the incentive story: a contractor who allocates more taxonomic effort to non-contract area samples is responding to the same asymmetric incentives that the model describes.

A related possibility is that the gap reflects differences in *reporting timing* rather than deliberate underinvestment in taxonomy. Contractors may submit prospectively-mined tract data to DeepData promptly (before taxonomic analysis is complete) in order to satisfy compliance obligations, while non-contract area samples carry no such pressure and may receive more time for refined identifications before upload. If prospectively-mined tract records are submitted earlier in the taxonomic workflow, they would mechanically show lower species-level reporting even if the underlying scientific effort is identical. This is plausible, but it still reflects

a difference in how contractors handle data from the two zones: the urgency to demonstrate compliance applies specifically to prospectively-mined tracts because of the permit at stake. In other words, even under this interpretation, the asymmetry in reporting precision traces back to the asymmetry in incentives between zones—the same fundamental mechanism the model describes, operating through the timing of data submission rather than the depth of taxonomic analysis. Moreover, if compliance-driven early submission were the primary driver, one might expect contractors to update their records over time; the persistence of the gap suggests that such updates, if they occur, do not fully close it.

A third possibility is that the gap reflects differences in the *composition* of organisms sampled across zones. If contract areas contain taxa that are inherently harder to identify—for example, because they are less well-studied or because no described species exist within those genera—the reporting gap could arise without deflation. The taxonomy robustness exercise (Figure 5) directly addresses this concern by progressively conditioning on taxonomy fixed effects from phylum through genus, comparing species reporting rates only among observations classified to the same taxon. The persistence of the gap after conditioning on genus means that compositional differences in sampled organisms cannot explain the pattern at any taxonomic level from kingdom through genus.

Finally, the dependent variable measures taxonomic resolution as *reported in DeepData*, which may not perfectly reflect the underlying taxonomic work. Lower species-level reporting could reflect incomplete database uploads, delays in taxonomic revision, or administrative bottlenecks rather than reduced scientific effort. Rabone, Horton, et al. (2023) document significant data quality issues in DeepData. While these limitations are real, they would need to affect contract area and non-contract area submissions *differentially* to explain the systematic gap documented here. Symmetric data quality issues would add noise but would not generate a systematic bias in one direction.

7 Policy Options

These findings do not suggest that the private requirement to report data is hopelessly flawed. It may still be the case that this solution is preferable to either a scenario where the seabed is mined without exploration or a scenario where the seabed is explored extensively by a third party prior to mining. In the former case, society might lose out on valuable scientific knowledge, genetic resources, or the opportunity to protect particular discoveries. In the latter case, the third party tasked with exploration would not internalize the significant returns to speed, and might indefinitely delay mining in service of perfect exploration. An ideal system balances the incentive to explore and the incentive to mine with the social optimum. Since the social

benefits of exploration are not well-compensated for, the ISA's solution of using the opportunity to mine as a carrot to incentivize exploration may well be the best that can be done, even in spite of the perverse incentives documented here. This is true even if one takes the more skeptical view that deep-sea mining may not be economically justified once full environmental costs are included (Sumaila et al. 2023): if mining ultimately does not proceed, the biodiversity data generated under the exploration requirement may prove to be the regime's most lasting contribution. What the analysis does show, however, is that there is a meaningful principal-agent problem at play. While the mining company does internalize the opportunity cost of delaying exploration, it does not act as a perfect agent of the ISA when it comes to exploration because it does not put equivalent weight on scientific discovery.

If this problem is considered significant enough to warrant policy changes, the ISA could consider a few different options. The exploitation regulations remain under negotiation, with numerous outstanding issues unresolved (Pickens et al. 2024), so there is an opportunity to address deflation incentives before the regulatory framework is finalized. The model's welfare analysis identifies two distinct wedges—a precision wedge and an exploration wedge—which call for distinct remedies.

The natural remedy would be to identify the uncontracted shirking opportunities and strengthen the reporting requirement to prevent them, effectively reducing the contractor's ability to exploit the precision wedge. Of course, this is easier said than done, as there are many different margins through which taxonomic deflation could occur. Because my results imply non-contract areas are reported with greater taxonomic precision than contract areas, greater attention to how exploration standards differ between zone types could be fruitful. However, the ISA should take great care to prevent the gaming of any updated rules: a requirement that explicitly compares samples from contract and non-contract areas could risk creating an incentive to similarly underreport from outside areas, whereas a requirement that explicitly targets reporting of specific taxonomic levels may induce false specificity. As always, when a measure becomes a target, it ceases to be a good measure.

A more targeted approach to the precision wedge would be to decouple discovery from denial. As the model shows, the contractor depresses taxonomic precision below the socially optimal level because precise identification risks triggering conservation review. Several concrete mechanisms could reduce the regulatory cost of a species discovery. *Safe harbor provisions* would guarantee that a contractor's mining approval cannot be revoked solely on the basis of species discoveries made during its own survey, removing the direct penalty for precise identification. However, such a guarantee could undermine conservation goals. *Sealed reporting* would require contractors to submit their raw biological specimens and taxonomic data to an independent scientific body, which would conduct or verify the taxonomic analysis separately from the permitting process.

The ISA would receive the final taxonomic assessments only after the permitting decision has been made, or would receive them in aggregated or anonymized form that cannot be traced to specific permit applications. This breaks the link between what the contractor discovers and what the regulator uses against it: contractors would have no reason to suppress precise identifications because those identifications could not influence their own permit outcome. The independent body would still make the full taxonomic record publicly available, ensuring that the scientific value of the data is preserved even if conservation becomes less targeted. These mechanisms address the incentive distortion for mining contractors by reducing the regulatory penalty of a species discovery, at the cost of limiting conservation options for the ISA.

A more costly approach would be for the ISA to conduct its own scientific exploration as an *audit of contractor-funded exploration*: this could leave the contractor with primary responsibility for exploration (which incentivizes timeliness) while creating a benchmark of expected reporting to incentivize exploratory effort. Stewart et al. (2026) note that the APEIs remain largely unsurveyed, with little knowledge of what species they harbor; independent baseline surveys of non-contract areas, and APEIs in particular, would both fill this gap and provide an external benchmark against which contractor-submitted data could be evaluated. Jażdżewska and Horton (2026) describe 24 new amphipod species from a single CCZ expedition, underscoring how much latent biodiversity can be discovered with focused scientific effort. The taxonomic capacity to identify CCZ fauna at high resolution exists; what is missing is the institutional mandate to apply it to contractor-collected samples. Such a program would require careful design to ensure it sufficiently penalizes shirking while not punishing companies for the inevitable difficulty and randomness involved with biological sampling in the deep sea. Another option would be *independent taxonomic review*, which would assign post-collection species identification to third-party specialists rather than leaving it to the contractor, eliminating the contractor's ability to influence the precision of reporting. These approaches could lead to more precise data, but at higher cost to the ISA.

