Unintended Consequences of a Clean Regulation: How the U.S. Ballast Water Act Led to Transboundary Invasion of Alien Species

F. Candau,[*][†] F. Lafferrere,[‡] J. Schlick[§]

April 23, 2026

Abstract

We study the transboundary ecological incidence of the U.S. Ballast Water Management Act (2012), which requires vessels to exchange coastal ballast water in the open ocean before entering U.S. ports. Using particle tracking through ocean currents, we trace where organisms released during mid-ocean exchange drift along Caribbean and Pacific coastlines. Coastlines that receive more of those organisms record significantly more invasive species after 2012. We also find that more species arrive, more of them become common, and no single species dominates.

JEL codes: Q56, Q58, F18, R41

Keywords: introduced species, ballast water, environmental regulation, maritime trade, transboundary externalities

[*]UPPA E2S, TREE, France. fabien.candau@univ-pau.fr

[†]We would like to express our gratitude to all the participants of several seminars for their comments (RITM, BSE, TREE) as well as participants at the 29th Annual Conference of the EAERE and at the Conference of the International Association of Maritime Economists at Bergen. We are grateful for the excellent working conditions provided by Paris Saclay and Pompeu Fabra University. Finally, we thank the Region Nouvelle-Aquitaine and the UPPA (project PORTeur) for funding this research.

[‡]UPPA E2S, TREE, France.

[§]Paris Saclay, RITM, France.

1 Introduction

Ballast water carries life across oceans. Every year, commercial vessels transfer approximately 3–5 billion metric tons of seawater for stability, together with the bacteria, plankton, invertebrates, and larval stages it contains (Molnar et al., 2008). When discharged, these organisms can establish in recipient ecosystems, damaging fisheries, aquaculture, tourism, and coastal infrastructure, with global costs estimated in the billions of dollars annually (Bax et al., 2003; Pimentel et al., 2005; Cuthbert et al., 2021). Ballast water is now recognized as one of the primary pathways through which marine species cross biogeographic boundaries (Molnar et al., 2008; Bax et al., 2003).

Countries have responded with regulation. But where regulation protects domestic ports, it may simultaneously disrupt foreign coastlines. When a country mandates that vessels exchange their ballast water mid-ocean before arrival, organisms are released not at the destination but along the approach route—in the waters of nations that did not set the rules. Whether ballast-water regulation protects one country's ecosystems at the expense of others is unknown. To our knowledge, no econometric study has traced the ecological consequences of a specific ballast-water regulation to outcomes at the coastlines of transit countries.

The United States is a natural setting for this question. As one of the world's largest trading nations, it generates thousands of vessel calls per year along Caribbean and Pacific shipping lanes that traverse the waters of dozens of small and mid-sized coastal states. In 2012, the U.S. Ballast Water Management Act required arriving vessels to exchange their ballast water in the open ocean before entering U.S. ports. Given the volume of U.S.-bound traffic, this compliance requirement could affect ecosystems far from U.S. shores.

In this paper, we show that the Act exported an invasive-species externality to transit-country coastlines through the very compliance mechanism it mandated. One key empirical challenge is how to connect those offshore events to coastal ecological outcomes. The organisms released there do not stay put; surface currents carry them hundreds of kilometers, and a discharge event south of Jamaica may deliver organisms to the Central American coast (Richardson, 2005). We address this by developing a Lagrangian identification strategy.[1] Using voyage-level ballast-water management reports from the U.S. National Ballast Information

[1] In fluid dynamics, a *Lagrangian* description tracks individual fluid parcels along their trajectories, in contrast to the *Eulerian* description, which records velocity at fixed grid points. Lagrangian particle tracking releases virtual tracer particles at a source location and integrates their positions forward in time using the local current velocity field. The term is named after Joseph-Louis Lagrange's 18th-century work on fluid mechanics and is unrelated to the constrained-optimization Lagrangian familiar in economics.

Clearinghouse, we identify the coordinates, date, and volume of each exchange event. We then release virtual particles at each location and advect them for 60 days through a monthly climatology of ocean-current fields, assigning each particle's final position to one of approximately 2,400 coastal segments (25 km arcs, buffered 100 km seaward) across 39 Caribbean and Pacific transit countries. The resulting treatment variable measures the drift-weighted volume of ballast-water exchange reaching each segment in each year.

We find that coastal segments with more drift exposure experience significantly higher invasive species occurrence, total abundance, and Shannon diversity in both the exchange and treatment eras. Three complementary composition indicators show that drift exposure increases species richness and the number of species reaching ecologically meaningful presence, while reducing single-species dominance (the Berger-Parker index falls significantly). A falsification test using monitoring effort as a placebo outcome finds no differential increase in survey activity at treated segments, making it unlikely that the result reflects changes in scientific survey activity rather than ecological change.

These findings matter for how countries design compliance pathways in environmental regulation. The International Maritime Organization (IMO) adopted exchange-based compliance as a transitional standard while treatment technology phased in. Our results show that this transitional measure shifted introduced-species exposure to transit countries that did not participate in the regulatory decision and cannot control the shipping traffic crossing their waters. Wang et al. (2021) document that the IMO Ballast Water Management Convention generates unequal economic burdens for Small Island Developing States; we show that introduced-species exposure is unequally distributed as well. As the IMO and individual countries design compliance frameworks for other maritime environmental regulations—scrubber discharge, greenhouse gas emissions, biofouling—the transboundary incidence of interim measures deserves explicit attention.

Our work departs from existing research in several ways. An environmental economics literature has established that trade is a vector for biological invasion and that policy design matters: McAusland and Costello (2004) show that optimal biosecurity policy combines tariffs with partner-specific inspections; Costello et al. (2007) confirm empirically that invasion risk varies across U.S. trading partners; Margolis et al. (2005) show that political-economy forces make it difficult to distinguish biosecurity tariffs from disguised protection. These papers model the trade-invasion link but do not examine the spatial consequences of a specific regulation on foreign ecosystems. To our knowledge, this is the first econometric estimate of that spatial incidence.

A biological and engineering literature evaluates ballast-water treatment technologies and finds that no single system eliminates all organisms (Bradie et al., 2018; Hess-Erga et al., 2019; Lakshmi et al., 2021). Separately, marine ecologists have used Lagrangian particle tracking to model the dispersal of ballast-water organisms through ocean currents (Lezama-Ochoa et al., 2022; Dobrzycka-Krahel et al., 2024), demonstrating the physical plausibility of the drift channel. Unlike these studies, which use connectivity modeling for risk assessment, we embed these methods within an econometric framework that allows causal inference about regulatory outcomes.

A complementary economics literature studies the local effects of maritime emissions regulation on air quality and human health. Hansen-Lewis and Marcus (2022) find that U.S. maritime emission-control areas significantly reduced coastal fine particulate matter, low birth weight, and infant mortality.[2] Our methodological kinship is direct—like them, we reject distance-based exposure and use a physical transport model to trace vessel-sourced pollution to its ultimate incidence, replacing their atmospheric aerosol transport model (CMAQ) with Lagrangian ocean-current tracking and their air-quality outcomes with coastal invasive-species outcomes. Where Hansen-Lewis and Marcus (2022) measure the domestic health incidence of a U.S. maritime regulation, we measure its transboundary ecological incidence—on the coastlines of countries that did not set the rule.

Section 2 presents the institutional background. Section 3 describes the data. Section 4 estimates the effect of exchange on introduced species at transit-country coastlines over 2007–2019, distinguishing the exchange era (2012–2015) from the treatment era (2016–2019). Section 5 decomposes the effect on community composition along three complementary dimensions. Section 6 concludes.

2 Background

The U.S. Coast Guard finalized its ballast-water management rule in 2012, establishing discharge standards for living organisms in ballast water released in U.S. waters.[3] The rule permits two compliance pathways—exchange (D1) and treatment (D2)—that generate distinct spatial predictions for introduced species risk. This section describes each margin and the empirical patterns they imply.

[2]In related work on vessel effects in California, Klotz and Berazneva (2022) also study the effect of air pollution, while Candau and Lafferrere (2024) analyze how vessel traffic affects California's underwater kelp forests.

[3]See the Federal Register notice "Standards for Living Organisms in Ships' Ballast Water Discharged in U.S. Waters" (2012): `https://www.federalregister.gov/documents/2012/03/23/2012-6579/standards-for-living-organisms-in-ships-ballast-water-discharged-in-us-waters`.

2.1 Exchange-based compliance (D1)

Under D1, vessels replace coastal ballast water with open-ocean water before discharging in port (typically targeting 95% tank-volume replacement; Miller et al., 2011). This reduces the density of coastal organisms in discharged ballast but necessarily entails mid-ocean discharge, so that some biological material is released along the vessel's approach route rather than at the destination port.

The spatial implication is that if exchange occurs within the Exclusive Economic Zones (EEZs) of transit countries, it can generate transboundary externalities. Figure 1 shows that U.S.-bound vessels routinely exchange ballast water within the EEZs of Caribbean and Central American countries. Descriptive evidence supports the relevance of this behavior; for several EEZs, U.S.-bound traffic represents a disproportionate share of the passing fleet, and exchange volumes shifted from U.S. waters toward transit-country EEZs after 2012 (Figures 7–8, Appendix A). The D1 analysis therefore estimates introduced alien species (IAS) outcomes along the coastlines of countries exposed to these exchange events.

A further complication is that exchange efficacy depends on environmental gradients. Exchange substantially reduces introduction risk when organisms face strong salinity or temperature shocks, but is less effective when gradients are weak; residual water and sediments in tanks are never fully eliminated (Gerhard et al., 2019). The net ecological effect of exchange is therefore an empirical question.

The transboundary externality depends not only on where vessels exchange but also on where the released organisms end up. Mid-ocean ballast discharge releases viable larvae, plankton, and invertebrates into the water column, where they drift with surface currents before potentially reaching coastal habitats. The Caribbean Current, Gulf Stream, and associated gyres can transport biological material hundreds of kilometers from the exchange location within weeks (Seebens et al., 2013; Lezama-Ochoa et al., 2022). Whether a given coastal area is exposed to ballast-water organisms therefore depends on the oceanographic connectivity between exchange locations and that coastline—not simply on geographic proximity.

Figure 1: BWE map in U.S. and transit countries

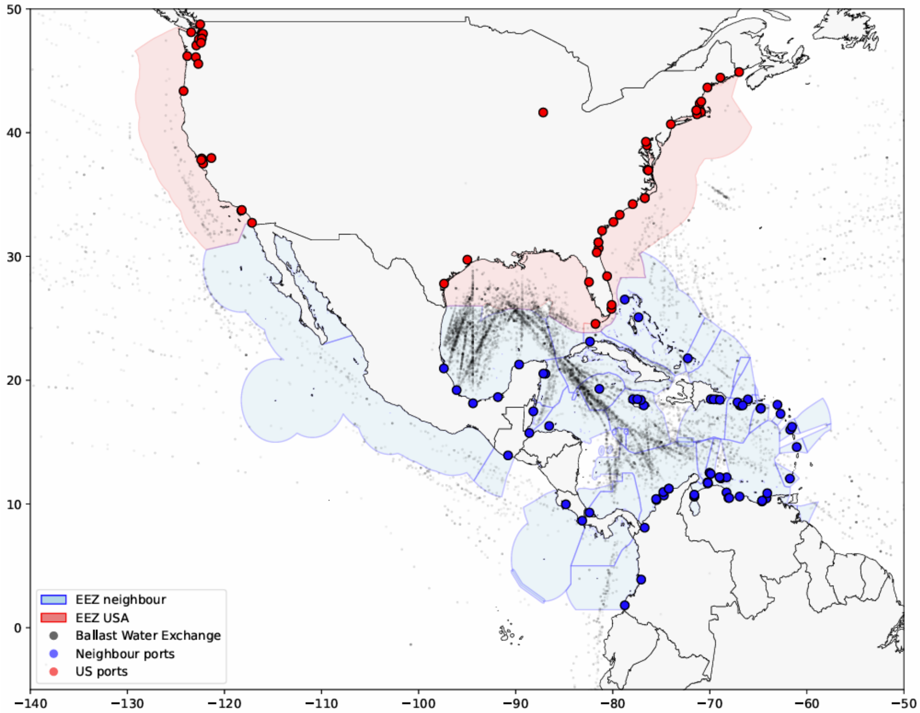

Note: Black dots show open-ocean ballast-water exchange (BWE) locations recorded by vessels bound for the United States, including events in international waters, U.S. coastal waters, and transit-country EEZs. U.S. ports are in red; transit-country ports in blue. Source: NBIC (defined in Section 3).

2.2 Treatment-based compliance (D2)

Under D2, ballast water must meet biological concentration limits before discharge, which in practice requires installing a ballast-water treatment system (BWTS).[4] Unlike exchange, treatment reduces biological pressure in the ballast stream itself, so that treated vessels discharge cleaner water at every subsequent port of call—generating potential positive ecological effects abroad.

BWTS adoption was effectively absent in U.S. arrivals before 2015 and then increased steadily. The IMO BWM Convention, which entered into force in September 2017, reinforced D2 incentives globally by tying BWTS installation to ships'

[4] 33 CFR § 151.2080; https://www.law.cornell.edu/cfr/text/33/151.2080.

certificate renewal schedules.[5]

Figure 2 shows the adoption trajectory, disaggregated by vessel route. For both Caribbean/Pacific-route vessels (those whose ballast-water exchange events fall predominantly inside our study-area bounding box) and Atlantic/Asia-Pacific-route vessels, adoption was essentially zero through 2015 and then rose steadily. Even by 2019, fewer than one in three vessels reported any effective BWTS use, and adoption on Caribbean/Pacific routes ran slightly behind adoption on Atlantic/Asia-Pacific routes. The gradual pace of the transition creates a tension in our empirical analysis. On the one hand, if adoption has been sufficient by 2019 to displace mid-ocean exchange along Caribbean/Pacific routes, the ecological externality at transit-country coastlines should attenuate after 2015. On the other hand, if adoption has been too slow to reduce exchange-based compliance on the routes traversing transit-country waters, mid-ocean exchange volumes above their EEZs should remain high—and the externality should persist or grow with trade volumes. Section 4 settles this question empirically.

