The Hidden Cost of Globalization in the Underwater World: Evidence from California's Kelp Forests

Fabien Candau[*] Florian Lafferrere[‡]

April 1, 2026

Abstract

This article examines how maritime regulation affects one of the most important aquatic ecosystems along the California coast: kelp forests. Combining satellite measures of kelp canopy and biomass with AIS vessel trajectories, we exploit variation in pre-policy shipping exposure across kelp cells. After California's Ocean-Going Vessel Fuel Rule (2009), kelp improves disproportionately where maritime traffic was highest, with estimated gains of approximately 15–17% per unit of log shipping exposure. Event studies show no differential pre-trends, with the treatment effect appearing sharply at implementation. We apply the same design to the 2011 boundary modification around the Channel Islands, which re-concentrated vessel traffic in specific corridors. Kelp declined most in the cells that had the highest pre-modification shipping exposure. JEL: Q57, R41, F18

Keywords: Kelp forests, maritime regulation, emission control areas, difference-in-differences, ecosystem services, AIS, shipping behaviour

1 Introduction

Contemporary globalization relies heavily on maritime routes that generate substantial negative externalities and hidden costs. The human health consequences of shipping-related air pollution are now well documented (e.g., Klotz and Berazneva, 2022; Hansen-Lewis and Marcus, 2022), but evidence on impacts on marine ecosystems remains comparatively limited at the spatial scale relevant to policy. A growing body of work shows

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

[†]We would like to thank David Nagy, Sandra Poncet, Pamina Koenig, Antoine Bouet, Aicha Hasni, Élise Labauvie, Isabelle Chort, and Christophe Maes for valuable comments and discussions. We also thank all the participants of several seminars (RITM, BSE, CEEM), particularly Raphael Chiappini, Tanguy Bernard and Sébastien Desbureaux as well as participants of the 29th Annual Conference of the EAERE. We are grateful for the excellent working conditions provided by the Atlantic Port of Bordeaux, Bordeaux School of Economics, and Pompeu Fabra University. Finally, we thank the Région Nouvelle-Aquitaine (project PORTeur) for providing all the funds necessary for this research.

[‡]UPPA E2S, TREE, France.

that environmental regulations can yield ecological co-benefits beyond their primary design targets (Shapiro and Walker, 2018; Keiser and Shapiro, 2019). Whether maritime emission controls similarly benefit nearshore ecosystems is an open question. This paper investigates it by quantifying how a major coastal shipping regulation and the associated reconfiguration of traffic patterns relate to the dynamics of kelp forests along the California coast. Our identification strategy isolates a specific physical channel—the spatial footprint of maritime activity—through which international trade affects nearshore ecosystems.

Kelp forests are large biogenic habitats that support coastal biodiversity and provide multiple ecosystem services (Dayton, 1985). As a foundation species (Dayton, 1985), kelp creates a habitat structure that functions as nursery grounds and supports complex food webs in Europe (Christie et al., 2009), as also documented in California (Steneck et al., 2002; Carr and Reed, 2019). Recent global accounting exercises suggest that kelp ecosystems contribute significant benefits through fishery production, nutrient cycling, and carbon sequestration (Eger et al., 2023; see also Costello et al., 2020 on the broader economics of marine food systems). Because kelp can respond quickly to changing local conditions through rapid growth and recruitment (Krumhansl et al., 2016; Carr and Reed, 2019), it allows us to detect ecosystem-level responses to the evolution of nearshore stressors over a relatively short horizon.

We study California's Ocean-Going Vessel Fuel Rule, which established a regulated zone in July 2009 and later modified the boundary in southern California at the end of 2011. The policy was designed to reduce sulfur-related emissions by requiring the use of low-sulfur marine distillate fuels within the zone.[1] For our setting, the rule altered vessel routing decisions (Klotz and Berazneva, 2022), thereby changing the spatial footprint of maritime activity near kelp forests. Nearshore traffic declines after 2009 and re-concentrates after the 2011 boundary change. These routing adjustments provide quasi-experimental variation that can be exploited to learn about the sensitivity of kelp forests to maritime exposure.