To address the exploration wedge, the ISA could incentivize discoveries, either directly by awarding prizes or otherwise allowing monetization of discoveries, or indirectly by strengthening the relationship between high-effort exploration and mining approval. Policies that raise the scientific return to precise identification, such as prizes or advanced commitments for discoveries or facilitating genetic resource markets, would raise both effort and precision in non-contract areas and compress the gap with prospectively-mined tract behavior. This would more closely approximate the economic optimum, where a single agent internalizes all of the relevant costs and benefits for both mining and exploration, but would be difficult to design appropriately and could backfire if it ends up rewarding contractors for continuing with current practice. The ISA could also reward or subsidize exploration by third parties, which could capture some of these benefits without their

risks. Importantly, the two classes of instruments are complements: reducing the regulatory penalty of species discovery amplifies the contractor's precision response to scientific incentives. However, these approaches have their own monetary costs or legal challenges that mean there are tradeoffs to their adoption.

These policy directions connect to the model's further implications. If deflation intensifies with tract value, monitoring and enforcement resources should be concentrated on the highest-value exploration contracts. If the regulatory-bite implication holds, then strengthening precautionary standards could inadvertently amplify deflation incentives unless accompanied by mechanisms that decouple discovery from denial. And the scientific-returns implication suggests that even modest investments in raising the payoff to precise identification (through data prizes, publication incentives, or access to genetic resource markets) could yield outsized improvements in reporting quality.

These policy directions have parallels in the literature on endangered species regulation. The ESA creates similar perverse incentives for private landowners who might harbor listed species (Brown and Shogren 1998; Langpap et al. 2018), and safe-harbor agreements have been shown to mitigate preemptive habitat destruction in that context (Byl 2019). Polasky and Doremus (1998) show that ESA policy gives landowners no incentive to cooperate with species information collection, and Innes et al. (1998) find that the design of takings compensation substantially affects information provision incentives. However, the deep-sea mining setting differs in an important respect: landowners under the ESA already hold property rights over their land, whereas mining contractors do not own the seabed and are applying to use a common heritage resource. There is less legal architecture to grant property rights over biological discoveries in the deep sea, which limits the applicability of compensation-based solutions that have proven effective on land. This makes the design of reporting incentives (sealed reporting, independent taxonomic review, discovery prizes) all the more important. Moreover, the entry into force of the Agreement on Biodiversity Beyond National Jurisdiction (BBNJ) introduces new obligations around environmental impact assessment and marine genetic resource benefit-sharing that interact with the ISA's existing regime in ways that remain unresolved (Willaert and Soete 2025). If the BBNJ framework creates additional regulatory consequences for species discoveries, it could amplify rather than reduce the incentive to deflate, making it all the more urgent to decouple discovery from denial before these regimes fully intersect.

Each option carries costs and risks. Overly prescriptive requirements could lead to false specificity, while excessive reliance on third-party surveys could delay mining without improving biodiversity knowledge. Nonetheless, without action, taxonomic deflation threatens the scientific integrity of deep-sea biodiversity assessments.

8 Conclusion

The motivating question of this paper has a partial answer: the imprecision of deep-sea mining data is not fully explained by scientific difficulty. The difficulty is real, but symmetric across zones; the asymmetry that the same contractors report less precise data from the tracts they intend to mine is the signature of taxonomic deflation. The counterfactual decomposition suggests that even by the most conservative estimate, strategic incentives account for roughly 7% of the missing species identifications in DeepData—around 8,000 observations under that conservative bound, and as many as 37,000 under the largest—all of which would carry a species name if discovery did not threaten the permit. The underlying reality is that no single entity has fully aligned incentives to balance exploration and exploitation of the deep sea, and the ISA's clever solution—using the prospect of mining as a carrot for biodiversity reporting—inevitably creates the precision wedge documented here. Whether that wedge is small enough to tolerate or large enough to fix is a choice the exploitation regulations now under negotiation will encode, and those choices will determine whether the deep-sea biodiversity record reflects what the data could show or only what mining contractors find it convenient to report.

References

Angrist, Joshua David, and Jörn-Steffen Pischke. 2008. *Mostly Harmless Econometrics : An Empiricist's Companion*. Princeton University Press.

Ardron, Jeff A. 2018. "Transparency in the Operations of the International Seabed Authority: An Initial Assessment." *Marine Policy* 95 (September): 324–31. https://doi.org/10.1016/j.marpol.2016.06.027.

Arrow, Kenneth J., and Anthony C. Fisher. 1974. "Environmental Preservation, Uncertainty, and Irreversibility." *The Quarterly Journal of Economics* 88 (2): 312–19. https://doi.org/10.2307/1883074.

Assis, Jorge, Salvador Jesús Fernández Bejarano, Vinícius W. Salazar, et al. 2024. "Bio-ORACLE V3.0. Pushing Marine Data Layers to the CMIP6 Earth System Models of Climate Change Research." *Global Ecology and Biogeography* 33 (4): e13813. https://doi.org/10.1111/geb.13813.

Bell, Katherine L. C., Kristen N. Johannes, Brian R. C. Kennedy, and Susan E. Poulton. 2025. "How Little We've Seen: A Visual Coverage Estimate of the Deep Seafloor." *Science Advances* 11 (19): eadp8602.

https://doi.org/10.1126/sciadv.adp8602.

Bennear, Lori S., and Sheila M. Olmstead. 2008. "The Impacts of the 'Right to Know': Information Disclosure and the Violation of Drinking Water Standards." *Journal of Environmental Economics and Management* 56 (2): 117–30. https://doi.org/10.1016/j.jeem.2008.03.002.

Blanchard, Catherine, Ellycia Harrould-Kolieb, Emily Jones, and Michelle L. Taylor. 2023. "The Current Status of Deep-Sea Mining Governance at the International Seabed Authority." *Marine Policy* 147 (January): 105396. https://doi.org/10.1016/j.marpol.2022.105396.

Bonifácio, Paulo, Lenka Neal, and Lénaïck Menot. 2021. "Diversity of Deep-Sea Scale-Worms (Annelida, Polynoidae) in the Clarion-Clipperton Fracture Zone." *Frontiers in Marine Science* 8.