[5]See IMO, "Implementing the BWM Convention": `https://www.imo.org/en/mediacentre/hottopics/pages/implementing-the-bwm-convention.aspx`.

Figure 2: BWTS adoption by vessel route, 2007–2019

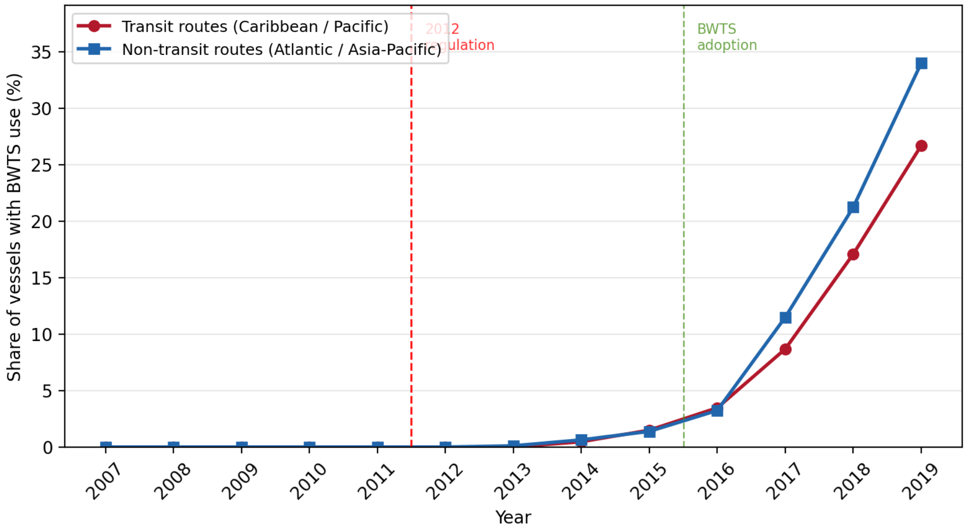

Note: Share of vessels with any effective BWTS use per year, among U.S. arrivals, split by vessel route class. A vessel is classified as "Caribbean/Pacific route" if more than 50% of its reported ballast-water exchange events fall inside the Caribbean / Pacific bounding box (0°–35°N, 120°–55°W); otherwise it is "Atlantic/Asia-Pacific route". This per-vessel classification is used only to disaggregate BWTS adoption in this figure and is distinct from our country-level sample of 39 transit countries defined in Section 3. Red dashed line: 2012 regulation. Green dashed line: start of noticeable BWTS adoption (2015). Source: NBIC tank-level ballast-water management reports.

3 Data

We combine three data sources. First, georeferenced occurrence records of introduced marine species from the Ocean Biodiversity Information System (OBIS), filtered using the World Register of Introduced Marine Species (WRiMS). Second, ship-level ballast-water management reports from the U.S. National Ballast Information Clearinghouse (NBIC), which collects and standardizes ballast-water management reporting forms submitted to the U.S. Coast Guard.[6] Third, a monthly

[6]NBIC was established in 1997 pursuant to the National Invasive Species Act of 1996. The stated aims include quantifying amounts and origins of ballast water discharged in U.S. coastal systems and documenting whether ballast water underwent open-ocean exchange or alternative treatment. A potential concern is that vessels might strategically misreport exchange volumes to appear compliant. This concern is mitigated by the USCG enforcement architecture: USCG inspectors con-

climatology of surface ocean currents built from the HYCOM GLBu0.08 reanalysis (Chassignet et al., 2007), which we use to simulate the drift of organisms released during mid-ocean ballast-water exchange (Section 4.1).

All spatial operations (buffers and spatial joins) are performed using geographic coordinates in a common reference system. Port locations are harmonized using standardized port names; ambiguous cases are resolved by manual inspection.

3.1 Dependent Variables

Introduced marine species. Our species outcomes are drawn from OBIS, filtered to species flagged as introduced in WRiMS. WRiMS is a thematic sub-register of the World Register of Marine Species (WoRMS) that records species introduced beyond their native range by human activities, excluding natural range expansions.[7] WRiMS covers all introduced marine species, including but not limited to species confirmed as ecologically harmful. We follow the convention in the invasion ecology literature and refer to these species as IAS throughout, noting that our estimates capture changes in introduced-species occurrence broadly rather than ecological damage specifically. We download occurrence records from the OBIS GeoParquet archive on AWS S3, retaining all georeferenced records with the `wrims=TRUE` flag for marine and brackish species. This yields a global panel of introduced-species occurrences linked to WoRMS taxonomy, ensuring nomenclatural consistency across time and data providers. Like all global occurrence compilations, the data reflect a combination of ecological processes and observation effort. Records are subject to geographic and temporal reporting heterogeneity. We therefore interpret estimates as relationships in observed occurrences rather than a complete census of introductions.

We construct three outcomes that capture complementary dimensions of invasion pressure at the segment level.[8] *Occurrence* is the total number of records of

duct regular boardings and can verify reported volumes against onboard ballast-water logs, tank gauges, and GPS-based position records. False reporting carries civil penalties of up to $35,000 per day per violation (33 CFR § 151.2080), and deliberate falsification constitutes a criminal offense under 18 U.S.C. § 1001. The USCG's National Vessel Documentation Center cross-references NBIC reports with independent voyage data, making systematic falsification difficult to sustain undetected. We therefore treat the NBIC reports as an accurate record of the management method declared at the time of port entry.

[7] WRiMS is accessible through the WoRMS platform (`marinespecies.org/introduced`). OBIS records carrying the `wrims=TRUE` flag are retained; purely natural colonizations are excluded by construction.

[8] See Appendix A, Table 3, for summary statistics.

introduced species s at segment i in year t. It measures the stock of established introduced species—the cumulative footprint of past introductions. *Total abundance* is the sum of occurrence counts across species at each segment-year, a scale measure of aggregate invasive pressure. *Shannon diversity* H' is computed over all introduced species detected at the segment-year and combines the number of species present with their relative abundances. A high H' indicates that many species contribute equally; a low H' indicates dominance by a few species.[9]

Together, the three outcomes characterize changes in the stock (occurrence), the aggregate count (abundance), and the community structure (Shannon) of introduced species. Two scope limitations of this outcome definition should be kept in mind throughout. First, all outcomes are built from recorded species observations; they measure detection rather than biomass, establishment, or ecological damage, and the coefficients should not be read as welfare-scale damage estimates. Second, because the species-level panel contains only segment-species-years where at least one occurrence is recorded, the occurrence and abundance coefficients measure the intensive margin (detection intensity given presence) rather than the extensive margin (presence versus absence). Shannon diversity and the composition indicators defined in Section 5 are segment-year outcomes that do include genuine zeros and therefore carry both margins.

3.2 From Mid-Ocean Exchange to Coastal Exposure

For transit countries, the central empirical challenge is connecting mid-ocean ballast-water exchange to coastal ecological outcomes. NBIC reports tell us exactly where and when each exchange event occurs—GPS coordinates, date, and volume—but the organisms released there do not stay put. Surface currents can carry them in very different directions. Assigning exchange to the nearest port or the nearest country would therefore misattribute exposure, because it ignores the physical transport that determines where organisms actually arrive.

We define as *transit countries* the 39 sovereign states and territories whose Exclusive Economic Zones lie on the maritime approach routes to U.S. Atlantic, Gulf, and Pacific ports — all Caribbean island states, the Caribbean coasts of Central America and northern South America, and the Pacific coasts of Central America plus Ecuador. We discretize their coastlines into approximately 2,400 coastal segments. Each segment is a 25 km arc of coastline, buffered 100 km seaward and clipped to the country's Exclusive Economic Zone. We measure both exposure

[9] Because WRiMS is a register of introduced coastal and benthic species, the index does not pick up mid-ocean pelagic drifters; Shannon therefore reflects shifts among established invaders, not new pelagic arrivals.

and invasive-species outcomes at this level.

On the exposure side, we trace where organisms released during mid-ocean exchange are carried by ocean currents, using Lagrangian particle tracking, a standard method in marine ecology (Lezama-Ochoa et al., 2022; Dobrzycka-Krahel et al., 2024). We release virtual particles at the coordinates of each NBIC-reported BWE event and assign each particle's final position to its nearest segment. On the outcome side, we assign WRiMS occurrence records directly to segments by the same nearest-segment spatial join. We exclude records within the U.S. Exclusive Economic Zone (VLIZ Maritime Boundaries v12) and within 30 km of the U.S. coast to prevent contamination of the transit-country sample.[10]

For each BWE event, we release 100 virtual particles at the recorded exchange coordinates with small spatial perturbations (± 0.1 degrees) and advect them for 60 days using a monthly climatology of surface currents built from the HYCOM GLBu0.08 reanalysis (Chassignet et al., 2007). The advection follows a forward-Euler scheme with 6-hour time steps and stochastic sub-grid diffusion to capture turbulent dispersion (Lezama-Ochoa et al., 2022). Because we use a monthly climatology rather than a year-specific reanalysis, inter-annual variability in surface currents does not enter the treatment variable; all time variation in BWE_{it}^{drift} comes from temporal variation in upstream ballast-water exchange events, not from temporal variation in how currents route them.

Three parameter choices deserve discussion. First, the 100 particles per event is a computational simplification of a much larger biological reality. A typical bulk carrier discharges 30,000–50,000 m^3 of ballast water containing on the order of 1,000 viable zooplankton per m^3 (Bradie et al., 2018), implying tens of millions of organisms per discharge. After exchange removes 80–95% of original organisms, each exchange event still releases on the order of 1.5 to 10 million residual organisms into the water column. We simulate 100 particles per exchange event (GPS coordinate). This implies a survival-to-arrival rate of roughly 1 in 15,000 to 1 in 100,000—meaning we assume that only 0.001–0.007% of released organisms survive the full drift and reach a coastal segment. For context, marine larval mortality rates are typically 90–99% over a few weeks, as most larvae are consumed by predators or starve. A 0.001% survival rate over 60 days of open-ocean drift is biologically plausible and extremely conservative. The resulting treatment variable should therefore be understood as a lower bound of the actual biological expo-

[10]The EEZ boundary starts at the 12 nautical-mile territorial sea baseline, leaving a gap between the coastline and the EEZ polygon that the 30 km buffer closes. Without this buffer, approximately 27,000 U.S. near-shore records survive the EEZ filter. The 30 km threshold covers the 22 km (12 NM) gap with a safety margin and does not reach any transit-country territory; the nearest (Bahamas) is approximately 80 km from the U.S. coast. We also filter out deep-ocean observations (bathymetry > 200 m).

sure. Second, the 60-day advection horizon is calibrated to the survival biology of ballast-water organisms. Ballast water carries a wide range of taxa—bacteria, phytoplankton, zooplankton, invertebrate larvae, and resting stages—of which planktonic larvae are among the most sensitive to drift duration. Most marine invertebrate larvae remain planktonic for 2–4 weeks, with planktotrophic species surviving up to two months (Shanks, 2009). More resistant forms, such as dinoflagellate cysts and diatom resting spores, can remain viable for years in ballast-tank sediments (Hallegraeff and Bolch, 1992). A 60-day horizon therefore captures the drift window for the majority of viable organisms, while remaining conservative relative to the most resilient taxa. At Caribbean Current speeds of 0.3–0.5 m/s (Richardson, 2005), 60 days covers 1,500–2,600 km, sufficient to traverse most of the Caribbean basin. Third, the sub-grid diffusion coefficient $K_h = 50$ m^2/s is at the upper end of values used in prior ballast-water studies (10 m^2/s in Lezama-Ochoa et al., 2022 and Dobrzycka-Krahel et al., 2024), but consistent with the Okubo scaling law at the 25–100 km resolution of our coastal segments (Okubo, 1971; Poje et al., 2014). The higher diffusion spreads particles more broadly across segments, making the treatment variable noisier and biasing estimates toward zero; our specification is therefore conservative with respect to this parameter.

The treatment variable for segment i in year t is the drift-weighted BWE volume:

$$BWE_{it}^{\text{drift}} = \sum_{e \in \mathcal{E}_t} w_{ei} \times \text{Volume}_e, \tag{1}$$

where w_{ei} is the fraction of particles from event e that reach segment i, Volume_e is the exchanged volume in metric tons, and $\mathcal{E}_t$ is the set of BWE events in year t. The variable is expressed in thousands of metric tons in the estimation tables. As a specification check, we compare drift exposure to a naive distance-based alternative that sums BWE volumes within a 300 km radius of each segment centroid, ignoring ocean currents. The naive measure produces insignificant estimates, supporting the current-based assignment (Appendix B).

Monitoring effort and the detection problem. A segment with zero IAS records poses an interpretation problem. Does it reflect genuine absence of invasive species, or simply the absence of marine biologists surveying there? WRiMS occurrence records exist only where scientists have collected samples, and monitoring effort is unevenly distributed across Caribbean and Pacific coastlines. If unmonitored segments are treated as zeros, the estimation confounds ecological processes with observation gaps.

To distinguish genuine zeros from unmonitored segments, we download the full OBIS database for the study region—3.75 million georeferenced marine and brackish species records (not limited to WRiMS introduced species)—and assign them to segments using the same spatial procedure. For each segment-year, we compute n_{it}^{native}, the count of non-invasive (non-WRiMS) species records. If $n_{it}^{native} > 0$, scientists surveyed the segment and recorded native species, so the absence of WRiMS records there is informative. A segment with $n_{it}^{native} = 0$ is unmonitored and excluded from the estimation sample. This procedure yields a panel of monitored segment-years, with genuine ecological zeros for segments where monitoring detected no invasive species. The variable $\log(n_{it}^{native})$ is also available as an explicit monitoring-intensity control; its role in the estimation is discussed in Section 4.1.

4 Exchange-Based Compliance and Transit-Country Coastlines

4.1 A Continuous Difference-in-Differences Design

Our goal is to estimate whether coastal segments with more drift exposure experience larger changes in introduced-species outcomes after the 2012 U.S. Ballast Water Management Act. We also test whether the externality persists as the fleet transitions from exchange to onboard treatment after 2015. Figure 1 shows the spatial distribution of BWE events across 39 transit countries; Figure 3 presents the directed acyclic graph (DAG) summarizing the causal structure of the design, with each back-door path from drift exposure to invasive species explicitly labeled by the device used to block it.