Our contribution is twofold. First, we combine quarterly satellite-based measures of kelp canopy and biomass with AIS-derived shipping exposure to implement a continuous difference-in-differences design. We exploit within-zone variation in pre-policy vessel traffic; kelp cells that were more exposed to shipping before 2009 should benefit more if the regulation effectively reduces nearshore pressures. We find evidence consistent with this prediction: the interaction between shipping exposure and the post-policy indicator is positive and significant across all specifications. Event-study estimates show no differential pre-trends in kelp outcomes across exposure levels, strengthening the causal

[1] Vessels entering the zone are required to switch from conventional heavy fuel oil — which can contain up to 3.5% sulfur by mass — to low-sulfur marine distillate fuels, capped at 0.1% sulfur by mass. Heavy fuel oil is a residual product of petroleum refining with high sulfur content, high viscosity, and high energy density; marine distillate fuels are lighter, cleaner fractions of crude oil with substantially lower sulfur concentrations. When burned, sulfur-rich fuels generate sulfur dioxide (SO2) and sulfate particulate matter (PM2.5), both of which can deposit into nearshore waters through dry and wet deposition. Under the California rule, vessels must complete the fuel switch before entering the regulated zone—typically within 24 nautical miles of the California baseline—thereby giving operators limited time to flush fuel lines and stabilize engine operations prior to entry.

interpretation.

Second, we apply the same continuous difference-in-differences framework to the 2011 boundary modification around the Channel Islands—a change that re-concentrated traffic in specific corridors. Using pre-modification shipping activity as a predetermined treatment intensity, we show that kelp cells more exposed to vessel traffic before 2011 experienced disproportionately worse outcomes after the boundary change. This result mirrors the first finding: cells that benefited most from the 2009 regulation saw kelp recover; cells most affected by the 2011 re-concentration of traffic saw kelp deteriorate. Taken together, the two designs trace out both sides of the same relationship between maritime exposure and kelp health.

This analysis relates to three strands of work. First, it brings marine forests into the economic analysis of deforestation. A growing literature studies how international economic activity drives terrestrial forest loss (e.g., Berman et al., 2023; Abman and Lundberg, 2020; Carreira et al., 2024; Nedoncelle et al., 2024), but underwater forests have received no attention in economics despite providing comparable ecosystem services—habitat provision, carbon sequestration, and nutrient cycling. By documenting how shipping regulation and the associated reconfiguration of maritime traffic affect kelp forests, we extend the deforestation literature to a new biome and a new mechanism: trade-related transport activity rather than land-use conversion. Second, it complements experimental and site-specific ecological studies on shipping-related stressors (e.g., Johansson et al., 2012) by providing a spatially comprehensive, policy-relevant analysis that follows a regulatory change. Third, it connects to the literature evaluating emission control policies in maritime settings by documenting ecosystem responses in coastal waters, beyond human-health endpoints (e.g., Klotz and Berazneva, 2022; Hansen-Lewis and Marcus, 2022).[2]

2 Background

The California–Baja California coast hosts extensive kelp forests dominated by canopy-forming species (giant kelp, *Macrocystis pyrifera*, and bull kelp, *Nereocystis luetkeana*). These species form surface (or near-surface) canopies that are detectable with satellite imagery (Bell et al., 2023).

Kelp performance is shaped by ocean temperature, nutrient availability, water light (or coastal darkening), as well as localized stressors that can damage fronds, inhibit photosynthesis, or alter community structure (Carr and Reed, 2019). More broadly, climatic variability is a first-order driver of kelp outcomes. Ocean climate oscillations influence sea surface temperatures, salinity, and nutrient availability over multi-year cycles (Castorani et al., 2022). Warm-water events can sharply reduce canopy extent, and the decline observed since 2014 has been linked to prolonged warm conditions and cascading ecological interactions (Rogers-Bennett and Catton, 2019; see also Krumhansl et al., 2016).

[2]More broadly, Keiser and Shapiro (2019) provide a precedent for our approach by estimating the ecological consequences of the Clean Water Act using large-scale environmental monitoring data—a parallel to our use of satellite-based kelp monitoring to evaluate maritime regulation.

These mechanisms motivate the inclusion of sea surface temperature and wave intensity controls in our empirical models.

Maritime activity affects kelp through multiple channels, but their impact are likely to grow with vessel traffic density — making busy corridors disproportionately hostile to kelp relative to quieter waters.