Brown, Gardner M., and Jason F. Shogren. 1998. "Economics of the Endangered Species Act." *Journal of Economic Perspectives* 12 (3): 3–20. https://doi.org/10.1257/jep.12.3.3.

Byl, Jacob P. 2019. "Perverse Incentives and Safe Harbors in the Endangered Species Act: Evidence From Timber Harvests Near Woodpeckers." *Ecological Economics* 157 (March): 100–108. https://doi.org/10.1016/j.ecolecon.2018.11.008.

Claridge, Michael F., Hassan A. Dawah, and Michael R. Wilson, eds. 1997. *Species : The Units of Biodiversity*. The Systematics Association Special Volume Series. Chapman & Hall.

Costello, Mark J., and Chhaya Chaudhary. 2017. "Marine Biodiversity, Biogeography, Deep-Sea Gradients, and Conservation." *Current Biology* 27 (11): R511–27. https://doi.org/10.1016/j.cub.2017.04.060.

de Marchi, Scott, and James T. Hamilton. 2006. "Assessing the Accuracy of Self-Reported Data: An Evaluation of the Toxics Release Inventory." *Journal of Risk and Uncertainty* 32 (1): 57–76. https://doi.org/10.1007/s10797-006-6666-3.

Doshi, Anil R., Glen WS Dowell, and Michael W. Toffel. 2013. "How Firms Respond to Mandatory Information Disclosure." *Strategic Management Journal* 34 (10): 1209–31.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2013. "Truth-Telling by Third-Party Auditors and the Response of Polluting Firms: Experimental Evidence from India." *The Quarterly Journal of Economics* 128 (4): 1499–545.

Frutos, Inmaculada, Stefanie Kaiser, Łukasz Pułaski, Maciej Studzian, and Magdalena Błażewicz. 2022. "Challenges and Advances in the Taxonomy of Deep-Sea Peracarida: From Traditional to Modern Methods." *Frontiers in Marine Science* Volume 9 - 2022.

George, Anna, and Richard Mayden. 2005. "Species Concepts and the Endangered Species Act: How a Valid Biological Definition of Species Enhances the Legal Protection of Biodiversity." *Natural Resources Journal* 45 (March): 369–407.

Ghezelayagh, Ava, Jeffrey W. Simmons, Julia E. Wood, et al. 2025. "Comparative Species Delimitation of a Biological Conservation Icon." *Current Biology* 35 (2): 398–406.e4. https://doi.org/10.1016/j.cub.2024.11.053.

Gonzalez-Lira, Andres, and Ahmed Mushfiq Mobarak. 2021. "Slippery Fish: Enforcing Regulation When Agents Learn and Adapt." *National Bureau of Economic Research Working Paper Series* No. 28610. https://doi.org/10.3386/w28610.

Harrington, Winston. 1988. "Enforcement Leverage When Penalties Are Restricted." *Journal of Public Economics* 37 (1): 29–53. https://doi.org/10.1016/0047-2727(88)90003-5.

Henry, Claude. 1974. "Investment Decisions Under Uncertainty: The "Irreversibility Effect"." *The American Economic Review* 64 (6): 1006–12.

Heyes, Anthony, and Neil Rickman. 1999. "Regulatory Dealing – Revisiting the Harrington paradox1We Are Grateful to Lars Hansen, Robert Cairns, Jeff Frank, Two Referees from This Journal and Seminar Participants at the University of London for Helpful Comments. The Usual Disclaimer Applies.1." *Journal of Public Economics* 72 (3): 361–78. https://doi.org/10.1016/S0047-2727(98)00098-X.

Innes, Robert, Stephen Polasky, and John Tschirhart. 1998. "Takings, Compensation and Endangered Species Protection on Private Lands." *Journal of Economic Perspectives* 12 (3): 35–52. https://doi.org/

10.1257/jep.12.3.35.

Isaac, Nick J. B., James Mallet, and Georgina M. Mace. 2004. "Taxonomic Inflation: Its Influence on Macroecology and Conservation." *Trends in Ecology & Evolution* 19 (9): 464–69. https://doi.org/10.1016/j.tree.2004.06.004.

Jaeckel, Aline L. 2017. *The International Seabed Authority and the Precautionary Principle: Balancing Deep Seabed Mineral Mining and Marine Environmental Protection*. Brill | Nijhoff. https://doi.org/10.1163/9789004332287.

Jażdżewska, Anna M., Karolina Biniek, Pedro Martínez Arbizu, and Annemiek Vink. 2025. "Hidden Behind the Scene: High Diversity, Low Connectivity of Deep-Sea Amphipoda in the Polymetallic Nodule Fields in the Northeast Pacific." *EGUsphere*, May, 1–50. https://doi.org/10.5194/egusphere-2025-1794.

Jażdżewska, Anna M., and Tammy Horton. 2026. "New Deep-Sea Amphipoda from the Clarion-Clipperton Zone: 24 New Species Described Under the Sustainable Seabed Knowledge Initiative: One Thousand Reasons Campaign." *ZooKeys* 1274 (March): 1–16. https://doi.org/10.3897/zookeys.1274.176711.

Jones, Daniel O. B., Maria Belen Arias, Loïc Van Audenhaege, et al. 2025. "Long-Term Impact and Biological Recovery in a Deep-Sea Mining Track." *Nature* 642 (8066): 112–18. https://doi.org/10.1038/s41586-025-08921-3.

Jones, Daniel O. B., Jennifer M. Durden, Kevin Murphy, et al. 2019. "Existing Environmental Management Approaches Relevant to Deep-Sea Mining." *Marine Policy* 103 (May): 172–81. https://doi.org/10.1016/j.marpol.2019.01.006.

Kennedy, Brian R. C., Kasey Cantwell, Mashkoor Malik, et al. 2019. "The Unknown and the Unexplored: Insights Into the Pacific Deep-Sea Following NOAA CAPSTONE Expeditions." *Frontiers in Marine Science* 6 (August). https://doi.org/10.3389/fmars.2019.00480.

Kennedy, Brian R. C., and Randi D. Rotjan. 2023. "Mind the Gap: Comparing Exploration Effort with Global Biodiversity Patterns and Climate Projections to Determine Ocean Areas with Greatest Exploration Needs." *Frontiers in Marine Science* 10 (November). https://doi.org/10.3389/fmars.2023.1219799.