Figure 3: Causal Structure of the Identification Strategy

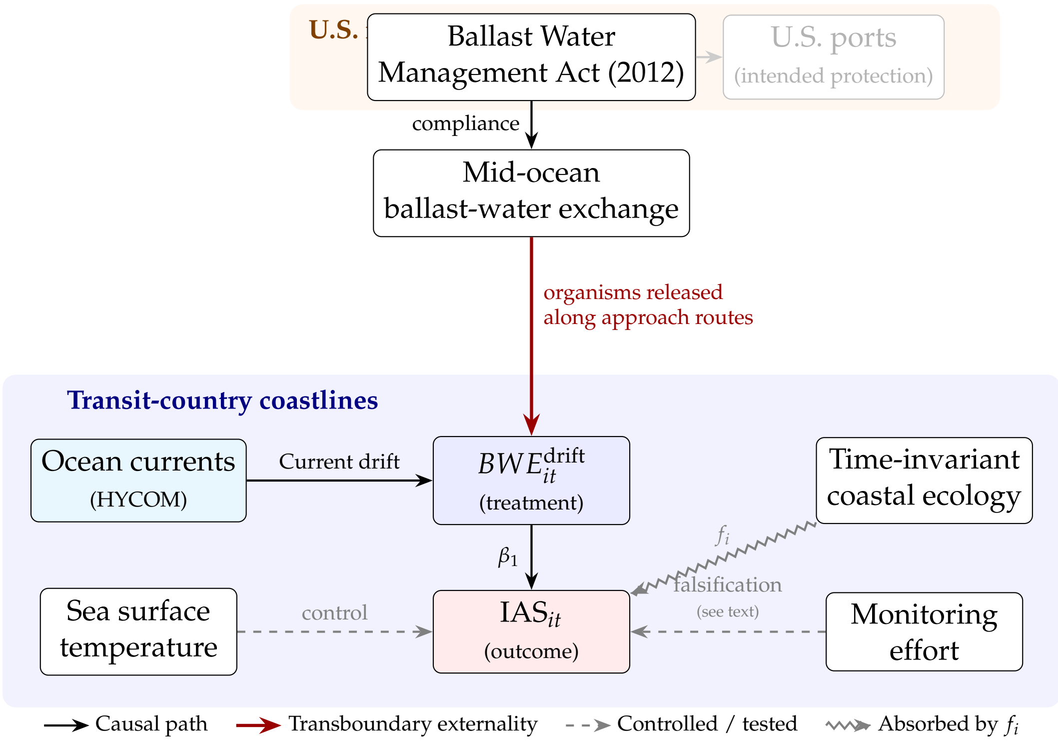

Note: Solid arrows represent causal paths. Dashed arrows indicate confounders controlled by covariates or assessed via falsification tests. Zigzag arrows denote paths absorbed by segment fixed effects.

The design is a continuous difference-in-differences (Callaway et al., 2024). It compares within-segment changes in invasive species outcomes across coastal locations that received different volumes of current-delivered ballast-water exchange. Two segments at similar latitudes, with similar baseline ecology, may be exposed to very different volumes of drifting organisms depending on where vessels chose to exchange and how ocean currents redistributed the discharge. After 2012, when the U.S. Ballast Water Management Act required all U.S.-bound vessels to exchange coastal ballast in the open ocean, segments downstream of intensified exchange activity should experience larger ecological changes than segments outside the drift corridors—if the regulation generates a transboundary externality.

We estimate the following PPML over 2007–2019:

$$\mathbb{E}[\mathrm{IAS}_{sit} \mid X_{sit}] = \exp\Big(\beta_1\,\mathrm{Exchange}_t \times \sqrt{BWE^{\text{drift}}_{it}} + \beta_2\,\mathrm{Treatment}_t \times \sqrt{BWE^{\text{drift}}_{it}} + \beta_3\sqrt{BWE^{\text{drift}}_{it}} + Z'_{it}\gamma + f_i + f_t + f_s\Big), \tag{2}$$

where i indexes coastal segments, t indexes years, and s indexes species (the species index is dropped for Shannon diversity and segment-year outcomes). $\mathrm{Exchange}_t = \mathbb{1}\{2012 \leq t \leq 2015\}$ captures the period of full exchange-based compliance, and $\mathrm{Treatment}_t = \mathbb{1}\{t \geq 2016\}$ captures the transition period 2016–2019, during which U.S.-bound vessels increasingly adopted onboard ballast-water treatment systems (Figure 2). BWE^{drift}_{it} is the drift-weighted volume of ballast-water exchange reaching segment i in year t (Equation 1), entered as a square root because the underlying distribution is heavily right-skewed (max approximately 980 times the median at the segment-year level).[11] We refer to $\sqrt{BWE^{\text{drift}}_{it}}$ as *drift exposure* in subsequent prose. Z_{it} is segment-year sea surface temperature, the only time-varying control. Segment fixed effects f_i block the back-door path through time-invariant heterogeneity in coastal ecology, monitoring infrastructure, and geographic exposure to shipping lanes (zigzag path in Figure 3); year fixed effects f_t absorb common annual shocks; species fixed effects f_s absorb cross-species differences in baseline detection. Standard errors are clustered at the country $\times$ basin level, which refines the country partition by splitting the seven countries whose Pacific and Atlantic/Caribbean coastlines are physically disconnected by the Central American isthmus.[12]

The coefficients β_1 and β_2 measure the change in the within-segment relationship between drift exposure and introduced-species outcomes during the ex-

[11] A Box-Cox analysis points to $\lambda = 0.5$. The square-root transformation compresses the upper tail and reduces the max-to-median ratio to approximately 30, preserving variation across the bulk of the distribution (Table 3).

[12] The seven countries are Mexico, Guatemala, Honduras, Nicaragua, Costa Rica, Panama, and Colombia. Ocean currents cannot cross land, so shocks on one side of the isthmus are uncorrelated with shocks on the other; clustering these countries as one unit would pool two independent sub-samples. The refinement follows the design-based clustering criterion of Abadie et al. (2023). The top five country-basin clusters account for 81.4% of the variance in $\sqrt{BWE^{\text{drift}}}$ and the effective number of clusters in the sense of Carter et al. (2017) is $G_{\text{eff}} = 4.7$ against a nominal $G = 34$, placing the data in the few-treated-clusters regime of MacKinnon and Webb (2020) in which wild cluster bootstrap is unreliable. We therefore complement the analytical cluster-robust standard errors with the cluster-jackknife CR3 correction of MacKinnon et al. (2023), which is robust to leverage imbalance. CR3 standard errors appear bracketed below each coefficient in the main-text tables; cluster counts are reported in the table bodies.

change era and the treatment era, respectively, relative to the pre-2012 baseline. Splitting the post-policy period allows the externality to differ between the years of full exchange compliance and the years during which the fleet transitioned to onboard treatment.

Estimation. We estimate all count-valued outcomes—here, occurrence, total abundance, and Shannon diversity—by Poisson Pseudo-Maximum Likelihood (PPML), which is consistent under arbitrary conditional variance structures and does not require the Poisson or negative binomial distribution to hold (Santos Silva and Tenreyro, 2006). The choice is driven by strong overdispersion in the outcomes. For instance, occurrence at the segment-species-year level is extremely right-skewed (median 5, maximum 7,100) with a pooled variance-to-mean ratio of 2,613. A full diagnostic—distribution, variance-mean relationship, and a discussion of PPML versus negative binomial—is provided in Appendix A (Figure 9).

Identification. The design rests on two identifying assumptions. Together with the fixed effects in Equation 2 and the SST control, they close every back-door path from BWE^{drift} to IAS in the DAG of Figure 3 and support a causal interpretation of β_1 and β_2.

The first is a *parallel trends assumption*: absent the 2012 policy, the within-segment relationship between drift exposure and introduced-species outcomes would have remained stable over time. The parallel-trends assumption is not directly testable, but the pre-2012 event-study coefficients are flat and indistinguishable from zero. Substantively, U.S.-bound vessels were not required to exchange mid-ocean before 2012, so there was no mechanism by which segments downstream of where exchange would later concentrate could have been differentially affected in anticipation of the rule.

The second is a *current transport exogeneity assumption*: the ocean currents that route organisms from exchange locations to coastal segments are not driven by factors that also affect introduced-species outcomes. Surface currents in the Caribbean and Pacific are governed by climatological physics, independent of any economic or ecological process at the segments they reach. The HYCOM reanalysis we use is a physical model fit to satellite altimetry and in situ observations, with no input from biological, regulatory, or shipping data. Vessels choose where to exchange; they cannot choose where released organisms drift.

Threats to identification

We discuss the three main threats to identification here; two further concerns (vessel anticipation and BWE-location exogeneity under a monthly current climatology) are addressed in Appendix B.

Port-level confounders. Segments exposed to more mid-ocean drift may also be closer to ports, and port-related pressures—hull fouling, cargo organisms, port pollution—could bias the estimated drift effect. Segment fixed effects absorb any time-invariant correlation between drift exposure and port proximity, and year fixed effects absorb common time trends in maritime traffic. For port-level pressure to bias β_1 or β_2, it would have to change differentially at high-drift segments precisely in 2012 and 2015. No parallel port regulation or pollution event is known to have coincided with these dates. As an additional check, we re-estimate Equation 2 on the subsample of segments with no port within 30 km and interact the treatment with a port-presence indicator in the full sample (Table 9, Appendix B).

Differential monitoring. Monitoring effort may change differentially at segments more exposed to drifting organisms, inflating the relationship between drift and recorded occurrences through surveillance rather than ecological change. Segment and year fixed effects absorb persistent and common-trend variation. The residual threat is a differential *change* in monitoring intensity at high-drift segments coinciding with 2012 or 2015, for which there is no natural channel: the 2012 U.S. rule was finalized by the Coast Guard with limited prior notice, and the IMO Ballast Water Management Convention did not enter into force until 2017. We test this directly with a falsification regression of monitoring intensity (n_{it}^{native}) on era-interacted drift exposure; neither coefficient is distinguishable from zero. As additional robustness, we include $\log(n_{it}^{native})$ as an explicit monitoring-intensity control in the three headline specifications (Table 10, Appendix B); all six era coefficients retain their sign and their statistical significance.

Overlap with the IMO BWM Convention. The IMO Ballast Water Management Convention entered into force in September 2017, inside our treatment era. Four features of the setting make Convention-driven contamination of β_2 unlikely. First, the IMO D-2 discharge standard is less stringent than the U.S. type-approval regime (Čampara et al., 2019), so a vessel already complying with the 2012 U.S. rule faced no additional obligation. Second, BWE_{it}^{drift} is constructed from NBIC reports of U.S.-bound voyages only. Convention-driven BWTS adoption by non-U.S. traffic therefore does not enter the treatment. Even for vessels serving both U.S. and non-U.S. routes, switching to BWTS on U.S.-bound legs would reduce BWE_{it}^{drift}, not inflate it. Third, the Convention had not entered into force in many of our 39 transit countries during 2017–2019 (including Belize, Colombia, Costa Rica, Cuba, Guatemala, Haiti, Nicaragua, and Venezuela), and even ratifying parties faced documented port-state-control and type-approval problems through 2020. Fourth, the event-study coefficients are flat from 2016 through 2019 rather than jumping at 2017, inconsistent with a step-wise Convention effect.

4.2 Results

Occurrence increases significantly in both the exchange era and the treatment era (Column 1, Table 1). Total abundance rises in both eras, with the treatment-era effect larger than the exchange-era one (Column 2). Shannon diversity also increases in both eras, again with a larger treatment-era effect (Column 3).

Across all three outcomes, the pattern is consistent. Downstream coastlines record invasive species more often (occurrence), carry a larger total invasive load (abundance), and host a more diverse pool of invaders (Shannon). The externality is therefore not a marginal species-specific phenomenon but a measurable shift in the ecological composition of transit-country coastlines.

Table 1: Effect of Exchange on Invasive Species at Transit-Country Coastlines

	(1) Occurrence	(2) Abundance	(3) Shannon H′
Exchange era (2012–15) × $\sqrt{BWE^{drift}}$	0.1335***	0.1071***	0.0782**
	(0.0414)	(0.0400)	(0.0346)
Jackknife SE	[(0.0248)]	[(0.0465)]	[(0.0619)]
Treatment era (2016–19) × $\sqrt{BWE^{drift}}$	0.1903***	0.2461***	0.1194***
	(0.0261)	(0.0329)	(0.0304)
Jackknife SE	[(0.0374)]	[(0.0664)]	[(0.0518)]
Observations	5,211	724	623
Mean dep. var.	60.886	386.570	0.984
Clusters	32	32	32
Controls (SST, $\sqrt{BWE^{drift}}$ level)	✓	✓	✓
Segment FE	✓	✓	✓
Year FE	✓	✓	✓
Species FE	✓		

Notes: PPML. Treatment: Era × $\sqrt{BWE^{\text{drift}}}$ (square root of drift-weighted BWE in thousands MT). Column (1): occurrence (segment-species-year level) with segment, year, and species FE. Column (2): total abundance = sum of occurrence counts (segment-year level). Column (3): Shannon diversity H′ (segment-year level). Controls (sea surface temperature and $\sqrt{BWE^{\text{drift}}}$ level) are included in every column; their coefficients are reported in Table 4 in Appendix B. Standard errors clustered at the country × basin level (see Section 4); cluster counts reported in the table body. Bracketed cells report cluster-jackknife (CR3) standard errors (MacKinnon et al., 2023). Singleton observations dropped (Correia, 2015). Time period: 2007–2019. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The event study in Figure 4 traces the dynamic effect across the three outcomes (occurrence, new arrivals, and Shannon diversity). All three show pre-2012 coefficients that are not significantly different from zero — supporting the parallel-trends assumption — and rise sharply after 2012, remaining positive across the treatment era.