Vessel operations generate discharges (greywater and blackwater) and associated chemical and nutrient loads that can contribute to eutrophication and oxygen depletion (Lindgren et al., 2016), with potential consequences for coastal vegetation and algal community composition (e.g., Filbee-Dexter and Wernberg, 2018).

A second channel operates through the combustion of heavy fuel oil (HFO), the dominant marine fuel prior to the ECA implementation. HFO combustion generates pyrogenic polycyclic aromatic hydrocarbons (PAHs) that deposit into nearshore waters and accumulate in coastal sediments (Yunker et al., 2002). Appendix A provides spatial evidence for this channel: across Southern California Bight monitoring stations, pyrogenic PAH concentrations in pre-ECA sediments are positively and significantly correlated with our AIS-derived measure of pre-policy vessel traffic intensity (Spearman $\rho = 0.55$, $p < 0.01$), confirming that the corridors most exposed to shipping are also the locations where HFO combustion byproducts were most elevated before enforcement.

Shipping can also introduce contaminants through antifouling practices. Biocidal compounds historically used to reduce hull fouling, including tributyltin (TBT) and copper-based formulations could enter coastal waters and affect marine organisms (Dafforn et al., 2011; Cima and Varello, 2022). Laboratory evidence confirms that antifouling compounds inhibit kelp photosynthesis and growth (Johansson et al., 2012), and field measurements in Southern California harbours document biocide concentrations in the range shown to cause sublethal effects in macroalgae (Schiff et al., 2011). Finally, maritime traffic can influence nearshore physical conditions (wake, turbulence, localized disturbance) and can concentrate stressors along heavily trafficked corridors.

A distinct, potentially important channel is the introduction of non-native species via hull fouling and ballast water. The literature has documented introductions linked to shipping, with consequences that vary across species and contexts (Williams and Smith, 2007). Evidence on the impacts of particular invaders on native kelp is mixed, including work on *Sargassum muticum* (Ambrose and Nelson, 1982; Britton-Simmons, 2004; Smith, 2015) and more recent assessments of *Undaria pinnatifida* (McHaskell, 2024).[3] Ballast water can also transport grazers and pathogens with ecosystem-wide effects (e.g., Phinney et al., 2001).

While our empirical designs do not identify the relative importance of these channels, they motivate the hypothesis that changes in the spatial distribution and intensity of shipping activity can matter for kelp outcomes. California's Ocean-Going Vessel Fuel Rule introduced a regulated zone extending 24 nautical miles off the coast starting in 2009 that enables us to test this hypothesis. Covered vessels must use low-sulfur marine distillate fuels within the zone. Compliance increases operating costs within the regulated area

[3] *Undaria pinnatifida* is a kelp from northeast Asia, established in California in the 2000s, and classified as one of the "100 of the world's worst invasive alien species" (in the IUCN's list). The invasion of this kelp may have an effect on the native sub-canopy (e.g., on *Pyrifera* and *Pterygophora californica*) but not on the native canopy kelp studied here, which are not found in the same habitats.

and can induce changes in routing and the spatial distribution of traffic. As we discuss in Section 3.5, voyage-level evidence confirms that vessels sharply adjusted routes following implementation (Klotz and Berazneva, 2022). Finally, in 2011, the southern boundary was modified around the Channel Islands, affecting where ships concentrate as they approach ports and creating additional variation in maritime exposure.

3 The 2009 Ocean-Going Vessel Fuel Rule

This section describes the data and the continuous difference-in-differences design used to estimate changes in kelp canopy and biomass associated with the 2009 implementation of the regulated zone.

3.1 Data

We combine satellite kelp observations, AIS vessel tracking data, and environmental covariates, all aggregated to a common spatial grid of 10 km hexagonal cells at quarterly frequency.

Kelp outcomes and shipping exposure. Kelp outcomes come from Bell et al. (2023), which provides quarterly cloud-free Landsat-based measures of floating kelp canopy at 30m resolution along the Pacific coast.[4] We use two variables: (i) kelp canopy area within each pixel (m^2) and (ii) kelp canopy biomass (kg) inferred from validated relationships between surface reflectance and in situ canopy and biomass measurements (Bell et al., 2023). Satellite detection captures surface (or near-surface) canopy and does not observe subsurface kelp.