Kim, Eun-Hee, and Thomas P. Lyon. 2011. "Strategic Environmental Disclosure: Evidence from the DOE's Voluntary Greenhouse Gas Registry." *Journal of Environmental Economics and Management* 61 (3): 311–26. https://doi.org/10.1016/j.jeem.2010.11.001.

Krusberg, Tilde, Lova Schildt, Jean-Baptiste Jouffray, Erik Zhivkoplias, and Robert Blasiak. 2024. "A Review of Marine Genetic Resource Valuations." *Npj Ocean Sustainability* 3 (1): 46. https://doi.org/10.1038/s44183-024-00081-7.

Kürzel, Karlotta, Stefanie Kaiser, Anne-Nina Lörz, et al. 2022. "Correct Species Identification and Its Implications for Conservation Using Haploniscidae (Crustacea, Isopoda) in Icelandic Waters as a Proxy." *Frontiers in Marine Science* Volume 8 - 2021.

Langpap, Christian, Joe Kerkvliet, and Jason F. Shogren. 2018. "The Economics of the U.S. Endangered Species Act: A Review of Recent Developments." *Review of Environmental Economics and Policy* 12 (1): 69–91. https://doi.org/10.1093/reep/rex026.

Levin, Lisa A., Diva J. Amon, and Hannah Lily. 2020. "Challenges to the Sustainability of Deep-Seabed Mining." *Nature Sustainability* 3 (10): 784–94. https://doi.org/10.1038/s41893-020-0558-x.

Lueck, Dean, and Jeffrey A. Michael. 2003. "Preemptive Habitat Destruction Under the Endangered Species Act." *The Journal of Law and Economics* 46 (1): 27–60. https://doi.org/10.1086/344670.

Lyon, Thomas P., and John W. Maxwell. 2011. "Greenwash: Corporate Environmental Disclosure Under Threat of Audit." *Journal of Economics & Management Strategy* 20 (1): 3–41. https://doi.org/10.1111/j.1530-9134.2010.00282.x.

Mu, Yingfei, Edward A Rubin, and Eric Zou. 2021. "What's Missing in Environmental (Self-)Monitoring: Evidence from Strategic Shutdowns of Pollution Monitors." *National Bureau of Economic Research Working Paper Series* No. 28735. https://doi.org/10.3386/w28735.

Niner, Holly J., Jeff A. Ardron, Elva G. Escobar, et al. 2018. "Deep-Sea Mining With No Net Loss of Biodiversity—An Impossible Aim." *Frontiers in Marine Science* Volume 5 - 2018.

Pickens, Chris, Hannah Lily, Ellycia Harrould-Kolieb, Catherine Blanchard, and Anindita Chakraborty. 2024. "From What-If to What-Now: Status of the Deep-Sea Mining Regulations and Underlying Drivers for Outstanding Issues." *Marine Policy* 169 (November): 105967. https://doi.org/10.1016/j.marpol.2023.105967.

Polasky, Stephen, and Holly Doremus. 1998. "When the Truth Hurts: Endangered Species Policy on Private Land with Imperfect Information." *Journal of Environmental Economics and Management* 35 (1): 22–47. https://doi.org/10.1006/jeem.1998.1021.

Rabone, Muriel, T Horton, D O B Jones, E Simon-Lledó, and A G Glover. 2023. "A Review of the International Seabed Authority Database DeepData from a Biological Perspective: Challenges and Opportunities in the UN Ocean Decade." *Database* 2023 (January): baad013. https://doi.org/10.1093/database/baad013.

Rabone, Muriel, Joris H. Wiethase, Erik Simon-Lledó, et al. 2023. "How Many Metazoan Species Live in the World's Largest Mineral Exploration Region?" *Current Biology* 33 (12): 2383–2396.e5. https://doi.org/10.1016/j.cub.2023.04.052.

Smith, Craig R., Malcolm R. Clark, Erica Goetze, Adrian G. Glover, and Kerry L. Howell. 2021. "Biodiversity, Connectivity and Ecosystem Function Across the Clarion-Clipperton Zone: A Regional Synthesis for an Area Targeted for Nodule Mining." *Frontiers in Marine Science* 8.

Stewart, Eva C. D., Helena Wiklund, Lenka Neal, et al. 2026. "Impacts of an Industrial Deep-Sea Mining Trial on Macrofaunal Biodiversity." *Nature Ecology & Evolution* 10 (2): 318–29. https://doi.org/10.1038/s41559-025-02911-4.

Sumaila, U. R., L. Alam, K. Pradhoshini, et al. 2023. "To Engage in Deep-Sea Mining or Not to Engage: What Do Full Net Cost Analyses Tell Us?" *Npj Ocean Sustainability* 2 (1): 19. https://doi.org/10.1038/s44183-023-00030-w.

Washburn, Travis W., Daniel O. B. Jones, Chih-Lin Wei, and Craig R. Smith. 2021. "Environmental Heterogeneity Throughout the Clarion-Clipperton Zone and the Potential Representativity of the APEI Network." *Frontiers in Marine Science* 8.

Wedding, L. M., A. M. Friedlander, J. N. Kittinger, et al. 2013. "From Principles to Practice: A Spatial Approach to Systematic Conservation Planning in the Deep Sea." *Proceedings of the Royal Society B: Biological Sciences* 280 (1773): 20131684. https://doi.org/10.1098/rspb.2013.1684.

Willaert, Klaas, and Anemoon Soete. 2025. "The Interaction Between the BBNJ Agreement and the International Deep Sea Mining Regime: More Questions Than Answers?" *Ocean Development & International Law* 56 (2): 176–213. https://doi.org/10.1080/00908320.2025.2472634.

Zakota, Mark. 2025. "EPA Scrutiny and Voluntary Environmental Disclosures." *Review of Accounting Studies* 30 (4): 3514–57. https://doi.org/10.1007/s11142-025-09901-8.

Zhivkoplias, Erik, Jean-Baptiste Jouffray, Paul Dunshirn, Agnes Pranindita, and Robert Blasiak. 2024. "Growing Prominence of Deep-Sea Life in Marine Bioprospecting." *Nature Sustainability* 7 (8): 1027–37. https://doi.org/10.1038/s41893-024-01392-w.

Appendix

The appendix is organized into three sections. **Extra Tables** and **Extra Figures** present material referenced from the main text that did not meet the bar for in-line inclusion. **Additional Analyses** presents exercises that are tangential to the central argument but that may be of interest for evaluating specific concerns.