Figure 4: Event Study: Transit-Country Coastlines (2007–2019)

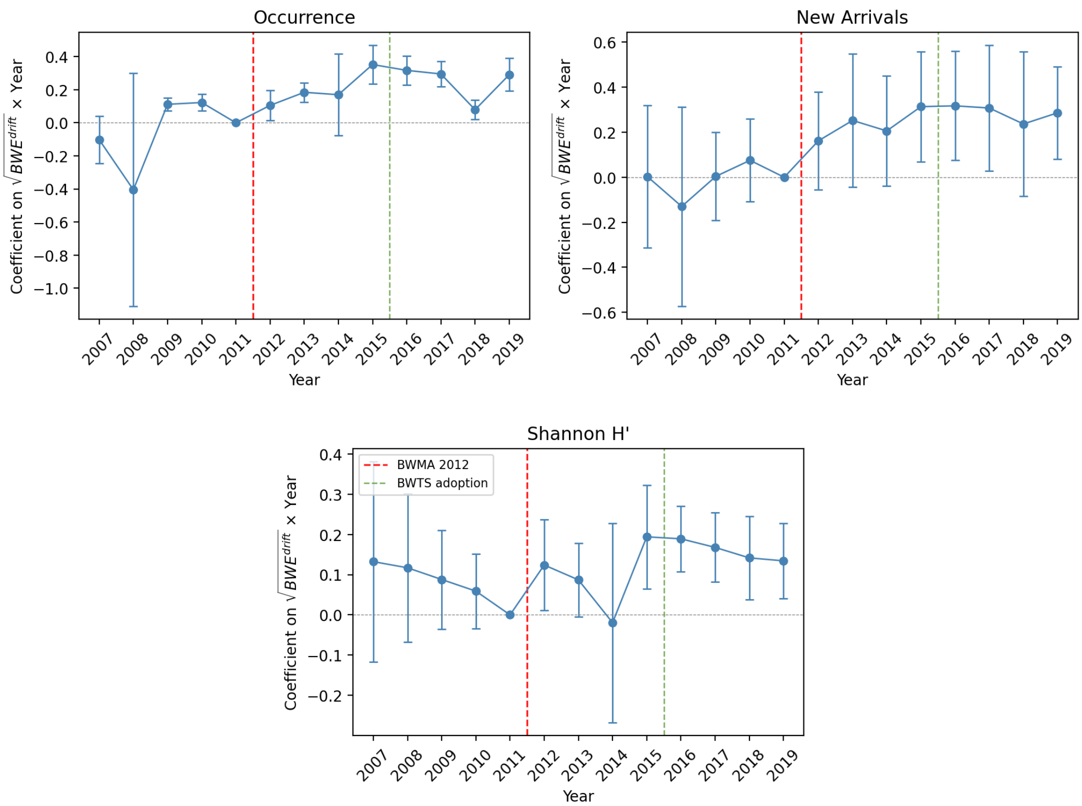

Note: PPML event-study estimates of $\mathbb{1}\{t = \tau\} \times \sqrt{BWE^{\text{drift}}}$, 2007–2019, with 2011 as the reference year. Top-left: occurrence (segment-species-year, segment + year + species FE). Top-right: new arrivals (segment-year, segment + year FE). Bottom-center: Shannon diversity H′ (segment-year, segment + year FE). Controls: $\sqrt{BWE^{\text{drift}}}$ level, SST. 95% confidence intervals, clustered at the country × basin level. Red line: 2012 regulation. Green line: beginning of BWTS adoption (2015).

4.3 Counterfactual

The coefficients in Table 1 are semi-elasticities; they are informative but do not directly answer the aggregate question: how many invasive species records at transit-country coastlines are attributable to the 2012 rule's shift in the drift channel? To translate the estimates into an aggregate, we compute, for each post-2012 transit-country segment-year, a counterfactual total abundance[13] in which the 2012 rule does not shift the drift–abundance relationship. In practice, we multiply the observed total abundance by $\exp(-\hat{\beta}_{\text{era}}\sqrt{BWE^{\text{drift}}})$, which nets out the era-interaction effect. This is the standard difference-in-differences counterfactual: it holds the pre-policy within-segment relationship between drift exposure and abundance fixed and asks how many post-2012 records would not have occurred had the policy-induced era shifts been zero.[14]

Between 2012 and 2019, we observe 241,315 total invasive species records at the 503 transit-country segment-years in our sample. Absent the 2012 rule's shift in the drift channel, the simulation implies that roughly 43,443 of these records—18.0% of the observed total (95% CI: 13.6%–21.5%)—would not have been observed. On average, about one in six of the invasive species records detected at transit-country coastlines during the post-policy period can be attributed to the policy-induced shift in the drift channel.

4.4 Why are treatment-era coefficients larger?

The treatment-era coefficients in Table 1 are larger than their exchange-era counterparts. This is a puzzle. The treatment era is when BWTS were being installed, so we might have expected exchange-based compliance—and the externality—to diminish. The opposite occurred because adoption was too slow. Fewer than 30% of Caribbean/Pacific-route vessels had operational BWTS by 2019 (Figure 2), and BWE volumes in the Caribbean and eastern Pacific continued to grow, from 48.5 million MT in 2015 to 61.3 million MT in 2019 (Figure 5).

[13]Total abundance is the natural unit for counting WRiMS records at the segment-year level. By construction, it is the sum of the species-level occurrence counts in Column 1 of Table 1, so the two outcomes carry the same underlying information and differ only in the aggregation level.

[14]The counterfactual does not require a stance on the pre-policy relationship between drift exposure and abundance. The difference between observed and counterfactual abundance, aggregated across segments and years, is the excess reported below. A 95% confidence interval is obtained by parametric bootstrap of the two era-interaction coefficients using the cluster-robust covariance matrix from the main specification (2,000 draws).

Figure 5: BWE Volumes in the Caribbean (2007–2019)

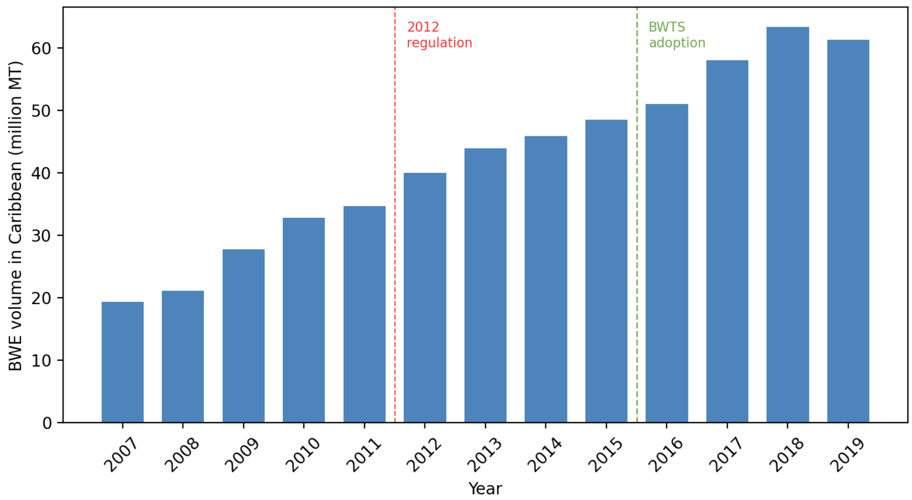

Note: Total ballast-water exchange volume (million MT) by U.S.-bound vessels exchanging inside the Caribbean / Pacific bounding box (0°–35°N, 120°–55°W), summed across all tank events per year. Red dashed line: 2012 regulation. Green dashed line: start of noticeable BWTS adoption (2015). Source: NBIC tank-level ballast-water management reports.

This resolves the tension described in Section 2 in favor of the second hypothesis. BWTS adoption through 2019 was insufficient to displace exchange-based compliance on the routes traversing transit-country waters.

4.5 Robustness

The Lagrangian treatment assignment is assessed by comparing it to a naive distance-based measure. When BWE^{drift} is replaced with the simple sum of BWE volumes within a 300 km radius of each segment's centroid (ignoring ocean currents), the occurrence coefficient drops to near zero (Table 5, Appendix B), suggesting that ocean-current transport, not geographic proximity, captures the relevant exposure channel. A leave-one-country-out diagnostic confirms that no single country drives the result for either occurrence (Table 7) or Shannon diversity (Table 8, both in Appendix B). Adding $\log(n^{native})$ as an explicit monitoring-intensity control preserves the sign and significance of all six era coefficients (Table 10, Appendix B). Interacting the treatment with a port-presence indicator reveals that the effect differs by segment type; details are in Appendix B (Table 9). Results are also robust to alternative advection horizons and diffusion coefficients: shortening the drift horizon from 60 to 50 days attenuates the coefficients but preserves the sign and statistical significance of all six era interactions, and lowering the sub-grid diffusion from $K_h = 50$ to $K_h = 30$ m^2/s leaves the estimates essentially unchanged (Table 6, Appendix B).

Two additional checks close the loop on the identification strategy. They serve different purposes and should be read separately. The first is a *randomization-inference* check — an exact inference procedure, not a placebo. Within each year, we permute drift exposure across segments, destroying any segment-specific spatial pattern while holding yearly totals fixed, and re-estimate Equation 2 $B = 1000$ times. Fisher permutation p-values for all six era coefficients are below one percent (Panel A, Table 11), sharper than the cluster-robust p-values in Table 1: the observed coefficients sit in the extreme tail of the permutation null, confirming that the segment-level spatial pattern of drift exposure carries the signal and that aggregate yearly BWE volume alone cannot reproduce the result (following Young, 2019). The second check is a genuine *placebo*: we re-run the full Lagrangian pipeline on a rotated HYCOM current field, negating the u and v components so that current magnitudes are preserved but every current vector points in the opposite direction. Under this rotation, particles released from observed BWE events drift away from the segments they reach in the data, breaking the physical routing from exchange locations to coastal exposure while preserving the particle-tracking machinery and the distribution of event volumes. If the drift channel is what identifies the externality, the rotated-field coefficients should collapse to near zero. Panel B of Table 11 reports the result. Five of the six era-interaction coefficients are indistinguishable from zero ($p > 0.25$ throughout), and the sixth (abundance, treatment era) is negative with $p = 0.10$, opposite in sign to its headline counterpart. The rotation therefore breaks the result entirely, confirming that the signal in Table 1 depends on the physical drift channel rather than on the particle-tracking machinery or the distribution of BWE volumes.

5 Does exchange diversify the invasive community?

The three outcomes in Section 4 measure the *magnitude* of invasive pressure: how many species, how abundant, how diverse. They do not tell us *which* species change. A rise in Shannon diversity could reflect the arrival of a few new species that displace the previously dominant assemblage, or a broader restructuring in which many species each gain a share. Distinguishing these two possibilities matters for the policy interpretation. If mid-ocean exchange moves a handful of rare species into a few segments, the ecological footprint is narrow and potentially manageable. If it diversifies the entire community—more species arriving, more of them establishing at meaningful levels, and single-species dominance breaking down—the footprint is broad and systemic. This section addresses the question directly.

5.1 Three dimensions of community change

We characterize composition change along three complementary axes. *Species richness* is the count of distinct WRiMS-listed species observed at segment i in year t. A rise in richness means the treated coastline is exposed to a wider pool of introduced species. Richness alone is noisy, however. It treats a species detected once as equivalent to one detected in large numbers. We complement it with the *number of common species*, defined as the count of species whose share of total segment-year abundance is at least 5%. This indicator measures how many species reach ecologically meaningful presence rather than appearing as transient or marginal detections. Finally, we use the *Berger-Parker dominance index*, the share of total abundance held by the single most abundant species:

$$BP_{it} = \max_{s} \frac{y_{s,it}}{\sum_{s'} y_{s',it}},$$

where $y_{s,it}$ is the WRiMS occurrence count of species s at segment i in year t. A high Berger-Parker means a single species monopolizes the invasive community; a decline means the concentration is broken as the community becomes more evenly distributed across species.

The three indicators answer three different questions. Richness asks whether *more* species arrive, the common-species count asks whether *the new arrivals matter*, and Berger-Parker asks whether *established dominance is broken*. A coherent diversification story requires movement in all three—richness up, common species up, Berger-Parker down—and no single indicator is decisive in isolation.

Richness and the number of common species are counts and are estimated by PPML. Berger-Parker is a proportion bounded in $[0, 1]$ and is estimated by OLS. All three specifications use segment and year fixed effects and the same era-interacted drift exposure as Equation 2, with sea surface temperature as the sole time-varying control.

5.2 Results

Table 2 reports the estimates. Species richness increases significantly during the treatment era (Column 1). Drift exposure introduces new species to transit-country coastlines. The number of common species—those holding at least 5% of total abundance at a segment—increases significantly in both eras (Column 2). The new arrivals do not remain rare but reach ecologically meaningful presence. The Berger-Parker dominance index decreases significantly in both eras (Column 3). The maximum species share falls, meaning no single species monopolizes the invasive community.

The pattern is consistent with mid-ocean exchange delivering a varied assemblage of coastal organisms from origin-port ecosystems, whose species composition differs from the locally dominant assemblage at transit-country coastlines.

Table 2: Community Composition at Transit-Country Coastlines

	(1) Richness	(2) N common	(3) Berger-Parker
Exchange era (2012–15) × $\sqrt{BWE^{drift}}$	0.0410	0.0539***	-0.0182***
	(0.0355)	(0.0165)	(0.0044)
Jackknife SE	[(0.0615)]	[(0.0222)]	[(0.0076)]
Treatment era (2016–19) × $\sqrt{BWE^{drift}}$	0.1144***	0.0685***	-0.0292***
	(0.0391)	(0.0232)	(0.0072)
Jackknife SE	[(0.0726)]	[(0.0272)]	[(0.0106)]
Observations	724	724	724
Clusters	32	32	32
Controls (SST, $\sqrt{BWE^{drift}}$ level)	✓	✓	✓
Segment FE	✓	✓	✓
Year FE	✓	✓	✓

Notes: Segment-year level. Column (1): richness = count of distinct invasive species (PPML). Column (2): number of common species = species with $\geq$ 5% share of total abundance (PPML). Column (3): Berger-Parker = maximum species share, a dominance index where higher values indicate greater single-species concentration (OLS). Controls (sea surface temperature and $\sqrt{BWE^{\mathrm{drift}}}$ level) are included in every column; their coefficients are reported in Table 4 in Appendix B. Segment and year FE. Standard errors clustered at the country × basin level; bracketed cells report cluster-jackknife (CR3) standard errors (MacKinnon et al., 2023). Singleton observations dropped (Correia, 2015). Time period: 2007–2019. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 6 tracks the three indicators year by year. The event-study specification replaces the era dummies with year-specific interactions of drift exposure, using 2011 as the reference year. The richness panel (top-left) has pre-period point estimates close to zero with confidence intervals covering zero, and rises after the start of visible BWTS adoption in 2015; every post-2015 coefficient is statistically significant and the level remains stable through 2019. The number of common species (top-right) has a flat pre-period, peaks in 2015, and attenuates toward zero by 2018–2019 while remaining directionally positive. This attenuation is consistent with the continuing Berger-Parker decline. As single-species dominance breaks down, abundance becomes more evenly distributed across species and fewer cross the 5% share threshold. The Berger-Parker panel (bottom-center) is the mirror image. Coefficients are flat before 2012, turn negative right after the policy, and remain stably negative through 2019. The three panels depict a persistent restructuring of the invasive community at treated segments rather than a one-off spike. More species arrive, more of them become common, and no single species dominates.