We overlay a uniform grid of 10 km hexagonal cells covering the study area (the 2009 regulated zone and the 2011 boundary modification area) and aggregate all variables to this common spatial unit at the quarterly frequency. [5]

For each year–quarter, the 30 m pixel-level kelp observations are converted to point geometries using their latitude and longitude coordinates and spatially joined to the hexagonal grid. Within each hexagon i and quarter t, we compute the sum and mean of kelp canopy area (m^2) and biomass (kg), as well as the count of kelp pixels.

Vessel trajectories are derived from georeferenced ship positions recorded by the Automated Identification System (AIS), made available by the U.S. Coast Guard.[6] We focus

[4]The underlying imagery consists of Landsat Collection 1 Level-2 Surface Reflectance products, which apply atmospheric correction for aerosols and water vapor *before* kelp classification. Images with more than 25% cloud cover are excluded, and remaining cloud and cloud-shadow pixels are masked using the quality-assurance band. Any localized change in atmospheric transparency—such as reduced sulfur aerosols from the fuel regulation—is therefore corrected at the reflectance stage rather than propagated into the kelp canopy estimates.

[5]Hexagonal grids are preferred over square or rectangular tessellations in spatial analysis because they provide equal area and equidistant centroids, minimise edge effects, and reduce sensitivity to the modifiable areal unit problem (Birch et al., 2007; Openshaw, 1984).

[6]AIS data accessed via MarineCadastre.gov, jointly maintained by the U.S. BOEM and NOAA.

on four major cargo vessel classes—tankers, container ships, general cargo ships, and bulk carriers—and summarize their activity at the hexagon–month level in terms of total distance traveled within the cell ("length", in meters) and number of unique vessel passages ("count"). These two measures capture complementary dimensions of shipping exposure. $Length_{it}$ proxies the cumulative intensity of vessel activity—and by extension the volume of associated discharges (exhaust emissions, antifouling leachates, greywater)—which scale with distance traveled. $Count_{it}$ captures the number of discrete vessel passages, reflecting the frequency of vessel-level events rather than their cumulative intensity. Together, they allow us to distinguish the intensive from the extensive margin of shipping exposure. To assign shipping exposure to kelp locations, we proceed as follows. For each hexagon containing kelp, we create a circular buffer of 20 km radius and identify all shipping hexagons that intersect this buffer. We then aggregate the shipping measures—summing length and count—across all identified shipping hexagons for each kelp hexagon i and quarter t. This yields two continuous exposure variables per kelp cell: total distance traveled by vessels in its vicinity ($Length_{it}$) and total number of vessel passages ($Count_{it}$).

Finally, we incorporate the cumulative human impact index of Halpern et al. (2009), which combines 25 anthropogenic stressors and 19 marine ecosystems to map cumulative pressure on the California Current at $\sim$1 km^2 resolution (see also Maxwell et al., 2013).[7] Using zonal statistics, we compute the mean value of this raster within each 10 km hexagonal cell, providing a time-invariant measure of baseline anthropogenic pressure at each kelp location.

Environmental controls. Kelp forests are affected by multiple interacting environmental drivers, including temperature, water motion, and nutrient availability (Carr and Reed, 2019; Castorani et al., 2022). We include three time-varying physical controls that are well-established determinants of kelp dynamics.

Sea surface temperature (SST) is sourced from MODIS monthly composites and aggregated to quarters. SST is a central determinant of kelp dynamics and is also informative about broader ocean conditions (e.g., Yang et al., 2023; Dunstan et al., 2018).[8] Wave intensity is drawn from ERA5 (Hersbach et al., 2023) and measured as significant wave height (the average height of the highest third of waves), aggregated to quarters; it captures physical stress associated with storms and strong sea states. Precipitation is drawn from ERA5 and aggregated to quarters; it proxies terrestrial nutrient runoff, which can affect kelp growth.

[7] For each pixel the index is $I_C = \sum_s \sum_e D_s E_e \mu_{se}$, where D_s is the normalized intensity of stressor s, E_e indicates the presence of ecosystem e, and μ_{se} is an expert-elicited vulnerability weight scored on a 0–4 scale based on spatial extent, frequency, functional impact, resistance, and recovery time.