Extra Tables

Table 3: Summary statistics for biological specimen records in the DeepData database. Contract Areas are contractor-allocated exploration tracts; APEIs are Areas of Particular Environmental Interest; Outside Areas are samples outside both zone types. Species reporting rate is the proportion of records that include a species-level identification.

Zone	Specimen Records	Contractors	Year Range	Species Reporting Rate
Contract Areas	149,795	12	2004–2021	16.4%
APEIs	366	6	2015–2018	41.5%
Outside Areas	10,806	7	2010–2021	44.4%
Total	160,967	12	2004–2021	18.3%

Table 4: Linear probability model estimates of species-level reporting. Dependent variable: indicator for species-level taxonomic identification. Comparison group is non-contract areas (APEIs and outside areas). Robust standard errors in parentheses.

	Contract Area	+ Contractor FE	+ Year FE	+ Environmental Controls
Intercept	0.443***	0.529***	1.233***	−23.842***
	(0.005)	(0.026)	(0.037)	(1.830)
Contract Area	−0.279***	−0.192***	−0.091***	−0.060***
	(0.005)	(0.005)	(0.006)	(0.006)
Depth (m)				0.000***
				(0.000)
Sea Surface Temp.				0.320***
				(0.022)
Surface Chlorophyll				−5.603***
				(0.507)
Dissolved Oxygen				3.496***
				(0.274)
Contractor FE	No	Yes	Yes	Yes
Year FE	No	No	Yes	Yes
Environmental Controls	No	No	No	Yes
Num.Obs.	160 967	160 967	160 967	150 104
R2	0.034	0.123	0.213	0.211

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

Table 5: Taxonomy fixed effects robustness. Dependent variable: species-level identification. Comparison group is non-contract areas (APEIs and outside areas). All specifications include contractor, year, and environmental controls. Each specification restricts the sample to observations with a valid identification at the relevant taxonomic level. Robust standard errors in parentheses.

	Baseline	+ Phylum FE	+ Class FE	+ Order FE	+ Family FE	+ Genus FE
Intercept	−23.842***	−13.079***	−6.905***	−1.407	−4.165**	−8.189***
	(1.830)	(1.629)	(1.630)	(1.933)	(1.926)	(1.444)
Contract Area	−0.060***	−0.051***	−0.051***	−0.043***	−0.049***	−0.068***
	(0.006)	(0.006)	(0.005)	(0.005)	(0.004)	(0.004)
Contractor FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Environmental Controls	Yes	Yes	Yes	Yes	Yes	Yes
Phylum FE	No	Yes	No	No	No	No
Class FE	No	No	Yes	No	No	No
Order FE	No	No	No	Yes	No	No
Family FE	No	No	No	No	Yes	No
Genus FE	No	No	No	No	No	Yes
Num.Obs.	150 104	139 331	131 863	85 873	73 042	63 989
R2	0.211	0.330	0.560	0.662	0.768	0.823

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

Sample Source Composition

Each biological specimen record in DeepData is reported by a specific mining contractor and is associated with a specific physical location. Using the official ISA contract polygon shapefiles—downloaded from the ISA's Maps repository at `isa.org.jm/exploration-contracts/maps/`—I match each observation to the contract polygon (if any) within whose boundary it falls. Combined with the identity of the *sampling* contractor (extracted from the `area_key` field), this yields four mutually exclusive observation types:

- **Own tract**: the observation falls inside the sampling contractor's own PMN polygon—this is the regulatory regime where a permit is at stake.
- **Other tract**: the observation falls inside *another* contractor's PMN polygon—e.g., a German (BGR) sample collected within Singapore's (OMS) tract during a joint or transit cruise. The sampling contractor has no permit-related incentive at this location.
- **APEI**: the observation falls inside an Area of Particular Environmental Interest, where mining is prohibited.
- **Outside all**: the observation falls in international waters outside any contract polygon and outside any APEI.

This decomposition is helpful both descriptively—it characterizes where the dataset's biological samples

actually come from—and methodologically, since the four types differ in their regulatory exposure. Table 6 reports counts and species-level reporting rates for each type. Roughly 93% of CCZ specimens are collected in the sampling contractor's own tract; 6% are collected by one contractor inside another's tract; and APEIs and truly-outside observations together make up less than 1%. The pattern in the species-reporting column is striking. Specimens collected by a contractor within their own tract are reported with species-level identification only 17.6% of the time. The same contractors, when they sample inside *another* contractor's tract, report at the species level 46.7% of the time. In APEIs, where mining is prohibited and there is no commercial interest at stake, contractor-collected specimens reach 39.9% species-level reporting. The sharp gap between "own tract" and "other tract" samples that share the same physical region motivates the within-sampler and within-polygon robustness exercises reported in Table **??** of the main text.

Table 6: Composition of biological observations by spatial source. "Own tract" = inside the sampling contractor's own PMN polygon; "Other tract" = inside another contractor's PMN polygon; "APEI" = inside an Area of Particular Environmental Interest; "Outside all" = outside any PMN polygon or APEI. Spatial classification uses official ISA shapefiles.

Sample source	Observations	Share of total	Species-level reporting
Own tract	140,255	93.2%	17.6%
Other tract	9,045	6.0%	46.7%
APEI	358	0.2%	39.9%
Outside all	801	0.5%	48.6%