Figure 6: Event Study: Composition Indicators (2007–2019)

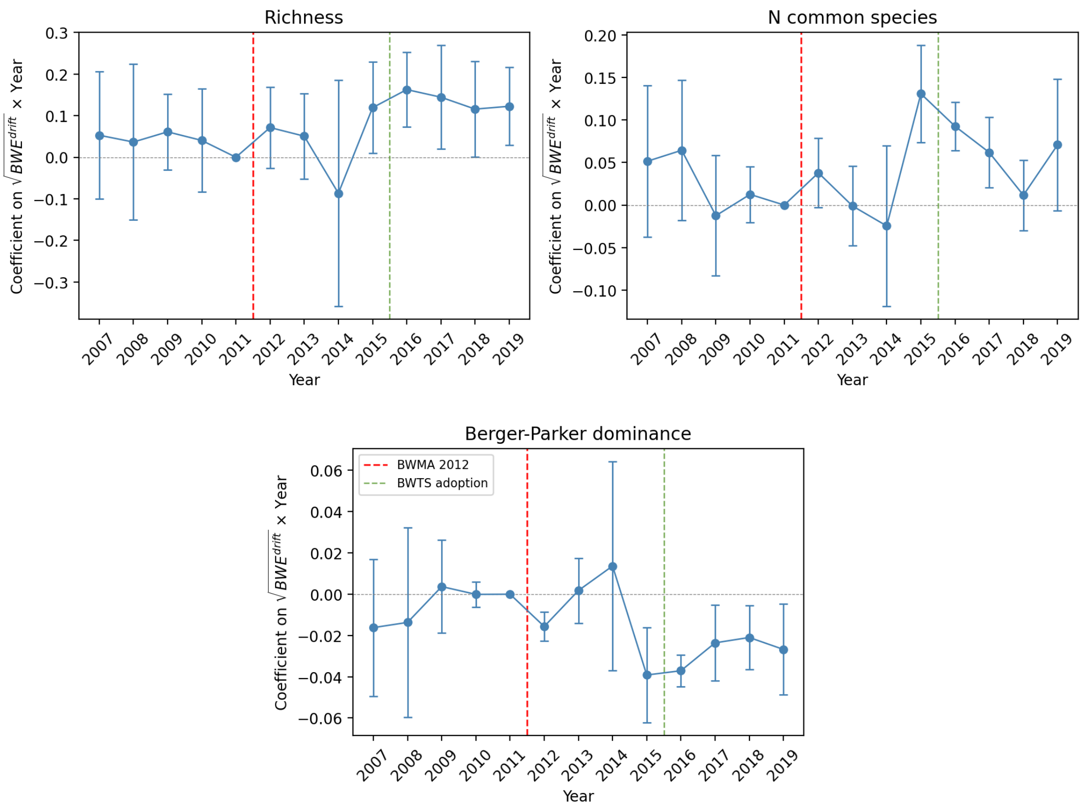

Note: Event-study estimates of $\mathbb{1}\{t = \tau\} \times \sqrt{BWE^{\text{drift}}}$, 2007–2019, with 2011 as the reference year. Top-left: richness (segment-year, PPML, segment + year FE). Top-right: number of common species, defined as species with $\geq$ 5% share of segment-year abundance (segment-year, PPML, segment + year FE). Bottom-center: Berger-Parker dominance index, the maximum species share (segment-year, OLS, segment + year FE). Controls: $\sqrt{BWE^{\text{drift}}}$ level, SST. 95% confidence intervals, clustered at the country $\times$ basin level. Red line: 2012 regulation. Green line: beginning of BWTS adoption (2015).

The invasion of transit-country coastlines can be illustrated by a species whose dispersal biology matches our Lagrangian treatment — the red lionfish (*Pterois volitans*), whose planktonic larvae drift with surface currents. A detailed analysis using Caribbean-native species as an internal control group is reported in Appendix C.

6 Conclusion

A large trading nation such as the United States bears a particular ecological responsibility. Its trade policy can shape biological outcomes along the shipping corridors it anchors. We study this responsibility by analyzing the transboundary ecological consequences of the U.S. Ballast Water Management Act (2012) at transit-country coastlines.

We find that mid-ocean exchange of ballast water—the dominant compliance margin after 2012—shifts ecological pressure to transit countries. Coastal segments with more drift exposure experience significant increases in occurrence, total abundance, and Shannon diversity. The change is compositional as well as aggregate; more species arrive, more become common, and no single species dominates. For lionfish, records rise to 6.1 times their pre-2012 level across 52 segments and 5 countries, with larvae riding the same surface currents our Lagrangian treatment captures.

Several caveats deserve attention. First, our ecological outcomes measure the count of introduced species but do not capture their impact on native biota or the economic damages they cause. Targeted field studies at transit-country coastlines with high drift exposure—combining species-level abundance surveys with environmental DNA sampling—could detect ecological changes that occurrence-count databases cannot resolve, and would complement the large-scale patterns documented here. Second, while the treatment-era coefficients suggest persistence, disentangling lagged effects of past exchange from contemporaneous effects of remaining exchange requires finer temporal variation than our annual panel provides.

Higher diversity of introduced species communities may alter trophic dynamics and reduce the resilience of coastal ecosystems to future disturbances, with implications for fisheries, aquaculture, and coastal tourism in transit countries. More broadly, the combination of oceanographic particle simulation with econometric methods to measure transboundary ecological incidence opens a research agenda on the spatial externalities of environmental regulation in the maritime sector.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used Claude (Anthropic) to assist with code development, data pipeline automation, and copy-editing of the manuscript. All analytical decisions, econometric specifications, interpretation of results, and scientific conclusions were made by the authors. The authors reviewed and edited all AI-assisted output and take full responsibility for the content of the publication.

References

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge (2023). "When Should You Adjust Standard Errors for Clustering?" In: *Quarterly Journal of Economics* 138.1, pp. 1–35.

Ahrenholz, Dean W. and James A. Morris (2010). "Larval duration of the lionfish, Pterois volitans along the Bahamian Archipelago". In: *Environmental Biology of Fishes* 88.4, pp. 305–309. DOI: 10.1007/s10641-010-9647-4.

Allison, Paul D. and Richard P. Waterman (2002). "Fixed-Effects Negative Binomial Regression Models". In: *Sociological Methodology* 32.1, pp. 247–265. DOI: 10.1111/1467-9531.00117.

Bax, Nicholas, Angela Williamson, Max Aguero, Exequiel Gonzalez Poblete, and Warren Geeves (July 2003). "Marine Invasive Alien Species: A Threat to Global Biodiversity". In: *Marine Policy* 27, pp. 313–323. DOI: 10.1016/S0308-597X(03)00041-1.

Bradie, Johanna et al. (Mar. 2018). "Detection of UV-treatment effects on plankton by rapid analytic tools for ballast water compliance monitoring immediately following treatment". In: *Journal of Sea Research* 133, pp. 177–184. ISSN: 1385-1101. DOI: 10.1016/j.seares.2017.09.002.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna (2024). *Difference-in-Differences with a Continuous Treatment*. NBER Working Paper 32117. National Bureau of Economic Research. DOI: 10.3386/w32117.

Čampara, Leo, Vlado Frančić, Lovro Maglić, and Nermin Hasanspahić (Aug. 2019). "Overview and Comparison of the IMO and the US Maritime Administration Ballast Water Management Regulations". In: *Journal of Marine Science and Engineering* 7.9, p. 283. ISSN: 2077-1312. DOI: 10.3390/jmse7090283.

Candau, Fabien and Florian Lafferrere (2024). *The Impact of International Trade on Maritime Ecosystems: Evidence from the California Emission Control Area*. SSRN Working Paper. DOI: 10.2139/ssrn.4984367. URL: https://ssrn.com/abstract=4984367.

Carter, Andrew V., Kevin T. Schnepel, and Douglas G. Steigerwald (2017). "Asymptotic Behavior of a t-Test Robust to Cluster Heterogeneity". In: *Review of Economics and Statistics* 99.4, pp. 698–709.

Chassignet, Eric P., Harley E. Hurlburt, Ole Martin Smedstad, George R. Halliwell, Patrick J. Hogan, Alan J. Wallcraft, Remy Baraille, and Rainer Bleck (2007). "The HYCOM (HYbrid Coordinate Ocean Model) data assimilative system". In: *Journal of Marine Systems* 65.1–4, pp. 60–83. DOI: 10.1016/j.jmarsys.2005.09.016.

Correia, Sergio (2015). "Singletons, Cluster-Robust Standard Errors and Fixed Effects: A Bad Mix". In: *Working Paper*.

Costello, Christopher, Michael Springborn, Carol McAusland, and Andrew Solow (2007). "Unintended Biological Invasions: Does Risk Vary by Trading Partner?" In: *Journal of Environmental Economics and Management* 54.3, pp. 262–276. DOI: 10.1016/j.jeem.2007.06.001.

Côté, Isabelle M., Stephanie J. Green, and Mark A. Hixon (2013). "Predatory fish invaders: Insights from Indo-Pacific lionfish in the western Atlantic and Caribbean". In: *Biological Conservation* 164, pp. 50–61. DOI: 10.1016/j.biocon.2013.04.014.

Cuthbert, Ross N. et al. (June 2021). "Global economic costs of aquatic invasive alien species". In: *Science of The Total Environment* 775, p. 145238. ISSN: 0048-9697. DOI: 10.1016/j.scitotenv.2021.145238.

Dobrzycka-Krahel, Aldona, Magda Bogalecka, and Maciej Szczepanek (2024). "Using biophysical modelling and marine connectivity to assess the risk of natural dispersal of non-indigenous species to comply with the Ballast Water Management Convention". In: *Biological Invasions* 26, pp. 3221–3237. DOI: 10.1007/s10530-024-03327-0.

Gerhard, William A., Kim Lundgreen, Guillaume Drillet, Raphael Baumler, Henrik Holbech, and Claudia K. Gunsch (Nov. 2019). "Installation and use of ballast water treatment systems – Implications for compliance and enforcement". In: *Ocean & Coastal Management* 181, p. 104907. ISSN: 0964-5691. DOI: 10.1016/j.ocecoaman.2019.104907.

Guimarães, Paulo (2008). "The Fixed Effects Negative Binomial Model Revisited". In: *Economics Letters* 99.1, pp. 63–66. DOI: 10.1016/j.econlet.2007.05.030.

Hallegraeff, Gustaaf M. and Christopher J. Bolch (1992). "Transport of Diatom and Dinoflagellate Resting Spores in Ships' Ballast Water: Implications for Plankton Biogeography and Aquaculture". In: *Journal of Plankton Research* 14.8, pp. 1067–1084.

Hansen-Lewis, Jamie and Michelle Marcus (2022). *Uncharted Waters: Effects of Maritime Emission Regulation*. NBER Working Paper 30181. National Bureau of Economic Research. DOI: 10.3386/w30181.

Hausman, Jerry, Bronwyn H. Hall, and Zvi Griliches (1984). “Econometric Models for Count Data with an Application to the Patents–R&D Relationship”. In: *Econometrica* 52.4, pp. 909–938. DOI: 10.2307/1911191.

Hess-Erga, Ole-Kristian, Javier Moreno-Andrés, Øivind Enger, and Olav Vadstein (Mar. 2019). “Microorganisms in ballast water: Disinfection, community dynamics, and implications for management”. In: *Science of The Total Environment* 657, pp. 704–716. ISSN: 0048-9697. DOI: 10.1016/j.scitotenv.2018.12.004.

Johnston, Matthew W. and Sam J. Purkis (2015). “Lionfish in the eastern Pacific: a cellular automaton approach to assessing invasion risk”. In: *Biological Invasions* 17.1, pp. 337–353. DOI: 10.1007/s10530-014-0696-z.

Klotz, Richard and Julia Berazneva (2022). “Local Standards, Behavioral Adjustments, and Welfare: Evaluating California’s Ocean-Going Vessel Fuel Rule”. In: *Journal of the Association of Environmental and Resource Economists* 9.3, pp. 383–424. DOI: 10.1086/717585.

Lakshmi, E., M. Priya, and V. Sivanandan Achari (Jan. 2021). “An overview on the treatment of ballast water in ships”. In: *Ocean & Coastal Management* 199, p. 105296. ISSN: 0964-5691. DOI: 10.1016/j.ocecoaman.2020.105296.

Lezama-Ochoa, Nerea, Gregório M. Ruiz, Derek Spielman, and Anson H. Hines (2022). “Dispersal models alert on the risk of non-native species introduction by ballast water in protected areas from the Western Antarctic Peninsula”. In: *Frontiers in Ecology and Evolution* 10. DOI: 10.3389/fevo.2022.782847.

MacIsaac, Hugh J., Emma M. De Roy, Brian Leung, Alice Grgicak-Mannion, and Gregory M. Ruiz (2016). “Possible ballast water transfer of lionfish to the eastern Pacific Ocean”. In: *PLoS ONE* 11.11, e0165584. DOI: 10.1371/journal.pone.0165584.

MacKinnon, James G., Morten Ørregaard Nielsen, and Matthew D. Webb (2023). “Cluster-Robust Inference: A Guide to Empirical Practice”. In: *Journal of Econometrics* 232.2, pp. 272–299.