[8] We do not include a separate upwelling index (e.g., CUTI or BEUTI; Jacox et al., 2018) because SST is strongly negatively correlated with upwelling intensity along the California coast—cold surface water is the primary physical signature of upwelling events—and our year–quarter fixed effects already absorb the dominant seasonal upwelling cycle. Any remaining bias would require a differential change in upwelling at the policy date that is spatially correlated with shipping exposure.

3.2 Specification

Our main design is a continuous difference-in-differences (Callaway et al., 2024) that exploits cross-sectional variation in pre-policy shipping exposure. The intuition is straightforward: kelp cells that were more exposed to vessel traffic before 2009 should benefit more from the implementation of the regulated zone if that regulation effectively reduces nearshore stressors. The treatment intensity is measured by pre-policy shipping activity derived from AIS data: logged mean vessel traffic (measured as total distance traveled) or logged mean vessel count within a 20 km buffer around each kelp hexagon during January–June 2009, before implementation.

We estimate:

$$Kelp_{it} = \exp(\beta\, Exposure_i \times Post2009_t + \delta\, Z_{it} + f_i + f_t) \cdot \eta_{it}, \tag{1}$$

where η_{it} is a multiplicative error term satisfying $E[\eta_{it} \mid Exposure_i, Z_{it}, f_i, f_t] = 1$, and $Kelp_{it}$ is the sum of kelp biomass (kg) or canopy area (m^2) in hexagonal cell i during quarter t. $Exposure_i$ is the time-invariant, predetermined measure of shipping activity described above (logged mean vessel traffic or logged mean vessel count, January–June 2009). $Post2009_t$ equals one from 2009Q3 onward and zero otherwise. Z_{it} is a vector of log-transformed environmental controls (log sea surface temperature, log precipitation, and log significant wave height).[9] f_i denotes kelp-cell fixed effects (one indicator per hexagonal cell) and f_t denotes year–quarter fixed effects. The cell fixed effects absorb any time-invariant determinant of kelp outcomes—including the main effect of $Exposure_i$—and the time fixed effects absorb shocks common to all cells, including the main effect of $Post2009_t$.

Because kelp outcomes are non-negative and include many zeros, we estimate Eq. (1) by Poisson pseudo-maximum likelihood (PPML; Santos Silva and Tenreyro, 2006).[10] Standard errors are clustered at the cell level (81 clusters).[11] The sample covers Southern California (latitudes 32°–36°N) from 2005 to 2016 [12]. The year–quarter fixed effects absorb any

[9] All controls are transformed as $\ln(1+x)$. For SST and significant wave height, which are strictly positive and large in magnitude, this is virtually identical to $\ln(x)$. For precipitation, the $\ln(1+x)$ form handles quarters with zero rainfall.

[10] We restrict the sample to cells recording at least two vessel passages per month on average during January–June 2009, excluding cells with negligible shipping exposure. Results are robust to alternative cutoffs (Figure 7) and to weighting by log vessel count (Appendix B).

[11] Because kelp cells are hexagons along a narrow coastal strip, adjacent cells may share ocean currents and temperature fields, inducing spatial correlation within clusters. Conley standard errors (Conley, 1999) which allow for spatial dependence decaying with distance, are the natural remedy in OLS settings but are not available for PPML: the Conley estimator relies on a sandwich formula for the linear GMM moment conditions, which does not directly extend to the nonlinear PPML first-order conditions. Implementing spatial HAC inference for Poisson models would require either a nonlinear GMM framework or a spatial bootstrap, neither of which is standard in the literature. We note this as a limitation and a direction for future methodological work.

[12] This geographic restriction also addresses the species composition of kelp forests along the Pacific coast. Giant kelp (*Macrocystis pyrifera*) dominates the Southern California coast and forms the detectable surface canopy measured by Landsat imagery (Bell et al., 2023). North of approximately 36–37°N, bull kelp (*Nereocystis luetkeana*) progressively replaces giant kelp as the dominant canopy species (Dayton, 1985; Steneck et al., 2002). Because these two species differ in growth form, canopy structure, and sensitivity to environmen-

coast-wide ecological shock, including the 2013–2016 marine heat wave that caused catastrophic kelp losses in northern California (Rogers-Bennett and Catton, 2019; McPherson et al., 2021). Such a shock would bias β only if its spatial footprint were systematically correlated with the pre-policy shipping-exposure cross-section, for which there is no ecological basis.