Extra Figures

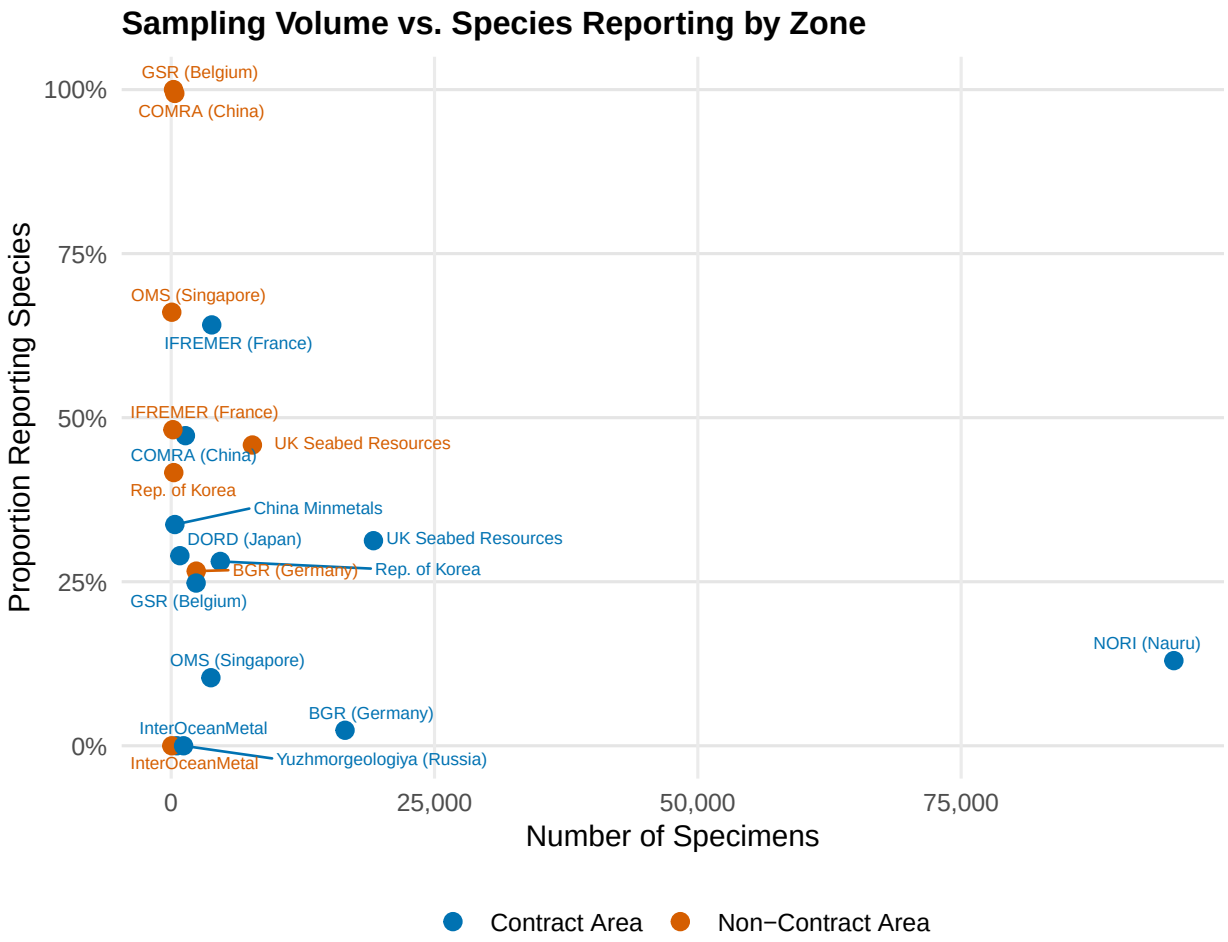

Figure 8: Sampling volume (number of biological specimens) versus species reporting rate by contractor and zone type. Each point represents one contractor's submissions from either contract areas (blue) or non-contract areas (vermillion). The negative cross-zone correlation—more samples but lower species reporting in contract areas—is consistent with the model's twin predictions and difficult to reconcile with a resource-constraints explanation.

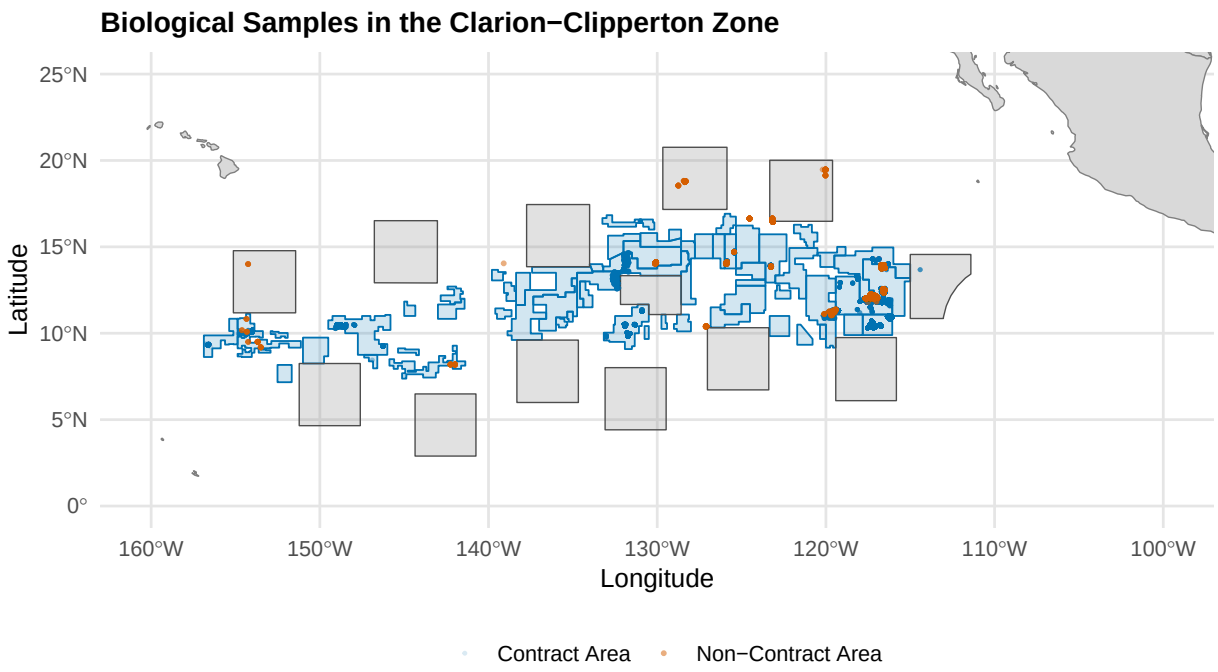

Figure 9: Sample locations overlaid on official ISA contract area polygons (blue) and APEIs (grey). Each point is one biological sample submitted to DeepData. Source for polygons: International Seabed Authority.