MacKinnon, James G. and Matthew D. Webb (2020). “Randomization Inference for Difference-in-Differences with Few Treated Clusters”. In: *Journal of Econometrics* 218.2, pp. 435–450.

Magurran, Anne E. (2004). *Measuring Biological Diversity*. Oxford: Blackwell Publishing.

Margolis, Michael, Jason F. Shogren, and Carolyn Fischer (2005). “How trade politics affect invasive species control”. In: *Ecological Economics* 52.3, pp. 305–313. DOI: 10.1016/j.ecolecon.2004.07.017.

McAusland, Carol and Christopher Costello (2004). “Avoiding Invasives: Trade-Related Policies for Controlling Unintentional Exotic Species Introductions”. In: *Journal of Environmental Economics and Management* 48.2, pp. 954–977. DOI: 10.1016/j.jeem.2003.11.002.

Miller, A. Whitman, Mark S. Minton, and Gregory M. Ruiz (Nov. 2011). “Geographic Limitations and Regional Differences in Ships’ Ballast Water Management to Reduce Marine Invasions in the Contiguous United States”. In: *BioScience* 61.11, pp. 880–887. ISSN: 0006-3568. DOI: 10.1525/bio.2011.61.11.7.
Molnar, Jennifer L, Rebecca L Gamboa, Carmen Revenga, and Mark D Spalding (Nov. 2008). “Assessing the global threat of invasive species to marine biodiversity”. In: *Frontiers in Ecology and the Environment* 6.9, pp. 485–492. ISSN: 1540-9309. DOI: 10.1890/070064.
Okubo, Akira (1971). “Oceanic Diffusion Diagrams”. In: *Deep-Sea Research* 18.8, pp. 789–802.
Pimentel, David, Rodolfo Zuniga, and Doug Morrison (Feb. 2005). “Update on the environmental and economic costs associated with alien-invasive species in the United States”. In: *Ecological Economics* 52.3, pp. 273–288. ISSN: 0921-8009. DOI: 10.1016/j.ecolecon.2004.10.002.
Poje, Andrew C. et al. (2014). “Submesoscale Dispersion in the Vicinity of the Deepwater Horizon Spill”. In: *Proceedings of the National Academy of Sciences* 111.35, pp. 12693–12698.
Richardson, Philip L. (2005). “Caribbean Current and Eddies as Observed by Surface Drifters”. In: *Deep-Sea Research Part II* 52.3–4, pp. 429–463.
Santos Silva, J.M.C. and Silvana Tenreyro (2006). “The Log of Gravity”. In: *The Review of Economics and Statistics* 88.4, pp. 641–658. DOI: 10.1162/rest.88.4.641.
Schofield, Pamela J. (2009). “Geographic extent and chronology of the invasion of non-native lionfish (Pterois volitans [Linnaeus 1758] and P. miles [Bennett 1828]) in the Western North Atlantic and Caribbean Sea”. In: *Aquatic Invasions* 4.3, pp. 473–479. DOI: 10.3391/ai.2009.4.3.5.
Seebens, Hanno, Michael T. Gastner, and Bernd Blasius (2013). “The risk of marine bioinvasion caused by global shipping”. In: *Ecology Letters* 16.6, pp. 782–790. DOI: 10.1111/ele.12111.
Shanks, Alan L. (2009). “Pelagic Larval Duration and Dispersal Distance Revisited”. In: *Biological Bulletin* 216.3, pp. 373–385.
Shannon, Claude E. (1948). “A Mathematical Theory of Communication”. In: *The Bell System Technical Journal* 27.3, pp. 379–423.
Wang, Haifeng, Andrew M. Countryman, James J. Corbett, and Mandana Saebi (2021). “Economic and Environmental Impacts of Ballast Water Management on Small Island Developing States and Least Developed Countries”. In: *Journal of Environmental Management* 300, p. 113816.
Young, Alwyn (2019). “Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results”. In: *Quarterly Journal of Economics* 134.2, pp. 557–598. DOI: 10.1093/qje/qjy029.

Appendix A: Setting, Data, and Sample Diagnostics

Maritime connectivity and BWE exposure

Figure 7 compares U.S.-bound traffic intensity within each transit-country EEZ to domestic port-call activity. For several EEZs, U.S.-bound traffic represents a disproportionate share of the passing fleet, motivating the transboundary externality channel.

Figure 7: Exposure to U.S.-bound traffic in transit-country EEZs

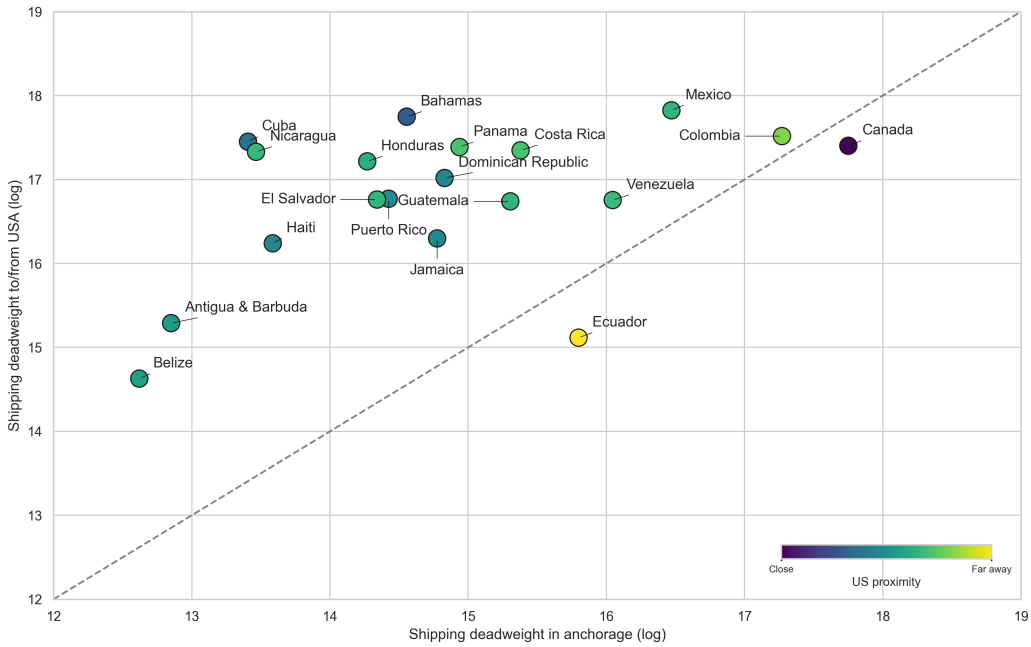

Note: Each point = ratio of U.S.-bound DWT passing through a country's EEZ to total DWT calling at domestic ports (log scale), averaged 2006–2012. Distances from CEPII gravity database.

Figure 8 shows that exchange volumes decreased in U.S. waters after 2012 while increasing in several transit-country EEZs (Cuba, Mexico, Cayman Islands), consistent with relocation of exchange activity toward transit-country waters.

Figure 8: Ballast Water Exchanges in Mid-Ocean by EEZ

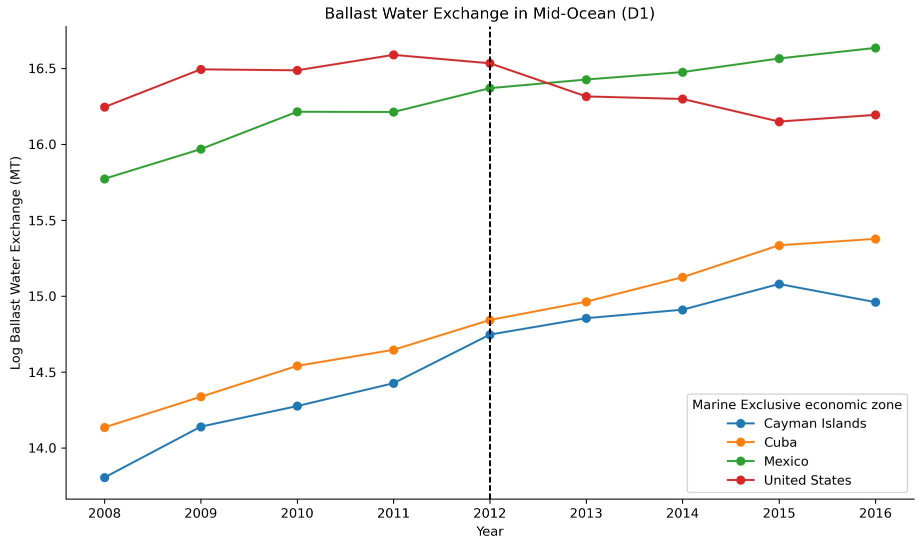

Note: Sum of reported mid-ocean exchange volumes (tankers, bulkers, container ships, general cargo) within each EEZ boundary.

Outcome definitions and descriptive statistics

Outcome definitions. Invasive species occurrence is a count at the segment-species-year level; all other outcomes (total abundance, Shannon diversity, richness, number of common species, and Berger-Parker) are at the segment-year level. *Total abundance* is the sum of WRiMS occurrence counts across species at each segment-year. The event study additionally reports a *new arrivals* panel, defined as the count of species detected for the first time at a segment in a given year. New arrivals appear only in the event study figure (Figure 4, top-right panel); they are not a column of the main table. The composition indicators (richness, number of common species, Berger-Parker) are defined in Section 5.

The *Shannon diversity index* H' is defined as

$$H'_{it} = -\sum_{s} p_{s,it} \ln(p_{s,it}),$$

where $p_{s,it}$ is the relative abundance of introduced species s among all introduced species recorded at segment i in year t (Shannon, 1948). A high H' indicates that many species contribute equally to the established community; a low H' indicates dominance by a few species. As a community-level measure, H' captures ecological restructuring that occurrence alone cannot detect. We compute H' on all introduced species detected at the segment-year without a persistence filter: monitoring gaps in OBIS mean that a species recorded in only one year is not necessarily transient, and restricting to species seen in multiple years would systematically exclude new arrivals—precisely the mechanism through which drift exposure is expected to act (Magurran, 2004).

Summary statistics and sample diagnostics. Table 3 reports summary statistics for the segment-level estimation panel (2007–2019). Singleton observations are dropped prior to estimation (Correia, 2015); singleton-adjusted estimation samples are reported in each table.

Table 3: Descriptive Statistics

	N	Mean	SD	Min	Median	Max	Zeros (%)	Miss. (%)
Occurrence (species-year)	5,877	57.06	386.09	1.00	5.00	7,100.0	0.0	0.0
Total Abundance	938	357.48	1,497.8	1.00	12.00	15,886.0	0.0	0.0
Shannon H′	938	0.9608	1.00	0.0000	0.6931	3.82	38.3	0.0
Richness	938	6.27	7.96	1.00	3.00	51.00	0.0	0.0
N common species (share $\geq$ 5%)	938	3.03	2.42	0.0000	2.00	16.00	0.4	0.0
Berger-Parker	938	0.6399	0.3385	0.0447	0.6667	1.00	0.0	0.0
BWE^{drift} (000s MT)	938	27.06	105.12	0.0000	1.40	1,286.1	10.8	0.0
$\sqrt{BWE^{drift}}$	938	2.82	4.37	0.0000	1.18	35.86	10.8	0.0
Exchange era $\times \sqrt{BWE^{drift}}$	938	0.9814	3.31	0.0000	0.0000	35.86	73.6	0.0
Treatment era $\times \sqrt{BWE^{drift}}$	938	1.01	2.84	0.0000	0.0000	32.49	70.6	0.0
Sea surface temp. (°C)	872	27.31	1.13	18.29	27.69	29.46	0.0	7.0

Notes: All statistics are computed on the 2007–2019 transit-country segment panel. Occurrence is at the segment-species-year level; all other outcomes, treatment variables, and controls are at the segment-year level. Zeros (%) = share of non-missing observations equal to zero. Miss. (%) = share of missing values.

Overdispersion and the choice of estimator

Figure 9 documents the distribution of the occurrence outcome and its variance-mean relationship. Panel (A) plots the histogram of raw occurrence counts at the segment-species-year level: the distribution is extremely right-skewed, with a median of 5 and a maximum of $7,100$. Panel (B) plots log conditional variance against log conditional mean across segment-year cells, with reference lines for Poisson (slope = 1) and negative binomial (slope = 2). The empirical slope is 2.22 and the pooled variance-to-mean ratio is $2,613$.

The Poisson reference line (slope = 1) is far from the data. Observed variances grow far more than proportionally with the mean. Negative binomial does better but is still not exact—the empirical slope of 2.22 exceeds the NegBin benchmark of 2. No standard parametric count distribution fits our data. This is the precise setting in which Santos Silva and Tenreyro (2006) recommend Poisson Pseudo-Maximum Likelihood. PPML treats the Poisson likelihood as a quasi-MLE. Its point estimates are consistent as long as the conditional mean is correctly specified, regardless of the true conditional distribution. Overdispersion is absorbed into the standard errors through the robust (sandwich) variance estimator, clustered at the country $\times$ basin level throughout.

A natural question is whether a negative binomial MLE would give more efficient estimates given the better fit of Panel (B). We prefer PPML for four reasons. First, NegBin is consistent only if the conditional distribution is genuinely negative binomial; Panel (B) shows that the empirical slope exceeds 2, so even NegBin is misspecified, and its MLE can be biased when the distributional assumption fails. PPML is robust to this. Second, the negative binomial fixed-effects estimator (Hausman et al., 1984) does not sweep out unit heterogeneity the way Poisson fixed effects do (Allison and Waterman, 2002; Guimarães, 2008); our design leans heavily on segment, species, and year fixed effects, exactly the setting in which PPML dominates. Third, overdispersion is already handled correctly in our specification through cluster-robust standard errors, so the efficiency loss from using a quasi-MLE is the appropriate price for consistency under distributional misspecification. Fourth, PPML with clustered standard errors is the default inference strategy for count outcomes in the trade and environmental economics literatures, which aligns our approach with the standard benchmark readers expect.