The coefficient of interest is β. A positive β estimate would indicate that kelp outcomes improved more strongly after the policy in cells that had been more exposed to vessel traffic beforehand — suggesting that the regulation was most beneficial precisely where maritime activity had been most intense.

This continuous design offers two advantages over a binary treated-versus-control comparison. First, it does not require selecting a specific control group outside the regulated zone; identification comes from variation in pre-policy exposure *within* the regulated area. Second, using a continuous treatment intensity exploits the full cross-sectional variation in exposure rather than collapsing it into a binary indicator, yielding a more powerful test of whether kelp forests that were more exposed to maritime traffic recover disproportionately after the implementation of the ECA.

As a complementary exercise, Appendix B reports a specification replacing our AIS-based exposure with the cumulative human impact index of Halpern et al. (2009) described in Section 3.1. Although the coefficient is also positive and significant, event-study diagnostics reveal pre-existing differential trends driven by non-shipping stressors bundled in the index (urbanization, fishing pressure, nutrient runoff). This motivates our preference for shipping-specific measures, which isolate the regulatory channel.

3.3 Results

Table 1 reports estimates of the continuous difference-in-differences using pre-policy shipping exposure as treatment intensity. Columns (1)–(2) measure exposure by total distance traveled by vessels ("Distance Traveled (km)"), with biomass and canopy area as dependent variables respectively. Columns (3)–(4) measure exposure by the number of vessel passages ("Vessel Count").

tal stressors (Carr and Reed, 2019), restricting the sample to 32°–36°N ensures that our analysis captures a spatially consistent kelp assemblage dominated by *Macrocystis pyrifera*

	Distance Traveled (km)		Vessel Count	
	(1) Biomass	(2) Area	(3) Biomass	(4) Area
Post 2009 × Exposure	0.1361***	0.1317***	0.1548***	0.1493***
	(0.0359)	(0.0352)	(0.0432)	(0.0424)
SST	-5.703***	-5.512***	-5.728***	-5.537***
	(1.145)	(1.136)	(1.150)	(1.142)
Precipitation	0.0764	0.0812	0.0846	0.0891
	(0.1003)	(0.0981)	(0.1012)	(0.0989)
Significant Wave Height	-3.267**	-3.102**	-3.358**	-3.189**
	(1.367)	(1.344)	(1.367)	(1.342)
Cell FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes
Observations	3880	3880	3880	3880

Notes: Unweighted PPML estimates of Eq. (1). The dependent variable is kelp biomass (kg) or canopy area (m^2). "Distance Traveled (km)" columns (1)–(2) measure exposure by logged mean distance traveled by vessels within a 20 km buffer around each kelp hexagon during January–June 2009; "Vessel Count" columns (3)–(4) use the corresponding logged mean vessel count. All specifications include kelp-cell and year–quarter fixed effects. Sample: Southern California (32°–36° N), 2005–2016, restricted to cells with at least two vessel passages per month on average (81 cells × 48 quarters, less 8 cell-quarters with missing covariates in 2015–2016 = 3,880 observations). Standard errors clustered at the cell level (81 clusters) in parentheses. Weighted estimates are reported in Appendix B.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1: Continuous difference-in-differences: Kelp outcomes and the 2009 ECA

Across all specifications, the interaction between post-2009 implementation and pre-policy shipping exposure is positive and statistically significant at the 1% level. The estimates are stable across treatment-intensity measures and dependent variables.

In columns (1)–(2), which use distance traveled by vessels, the coefficient of approximately 0.136 implies that a one-unit increase in log pre-policy distance traveled is associated with an $\exp(0.136) - 1 \approx 14.6\%$ larger post-2009 increase in expected kelp biomass and canopy. In columns (3)–(4), the coefficient on vessel count is approximately 0.155, implying a roughly 16.7% relative increase per unit of log vessel count. The consistency of results across cumulative distance and number of passages strengthens robustness: the former captures the cumulative intensity of emissions and discharges, the latter the frequency of discrete exposure events, and both yield a similar estimate of the post-policy effect.