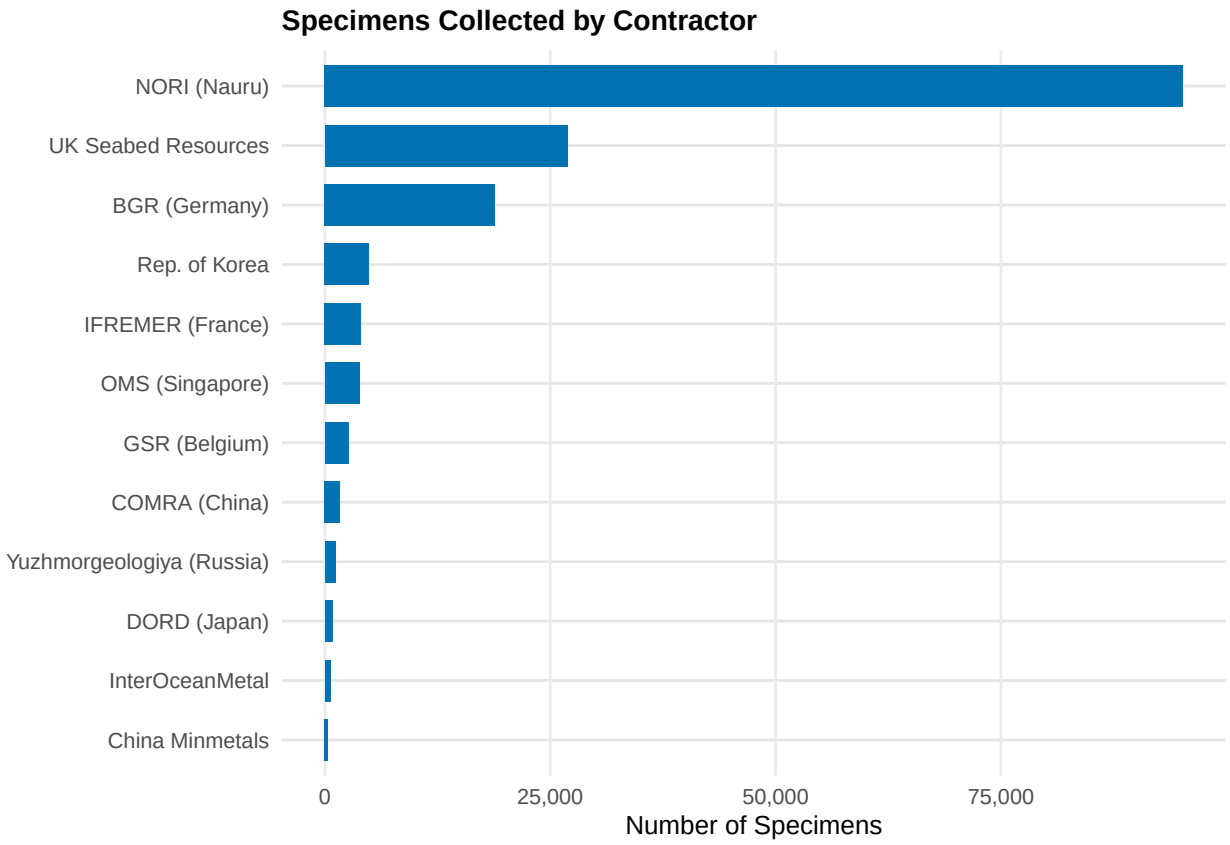

Figure 10: Number of biological specimens collected by each contractor in the CCZ.

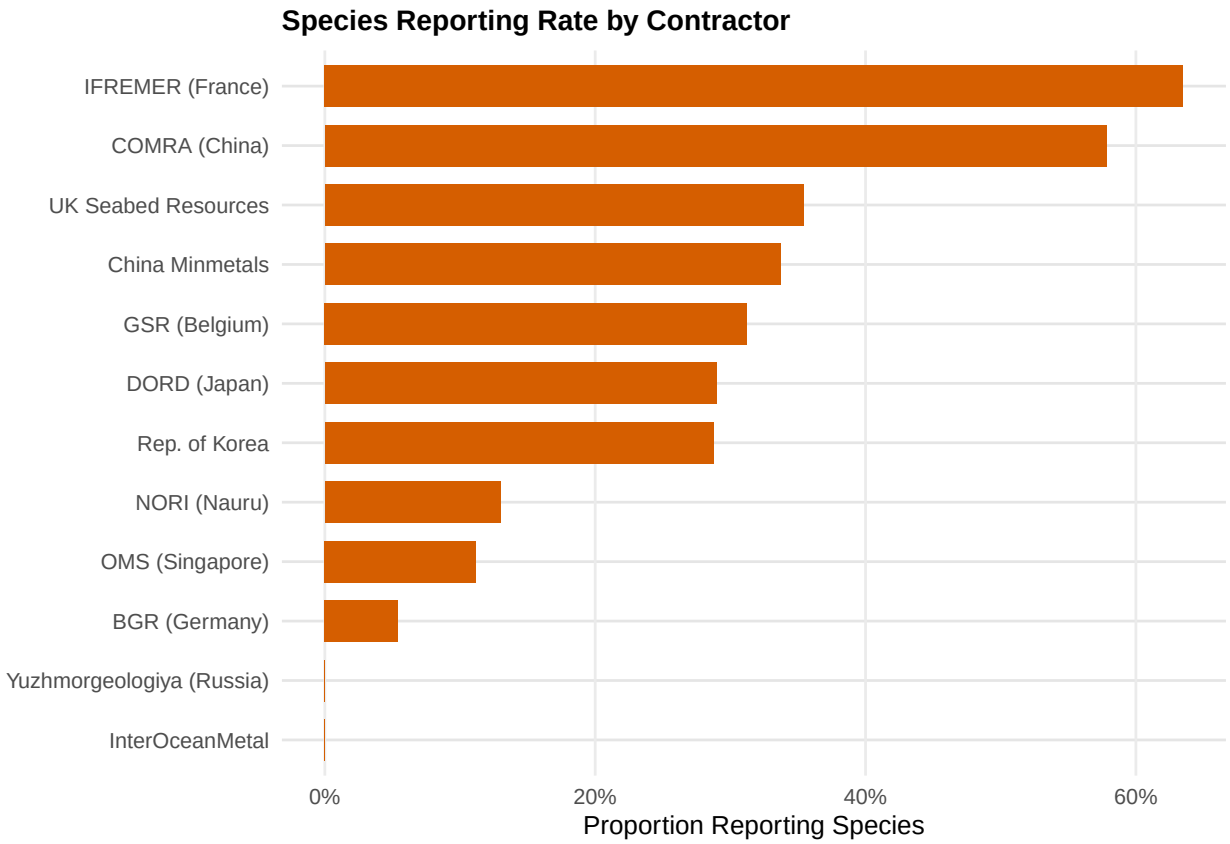

Figure 11: Proportion of observations reporting species-level identification by contractor.