Figure 9: Distribution and Overdispersion of the Occurrence Outcome

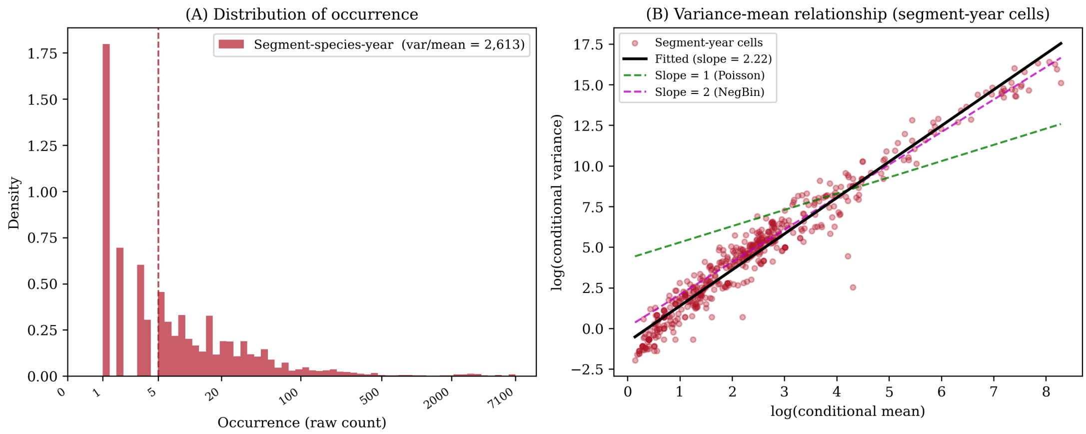

Notes: Left panel: histogram of raw occurrence counts (x-axis on log scale) for the segment-species-year estimation sample (2007–2019). Dashed vertical line marks the median. Right panel: log conditional variance versus log conditional mean of occurrence across segment-year cells; fitted slope shown alongside reference lines for Poisson (slope = 1) and negative binomial (slope = 2). The fitted slope of 2.22 exceeds the NegBin benchmark of 2 and is far above the Poisson benchmark of 1, indicating that no standard parametric count distribution fits the data. PPML is consistent under this form of misspecification; see discussion in the text.

Appendix B: Robustness and Validation

All robustness tables in this appendix use the main merged specification (Equation 2, drift exposure interacted with exchange-era and treatment-era dummies, 2007–2019). They assess the sensitivity of the main identification to alternative treatment construction, country composition, sample restrictions, and monitoring controls.

Control coefficients

Table 4 reports the coefficients on the two controls included in every column of Tables 1 and 2: $\sqrt{BWE^{\text{drift}}}$ (levels) and sea surface temperature.

Table 4: Control Coefficients for Tables 1 and 2

	(1) Occurrence	(2) Abundance	(3) Shannon H′	(4) Richness	(5) N common	(6) Berger-Parker
$\sqrt{BWE^{drift}}$ (levels)	-0.2545***	-0.2837***	-0.0756**	-0.0549	-0.0170	0.0296**
	(0.0351)	(0.0390)	(0.0354)	(0.0525)	(0.0522)	(0.0134)
Sea surface temp.	-2.7492**	-0.4505	-0.7412*	-0.7044*	-0.0919	0.1277
	(1.2379)	(1.4612)	(0.4309)	(0.3849)	(0.1771)	(0.1086)
Observations	5,211	724	623	724	724	724
Era × $\sqrt{BWE^{drift}}$	✓	✓	✓	✓	✓	✓
Segment FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Species FE	✓					

Notes: Coefficients on the two controls included in every column of Tables 1 and 2. Columns (1)–(3) reproduce the specifications of Table 1 (occurrence, total abundance, Shannon diversity); Columns (4)–(6) reproduce the specifications of Table 2 (richness, number of common species, Berger-Parker). Each column includes the exchange-era and treatment-era $\sqrt{BWE^{\text{drift}}}$ interactions (reported in the main-text tables) as well as segment and year fixed effects; Column (1) additionally includes species fixed effects. Columns (1)–(2) and (4)–(5) are PPML; Column (3) is PPML on Shannon H′; Column (6) is OLS. Standard errors clustered at the country × basin level. Time period: 2007–2019. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Treatment construction

Table 5 compares Lagrangian drift exposure with a naive distance-based measure. The naive specification sums BWE volumes within a 300 km radius of each segment's centroid, ignoring ocean currents, and is also entered as a square root. The Lagrangian coefficients are substantially larger than their naive counterparts for both eras, supporting the current-based treatment assignment as the relevant exposure channel.

Table 5: Lagrangian vs. Naive Treatment Assignment

	(1) Lagrangian	(2) Naive 300km
Exchange era (2012–15) $\times \sqrt{BWE}$	0.1335***	-0.0184
	(0.0414)	(0.0273)
Treatment era (2016–19) $\times \sqrt{BWE}$	0.1903***	-0.0001
	(0.0261)	(0.0233)
Observations	5,211	5,211
Clusters (country $\times$ basin)	33	33
Segment FE	✓	✓
Year FE	✓	✓
Species FE	✓	✓

Notes: PPML at segment-species-year level. Column (1): Lagrangian drift-weighted $\sqrt{BWE^{\text{drift}}}$. Column (2): $\sqrt{\text{BWE}}$ summed within 300 km of each segment centroid, ignoring ocean currents. Both columns use segment, year, and species FE. Standard errors clustered at the country $\times$ basin level. Time period: 2007–2019. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6 reports the sensitivity of the Lagrangian treatment to two key parameters of the particle-tracking model: the advection horizon and the sub-grid diffusion coefficient. Alternative 1 shortens the drift horizon from 60 to 50 days; this attenuates the coefficients but preserves the sign and statistical significance of all six era interactions. Alternative 2 lowers K_h from 50 to 30 m^2/s (weaker diffusion, within the 10–50 range used in the ballast-water literature); the estimates are essentially unchanged.

Table 6: Lagrangian Sensitivity: Advection Horizon and Diffusion

	(1) Occurrence	(2) Abundance	(3) Shannon H′
Baseline (60d)			
Exchange era × $\sqrt{BWE^{drift}}$	0.1281*** (0.0384)	0.1036*** (0.0385)	0.0828** (0.0336)
Treatment era × $\sqrt{BWE^{drift}}$	0.1866*** (0.0250)	0.2436*** (0.0321)	0.1162*** (0.0297)
50 days ($K_h = 50$)			
Exchange era × $\sqrt{BWE^{drift}}$	0.0851*** (0.0306)	0.0497* (0.0280)	0.0704* (0.0371)
Treatment era × $\sqrt{BWE^{drift}}$	0.1491*** (0.0249)	0.1964*** (0.0304)	0.0963*** (0.0299)
$K_h = 30$ (60 days)			
Exchange era × $\sqrt{BWE^{drift}}$	0.1232*** (0.0356)	0.0957*** (0.0334)	0.0811** (0.0339)
Treatment era × $\sqrt{BWE^{drift}}$	0.1838*** (0.0248)	0.2370*** (0.0308)	0.1145*** (0.0293)
Segment FE	✓	✓	✓
Year FE	✓	✓	✓
Species FE	✓		

Notes: PPML. Baseline: 60-day Lagrangian particle advection with $K_h = 50$ m^2/s sub-grid diffusion. Alternative 1: 50-day advection (shorter drift horizon). Alternative 2: $K_h = 30$ m^2/s (weaker diffusion, within the 10–50 range used in the ballast-water literature; Lezama-Ochoa et al., 2022; Dobrzycka-Krahel et al., 2024). Columns (1)–(3): occurrence, total abundance, Shannon diversity, each with exchange-era and treatment-era $\sqrt{BWE^{\text{drift}}}$ interactions. Segment, year, and species FE (occurrence) or segment and year FE (abundance, Shannon). SE clustered at country × basin level. Time period: 2007–2019. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Sample robustness

Tables 7 and 8 report leave-one-country-out diagnostics for the segment occurrence and Shannon specifications, with both the exchange-era and treatment-era coefficients reported. The coefficients remain stable across exclusions, confirming that no single country drives the main findings.

Table 7: Leave-One-Country-Out: D1 Segment Occurrence

Dropped	Obs. dropped	Exchange $\hat{\beta}$	(SE)	Treatment $\hat{\beta}$	(SE)
Full sample	—	0.1335***	(0.0414)	0.1903***	(0.0261)
PRI	1,572	0.1131	(0.1453)	0.0926	(0.1009)
COL	1,238	0.0184	(0.0538)	0.1025	(0.1047)
VIR	990	0.1368***	(0.0421)	0.1976***	(0.0282)
MEX	635	0.1536***	(0.0432)	0.1975***	(0.0220)
VEN	555	0.1338***	(0.0420)	0.1914***	(0.0262)
MTQ	191	0.1344***	(0.0413)	0.1905***	(0.0259)
GLP	156	0.1332***	(0.0413)	0.1900***	(0.0261)
PAN	117	0.1331***	(0.0412)	0.1903***	(0.0261)
AIA	48	0.1335***	(0.0415)	0.1903***	(0.0261)
SXM	23	0.1336***	(0.0415)	0.1904***	(0.0260)

Notes: Each row drops one country and re-estimates the merged-spec PPML on the remaining sample. Treatment: exchange-era and treatment-era interactions with $\sqrt{BWE^{\text{drift}}}$. Segment, year, and species FE. Controls: $\sqrt{BWE^{\text{drift}}}$ level, SST. Standard errors clustered at the country $\times$ basin level. Time period: 2007–2019. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Leave-One-Country-Out: D1 Segment Shannon Diversity

Dropped	Obs. dropped	Exchange $\hat{\beta}$	(SE)	Treatment $\hat{\beta}$	(SE)
Full sample	—	0.0782**	(0.0346)	0.1194***	(0.0304)
MEX	240	0.0154	(0.0428)	0.1177**	(0.0569)
COL	219	0.0710*	(0.0375)	0.0986***	(0.0285)
VEN	139	0.0637*	(0.0354)	0.1057***	(0.0282)
VIR	103	0.0991***	(0.0278)	0.1318***	(0.0298)
PRI	81	0.0995***	(0.0330)	0.1343***	(0.0404)
PAN	25	0.0798**	(0.0361)	0.1165***	(0.0292)
BHS	15	0.0852**	(0.0351)	0.1233***	(0.0313)
HND	10	0.0782**	(0.0346)	0.1194***	(0.0304)
GLP	10	0.0783**	(0.0351)	0.1235***	(0.0309)
MTQ	10	0.0844**	(0.0341)	0.1243***	(0.0308)

Notes: Each row drops one country and re-estimates the merged-spec PPML on Shannon diversity. Treatment: exchange-era and treatment-era interactions with $\sqrt{BWE^{\text{drift}}}$. Segment and year FE. Controls: $\sqrt{BWE^{\text{drift}}}$ level, SST. Standard errors clustered at the country $\times$ basin level. Time period: 2007–2019. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9 decomposes the effect by port presence. Column (1) interacts each era treatment with a port indicator; the base coefficients capture the effect at segments without ports, while the interactions capture the differential effect at port segments. Column (2) restricts the sample to no-port segments only. The exchange-era coefficient remains significant in the port-interaction specification. The treatment-era coefficient and the no-port-subsample coefficients are smaller and imprecisely estimated—consistent with the reach of the drift channel being concentrated on segments near transit-country ports, where particles delivered by mid-ocean exchange tend to accumulate downstream of commercial shipping lanes.

Table 9: D1 Segment Robustness: Interaction and No-Port Subsample

	(1) Interaction	(2) No-Port
Exchange era (2012–15) × $\sqrt{BWE^{drift}}$	0.1176*	-0.0116
	(0.0698)	(0.0538)
Treatment era (2016–19) × $\sqrt{BWE^{drift}}$	0.1223	0.0215
	(0.1146)	(0.0260)
× has_port (Exchange)	0.0117	
	(0.0740)	
× has_port (Treatment)	0.0700	
	(0.1177)	
Observations	5,211	499
Clusters (country × basin)	33	13
Segment FE	✓	✓
Year FE	✓	✓
Species FE	✓	✓

Notes: PPML. Occurrence (segment-species-year). Column (1): full sample with era × port interactions — base era coefficients capture the effect at no-port segments, interactions capture the differential effect at port segments. Column (2): no-port subsample (segments without any port within 30 km). Segment, year, and species FE. Controls: $\sqrt{BWE^{\mathrm{drift}}}$ level, SST. SE clustered at country × basin level. Time period: 2007–2019. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Monitoring control

Table 10 re-estimates the three headline specifications with $\log(n_{it}^{native})$, the log count of non-WRiMS OBIS records at each segment-year, added as an explicit control for survey effort. All six era coefficients retain their sign and their statistical significance. Exchange-era coefficients are nearly unchanged across all three outcomes. Treatment-era coefficients attenuate: occurrence and abundance drop by roughly 50% while remaining significant at the 1% level, and Shannon attenuates by about 18% while remaining significant at the 5% level. The attenuation is consistent with $\log(n^{native})$ capturing shared sampling-effort variation that the segment and year fixed effects do not absorb on their own, and the monitoring control itself carries a significantly positive sign in every column, as expected for a valid measure of survey intensity.