Environmental controls enter with expected signs. Sea surface temperature is strongly and negatively associated with kelp outcomes across all columns, consistent with the well-documented sensitivity of kelp to warm-water conditions (Carr and Reed, 2019; Rogers-Bennett and Catton, 2019). Significant wave height is also negatively associated with kelp, consistent with physical damage to fronds and holdfasts during strong sea states. Precipitation is positively associated with kelp, possibly reflecting increased terrestrial nutrient runoff that can support kelp growth under moderate conditions.

3.4 Dynamic effects

To assess whether the parallel-trends assumption holds and to trace the evolution of treatment effects over time, we estimate an event-study specification:

$$Kelp_{it} = \exp\left(\sum_{\ell \neq 2008} \beta_\ell \, Exposure_i \times \mathbf{1}_{t \in \ell} + \delta \, Z_{it} + f_i + f_t \right) \cdot \eta_{it}, \tag{2}$$

where ℓ indexes event-study periods and $\mathbf{1}_{t \in \ell}$ equals one when quarter t falls in period ℓ. Each calendar year constitutes a single period, except 2009, which is split into H1 (Q1–Q2, pre-enforcement) and H2 (Q3–Q4, post-enforcement). The reference period is 2008. This specification interacts the time-invariant treatment intensity with period indicators, allowing us to trace out the differential evolution of kelp outcomes across cells with different levels of shipping exposure.

Figure 1 reports the estimated coefficients $\hat{\beta}_\ell$ for distance traveled interacted with period indicators, estimated by PPML with the same controls and fixed effects as Eq. (1).

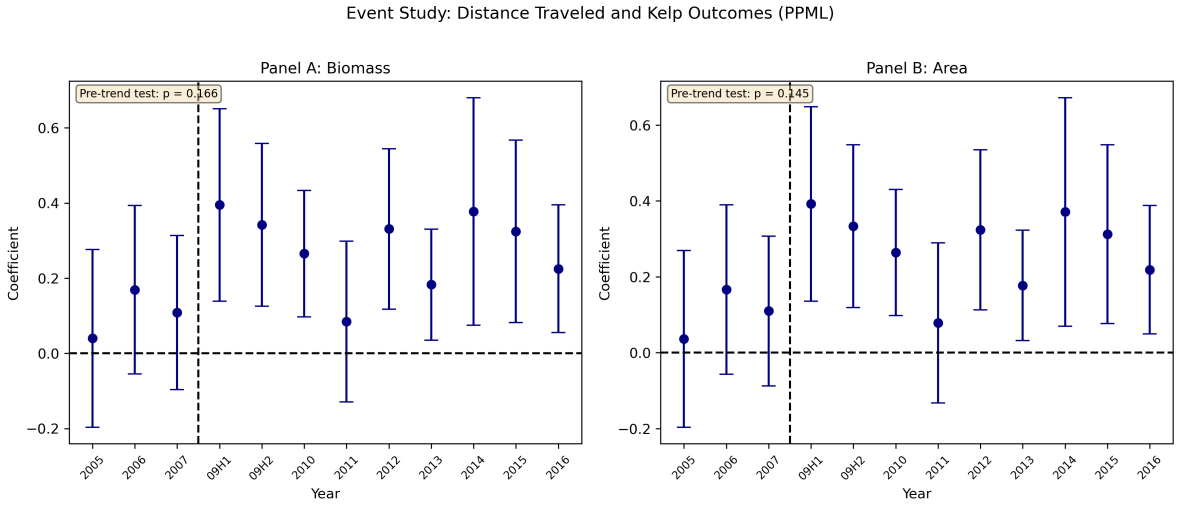

Note: Each panel plots the coefficients from Eq. (2), where distance traveled (logged mean, January–June 2009) is interacted with year indicators, estimated by PPML. The reference year is 2008. The vertical dashed line marks the start of enforcement (2009Q3). Error bars are 95% confidence intervals based on standard errors clustered at the cell level. Panel A: kelp biomass; Panel B: kelp canopy area.