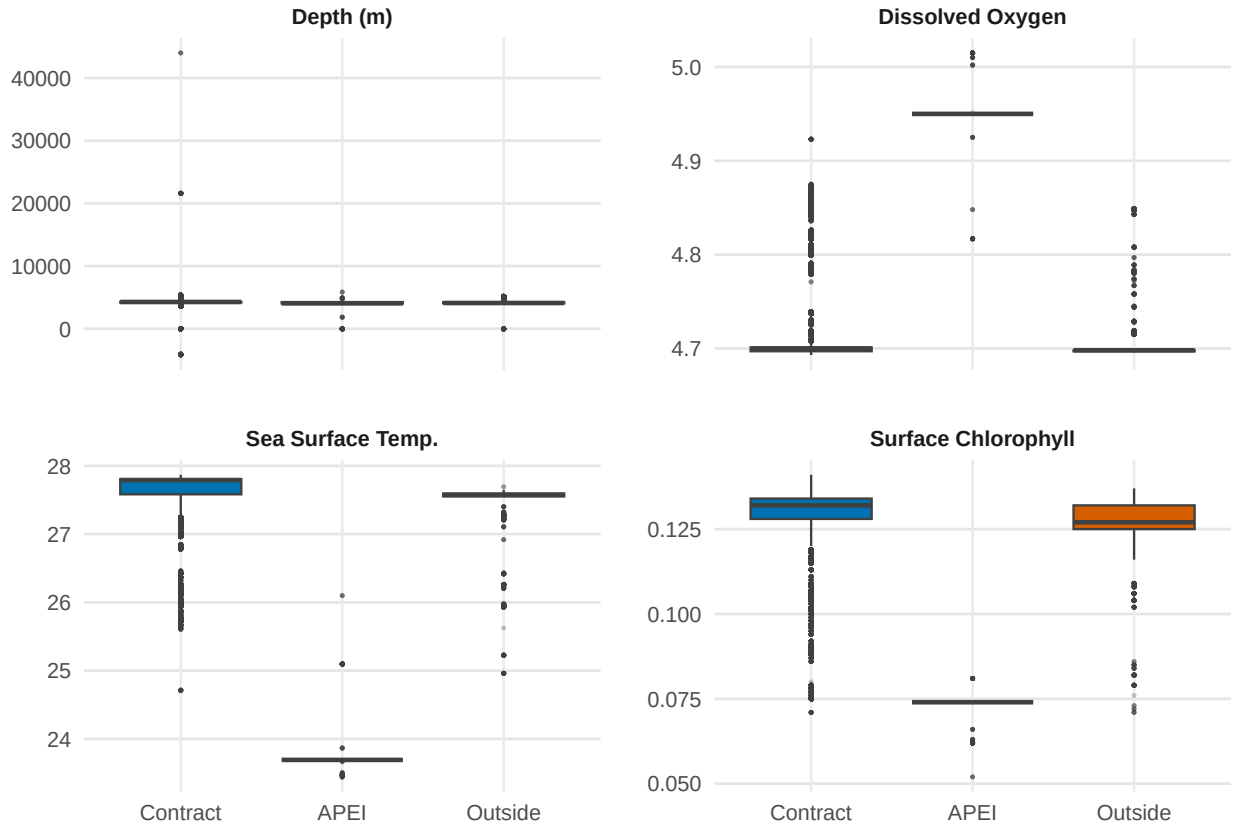

Figure 12: Distributions of the four environmental covariates used in the main analysis, by zone type. Bio-ORACLE values are matched to each sample's longitude and latitude. The contract-area and APEI distributions overlap heavily, consistent with the ISA's representativity criterion for APEI siting; outside areas show somewhat broader spread on most variables. The remaining marginal heterogeneity is absorbed by the environmental controls in the main regression.

Additional Analyses

APEI-Only Comparison

Tables 7 and 8 replicate the main regression analyses restricting the comparison group to Areas of Particular Environmental Interest only, excluding outside areas. The APEI-only sample is the cleanest comparison: APEIs are explicitly designated conservation zones with no mining permit at stake, and all APEI samples in the dataset are contractor-submitted, holding fixed differences in reporting infrastructure and scientific methodology. The sample is considerably smaller (APEIs contribute only 374 records), but the directional results are consistent with the broader analysis.

Table 7: APEI-only comparison. Linear probability model estimates restricting the sample to contract areas and APEIs. Dependent variable: indicator for species-level taxonomic identification. Robust standard errors in parentheses.

	APEI	+ Contractor FE	+ Year FE	+ Environmental Controls
Intercept	0.164***	0.337***	1.077***	−46.931***
	(0.001)	(0.025)	(0.034)	(2.099)
APEI	0.243***	0.088***	0.063*	0.822***
	(0.025)	(0.026)	(0.034)	(0.088)
Depth (m)				0.000***
				(0.000)
Sea Surface Temp.				0.762***
				(0.034)
Surface Chlorophyll				−10.259***
				(0.516)
Dissolved Oxygen				5.863***
				(0.288)
Contractor FE	No	Yes	Yes	Yes
Year FE	No	No	Yes	Yes
Environmental Controls	No	No	No	Yes
Num.Obs.	150 161	150 161	150 161	139 675
R2	0.001	0.099	0.219	0.220

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

Table 8: APEI-only taxonomy fixed effects robustness. Sample restricted to contract areas and APEIs. All specifications include contractor, year, and environmental controls. Robust standard errors in parentheses.

	Baseline	+ Phylum FE	+ Class FE	+ Order FE	+ Family FE	+ Genus FE
Intercept	−46.931***	−30.766***	−22.993***	−18.771***	−12.226***	−5.690***
	(2.099)	(1.899)	(2.020)	(2.324)	(2.421)	(1.885)
APEI	0.822***	0.468***	0.777***	0.853***	0.663***	0.421***
	(0.088)	(0.067)	(0.084)	(0.089)	(0.104)	(0.078)
Contractor FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Environmental Controls	Yes	Yes	Yes	Yes	Yes	Yes
Phylum FE	No	Yes	No	No	No	No
Class FE	No	No	Yes	No	No	No
Order FE	No	No	No	Yes	No	No
Family FE	No	No	No	No	Yes	No
Genus FE	No	No	No	No	No	Yes
Num.Obs.	139 675	128 942	121 777	77 584	67 253	58 984
R2	0.220	0.339	0.576	0.657	0.759	0.825

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

The pattern in the species-reporting column is striking. Specimens collected by a contractor within their own tract are reported with species-level identification only 17.6% of the time. The same contractors, when they sample inside *another* contractor's tract, report at the species level 46.7% of the time. In APEIs, where

mining is prohibited and there is no commercial interest at stake, contractor-collected specimens reach 39.9% species-level reporting. The sharp gap between "own tract" and "other tract" samples that share the same physical region motivates the within-sampler and within-polygon robustness exercises reported in Table **??** of the main text.