Table 10: Monitoring-Control Robustness: Occurrence, Abundance, and Shannon

	(1) Occurrence	(2) Abundance	(3) Shannon H′
Exchange era (2012–15) × $\sqrt{BWE^{drift}}$	0.1379***	0.1248***	0.0731**
	(0.0482)	(0.0417)	(0.0289)
Treatment era (2016–19) × $\sqrt{BWE^{drift}}$	0.1073**	0.1033**	0.0981***
	(0.0419)	(0.0409)	(0.0200)
$\sqrt{BWE^{drift}}$ (levels)	-0.2643***	-0.3109***	-0.0800**
	(0.0442)	(0.0415)	(0.0327)
Sea surface temp.	-2.2585	1.5697***	-0.3775
	(1.8488)	(0.3990)	(0.3105)
$\log(n^{native})$	0.2712***	0.4629***	0.2093***
	(0.0869)	(0.0669)	(0.0674)
Observations	5,211	724	623
Clusters (country × basin)	32	32	32
Segment FE	✓	✓	✓
Year FE	✓	✓	✓
Species FE	✓		

Notes: PPML. Column (1): occurrence (segment-species-year), segment + year + species FE. Column (2): total abundance (segment-year), segment + year FE. Column (3): Shannon diversity H' (segment-year), segment + year FE. All columns add $\log(n_{it}^{native})$, the log count of non-WRiMS marine and brackish species records from OBIS, as an explicit monitoring-intensity control. Controls: $\sqrt{BWE^{\text{drift}}}$ level, sea surface temperature. SE clustered at country × basin. Time period: 2007–2019.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Inference validation

Table 11: Randomization Inference and Rotated-Currents Placebo

	(1) Occurrence	(2) Abundance	(3) Shannon H′
Panel A: Randomization inference (Fisher permutation, exact p-values)			
Exchange era (2012–15) p-value	0.001	0.003	0.001
Treatment era (2016–19) p-value	0.001	0.001	0.001
Panel B: Rotated-currents placebo (180° rotation of HYCOM u and v)			
Exchange era (2012–15) × $\sqrt{BWE^{\text{drift}}}$	-0.0016	+0.0041	-0.0001
(p-value)	[0.910]	[0.877]	[0.996]
Treatment era (2016–19) × $\sqrt{BWE^{\text{drift}}}$	-0.0146	-0.0390	+0.0123
(p-value)	[0.287]	[0.095]	[0.771]

Notes: Two distinct checks with different null hypotheses. Panel A is an exact-inference procedure, not a placebo: within each year, drift exposure $\sqrt{BWE_{it}^{\text{drift}}}$ is permuted across segments, keeping yearly totals, segment/year/species fixed effects, and SST_{it} at their observed values; Equation 2 is re-estimated; the reported p-value is the fraction of $B = 1000$ permutation draws whose era-interaction coefficient has magnitude at least as large as the observed one (Young, 2019). A small p-value *rejects* the null that the segment-level spatial pattern of drift is uninformative — it validates the paper's identification. Panel B is a placebo in the standard sense: we re-run the full Lagrangian pipeline with the HYCOM current field rotated 180° (both u and v negated) and re-estimate Equation 2. The rotation preserves current magnitudes, event locations, event volumes, and segment geography; it breaks only the physical routing from each BWE event to its coastal-segment exposure. Under the null that the physical drift channel is load-bearing, the rotated-field era coefficients should collapse toward zero.

Additional threats to identification

Two further threats, discussed briefly here, supplement the three main threats in Section 4.

Vessel anticipation. Vessel operators may have begun relocating ballast-water exchange mid-ocean in anticipation of the 2012 rule, which would push the treatment signal into the pre-period and bias the post-2012 break downward. This is unlikely on incentive grounds: exchange is costly, and pre-2012 behavior faced no penalty, so operators had no reason to incur that cost before the rule took effect. The event study confirms this empirically: pre-2012 coefficients remain close to zero through 2011, with the break occurring sharply at the policy date rather than drifting up before it.

BWE-location exogeneity under a monthly climatology. Because $BWE_{it}^{\mathrm{drift}}$ is built from a monthly current climatology rather than year-specific currents, all time variation in the treatment variable flows from changes in the upstream distribution of BWE locations. Vessels choose those locations. The identifying assumption is that, conditional on segment and year fixed effects, the post-2012 redistribution of BWE coordinates is not driven by segment-level time-varying factors that also affect introduced-species outcomes. This is satisfied by design: a remote transit segment's IAS trajectory has no plausible channel through which it could influence a U.S.-bound vessel's mid-ocean exchange coordinates, which are chosen under the 200-nautical-mile rule and constraints of route and weather. The segment fixed effects absorb all time-invariant correlation, the year fixed effects absorb the aggregate trend in Caribbean/Pacific exchange activity, and the falsification regression on n_{it}^{native} rules out a surveillance-driven version of the concern.

Appendix C: Lionfish and the Drift Channel

This appendix presents a species-specific illustration of the transit-country externality documented in Section 4. The red lionfish (*Pterois volitans*) is a natural test case: its planktonic larvae drift with Caribbean surface currents for roughly a month, matching the physical channel our Lagrangian treatment captures. We use a triple-difference specification with Caribbean-native species as an internal control group.

Scope and interpretation. Two clarifications apply throughout this appendix. First, the triple-difference is identified off within-segment-year cross-species variation, a different source of variation from the headline Section 4 design; we therefore treat it as a complementary illustration of the drift channel mechanism rather than a replication of the main results. Second, the parameterization in Equation 3 omits the two-way interactions Era $\times$ NonNative. A fully-saturated alternative that adds these interactions yields a qualitatively similar result for the broader 16-species non-native group (Column 2 below) but attenuates and imprecisely estimates the lionfish-specific coefficient (Column 1). We report the under-saturated specification because our object of interest is the differential response to drift exposure, not the level difference between native and non-native detection rates. The Column 2 broader-non-native differential survives the saturated specification and is the methodologically robust finding; Column 1 should be read as descriptive illustration of the lionfish case rather than an independent causal estimate.

Lionfish as a natural test case. The red lionfish, a venomous Indo-Pacific reef fish native to the western Pacific and eastern Indian Ocean, has staged what is now recognized as the fastest invasion by a marine fish in recorded history. The first record in the Western Atlantic dates from 1985, when individuals were sighted near Dania Beach, Florida, as the residue of an aquarium release (Schofield, 2009). Establishment was slow through the 1990s, but from the early 2000s the invasion accelerated sharply.[15] The species now occupies approximately 7.3 million km^2 of the Western Atlantic and Caribbean, frequently at population densities exceeding those in the native range, and is responsible for documented declines in native reef-fish recruitment (Côté et al., 2013).

Two features of lionfish biology make it a natural test case for the mechanism we identify. The first is extreme reproductive output combined with small early-stage larvae that pass through standard ballast-tank intake screens.[16] Ballast water is therefore not a hypothetical vector for lionfish; MacIsaac et al. (2016) document 27 Caribbean-to-Pacific ballast transfers from lionfish-colonized ports between 2006 and 2013. The second is a long pelagic larval duration of about a month, long enough for Caribbean surface currents to carry larvae hundreds of kilometers.[17] Johnston and Purkis (2015) find ocean current to be the single most influential parameter governing lionfish dispersal. The physical channel that governs lionfish dispersal in the open ocean—planktonic larvae riding surface currents—is precisely the channel our Lagrangian treatment variable measures.

Empirical strategy: native species as an internal placebo. A feature of the WRiMS taxonomy lets us sharpen identification beyond Sections 4 and 5. WRiMS flags any species introduced *anywhere* in the world, so the panel contains both genuine Caribbean/Pacific non-natives (lionfish, invasive corals, ballast-water bivalves) and Caribbean-native species flagged because they are invasive in other oceans (large crabs, sharks, toxic or armored reef fishes). Native species are much less likely to respond to the drift channel than introduced species, because U.S.-bound mid-ocean exchanges primarily discharge water loaded in distant foreign ports rather than in the Caribbean. Any residual re-seeding of natives via Caribbean-origin exchanges would bias the triple-difference coefficient toward zero, so the estimate is conservative. Native species therefore serve as an internal placebo for

[15] Lionfish were recorded in Bermuda in 2000, established a breeding population in the Bahamas in 2004, and from there expanded across the Caribbean: Turks and Caicos (2006), then Colombia, the Cayman Islands, Jamaica, Puerto Rico, Haiti, Belize, and the Dominican Republic (all 2008). Mexico, Central America, and northern South America followed within a few years.

[16] Mature females spawn every three to four days in warm waters, releasing each time two gelatinous egg masses containing up to 25,000 eggs, for a total annual output on the order of two million eggs per female. Early larvae measure 1.5 to 11 mm in body length, well below the 15 to 25 mm mesh of the sea-chest intake screens that protect vessel ballast pumps.

[17] Ahrenholz and Morris (2010) report 20 to 35 days planktonic, mean 26.

any segment-year-level confounder. The estimating equation is

$$\begin{aligned}\log \mathbb{E}[\mathrm{IAS}_{sit} \mid X_{sit}] = \gamma_1\, \mathrm{Exchange}_t \times \sqrt{BWE_{it}^{\mathrm{drift}}} \times \mathrm{NonNative}_s \\ + \gamma_2\, \mathrm{Treatment}_t \times \sqrt{BWE_{it}^{\mathrm{drift}}} \times \mathrm{NonNative}_s \\ + \gamma_3 \sqrt{BWE_{it}^{\mathrm{drift}}} \times \mathrm{NonNative}_s + f_{it} + f_s,\end{aligned} \tag{3}$$

where $\mathrm{NonNative}_s \in \{0, 1\}$ indicates membership in the non-native treatment group, f_{it} are segment-by-year fixed effects that absorb everything varying at the segment-year level—drift exposure itself, SST, monitoring intensity, and any shock common to all species at that segment-year—and f_s are species fixed effects. The coefficients γ_1 and γ_2 are the differential response of non-native species to drift exposure, relative to natives, during each policy era; γ_3 captures any always-on differential. We classify 16 species as unambiguous Caribbean-region non-natives and 21 Caribbean-native species (not known to be lionfish prey) as the control group. Species with contested Caribbean-native status (such as *Percnon gibbesi* and *Perna perna*, both treated as introduced by some authorities but native by others) are excluded from the treatment group; lionfish prey, freshwater tilapia, and unclassified plankton are also dropped.

Identifying assumptions. The design rests on two assumptions. The first is that native species are a valid placebo for the drift channel. This would fail if lionfish predation were itself depressing native occurrence at high-drift segments post-2012, which would mechanically inflate γ_1 and γ_2. We address this directly by excluding from the control group every species documented as lionfish prey (small reef fish—damselfishes, gobies, blennies, juvenile surgeonfishes, and angelfishes). The retained controls are large crabs, sharks, rays, toxic or armored reef fishes, and native filter-feeders—species whose adults are not routinely consumed by lionfish. The second assumption is that, absent the 2012 rule, the relative response of non-native and native species to drift exposure would have remained stable. This is the triple-differences analog of parallel trends. The segment-by-year fixed effects absorb any shock that moves natives and non-natives together at a given coastline in a given year, so a violation would require a shock that hits non-natives differentially, specifically at high-drift segments, specifically in 2012 or 2015. No such event is known to us. The 2012 rule targeted ballast-water discharge at U.S. ports and has no direct differential effect on Caribbean natives relative to non-natives, and no regional ecological or regulatory event in 2012 or 2015 selectively targets one group over the other.

Results. Table 12 reports the triple difference. Column 1 restricts the non-native group to *Pterois volitans* alone; Column 2 extends it to the full set of 16 non-native species for external validity.

Table 12: Triple-Difference: Lionfish and Non-Native Species vs Native Placebos

	(1) Lionfish only	(2) All non-natives
Exchange era (2012–15) × $\sqrt{BWE^{drift}}$ × NonNative	0.2206***	0.2505***
	(0.0139)	(0.0179)
Treatment era (2016–19) × $\sqrt{BWE^{drift}}$ × NonNative	0.2214***	0.1575***
	(0.0086)	(0.0502)
$\sqrt{BWE^{drift}}$ × NonNative (base)	-0.2087***	-0.3222***
	(0.0247)	(0.0630)
Observations	1,150	1,338
Clusters (country × basin)	24	26
Segment × Year FE	✓	✓
Species FE	✓	✓

Notes: PPML at the segment-species-year level. The non-native treatment group is *Pterois volitans* alone in Column (1) and the full set of 16 unambiguous Caribbean-region non-natives in Column (2). The control group, identical in both columns, is 21 Caribbean-native species not known to be lionfish prey (large crabs, sharks, toxic or armored reef fishes, filter feeders). Segment-by-year fixed effects absorb all segment-year-level variation, including $\sqrt{BWE^{\text{drift}}}$ itself and SST. Species fixed effects included. Standard errors clustered at the country × basin level. CR3 cluster-jackknife rows are reserved for Tables 1 and 2: the triple-difference identifies γ_1 and γ_2 off within-segment-year cross-species variation rather than cross-cluster variation, so the leverage-concentration concern that motivates CR3 in the headline specifications is not binding here. Singleton observations dropped (Correia, 2015). Time period: 2007–2019. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

For lionfish specifically (Column 1), the exchange- and treatment-era differentials are significant at the 1% level and essentially identical in magnitude. The always-on base differential γ_3 is significantly negative. In the pre-policy years, lionfish responded less positively to drift exposure than the native placebo group, consistent with a species that had not yet fully colonized the high-drift segments by 2011. After 2012, the post-policy shifts are large enough to close this pre-period gap exactly—lionfish occurrence becomes as responsive to drift exposure as native occurrence, as the species expands into the coastlines downstream of mid-ocean exchange. Column 2 confirms that the pattern is not driven by lionfish alone. For the broader 16-species non-native group, both era differentials are also significant at the 1% level. Column 2 differs from Column 1 in a telling way. The Column 2 exchange-era differential is larger than lionfish's, while the treatment-era differential is smaller. Lionfish keep their accelerated colonization through 2019, whereas the pooled non-native group includes slower-establishing taxa (hull-fouling bivalves, tunicates, barnacles) whose relative response peaks earlier in the exchange era and moderates afterwards. The two columns therefore tell compatible but distinct stories: a sharp and sustained lionfish signal, and a broader invader signal concentrated in the immediate post-2012 period.

Robustness checks. Two checks support the design. First, dropping the two most-recorded natives (88% of control-group records)[18] leaves every coefficient unchanged at the third decimal; species fixed effects fully absorb dominant-native variation. Second, excluding species of contested Caribbean-native status from the non-native group strengthens the Column 2 coefficients, so the result is not sensitive to classification judgments at the margin.

Lionfish records at our transit-country segments are 6.1 times the pre-2012 level.[19] The triple-difference shows that this increase is not a byproduct of monitoring trends. Conditional on segment and year, lionfish tracks drift exposure sharply after 2012 while native placebos do not. The 2012 rule did not cause the lionfish invasion, which was already underway; it amplified the drift corridors along which lionfish larvae reach transit-country coastlines.

[18] *Callinectes bocourti* and *C. sapidus*.

[19] Raw counts: 1,090 in 2007–2011, 1,650 in 2012–2015, 4,946 in 2016–2019, spanning 52 segments across 5 countries.