Figure 1: Event study: Distance traveled and kelp outcomes (PPML)

Two features of Figure 1 merit attention. First, all pre-policy coefficients (2005, 2006, and 2007) are close to zero and statistically insignificant, and a joint Wald test fails to reject the null of no pre-trend, though the power of this test is limited with only three pre-treatment periods. Second, the coefficients turn positive and significant starting in 2009H1—the six months immediately before enforcement—and remain so through the end of the sample (2016), with a temporary dip in 2011. The fact that the effect emerges already in 2009H1 is consistent with some operators switching to compliant fuel during

the compliance ramp-up period preceding the July 2009 enforcement date, though mild anticipation or a coincident shock cannot be fully ruled out. The coefficients are similar in magnitude in 2009H1 and 2009H2, suggesting a relatively smooth onset rather than a discrete jump at enforcement. A placebo test setting the treatment date at 2007Q3 on the pre-policy sample alone confirms no effect (Appendix C). The next subsection examines what these patterns imply for identification.

3.5 Threats to identification

The identifying assumption behind Eq. (1) is that, absent the policy, kelp cells with different levels of pre-policy shipping exposure would have followed parallel trajectories. We now assess this assumption in light of the results.

Pre-trends. As documented above, the event study (Figure 1) shows no systematic differential pre-trend (joint Wald test $p = 0.17$; see also Rambachan and Roth, 2023 on the interpretation of pre-trend tests). Appendix D reports a balance table showing that pre-treatment environmental covariates (SST, precipitation) are comparable across treatment-intensity quartiles, while kelp outcomes and wave height differ—consistent with the geographic concentration of shipping near specific coastal segments. A placebo test using a fake treatment date reinforces this conclusion (Appendix C).

Traffic response to the policy. Our identification strategy treats shipping exposure as a predetermined intensity measure and tests whether kelp outcomes improve differentially in more-exposed cells after the policy date. A natural question is whether the policy actually altered vessel behavior. Klotz and Berazneva (2022) provide direct, voyage-level evidence. Using one-minute AIS data for roughly 20,000 voyages per year off the U.S. west coast, they document a sharp discontinuity at the July 2009 enforcement date: container ships reduced fuel consumption within the ECA by 33–52%, primarily by rerouting outside the regulated boundary—exiting the zone as quickly as possible and traveling south of the Channel Islands rather than through the Santa Barbara Channel. Speed reductions within the ECA reinforced the effect. The 2011 boundary modification around the Channel Islands partially reversed the avoidance pattern, drawing some traffic back into the channel. These voyage-level behavioral shifts are precisely the variation in shipping intensity that our cell-level treatment variable captures, and they confirm that the policy generated a meaningful change in nearshore maritime activity at both dates exploited in our analysis.

Financial crisis. The most salient contemporaneous shock is the 2008–09 global financial crisis, which reduced trade volumes and port throughput at Southern California ports by roughly 20% in 2009. If reduced port-related economic activity differentially improved conditions in high-shipping cells through channels unrelated to the ECA (e.g., reduced industrial discharge or ancillary vessel traffic), it could confound our estimates. Appendix E addresses this concern directly by constructing a cell-level measure of port economic activity from the National Ballast Information Clearinghouse (NBIC) vessel arrival database

and adding it as a time-varying control. The shipping-exposure coefficient is virtually unchanged (at most 2% attenuation). This coefficient stability when controlling for port activity is the most direct evidence that the crisis does not confound the ECA effect.

Anticipation. Because the regulation was announced before the July 2009 enforcement date, vessel operators may have begun rerouting in the preceding months, potentially contaminating our January–June 2009 treatment measure. There are strong economic reasons to discount this concern. The regulation requires the use of low-sulfur marine fuel within the ECA, which is substantially more expensive than the heavy fuel oil used outside it (Klotz and Berazneva, 2022). Switching to compliant fuel before enforcement was required would raise operating costs with no offsetting benefit, since penalties for noncompliance did not apply until July 2009. Vessel operators therefore had a clear incentive to maintain pre-policy routing patterns until the enforcement date. The event study is broadly consistent with this reasoning: the 2009H1 coefficient is positive and statistically significant, but similar in magnitude to 2009H2, suggesting a smooth onset of the effect during the compliance ramp-up period rather than anticipatory rerouting. This pattern is compatible with some operators beginning to use low-sulfur fuel in the months before mandatory enforcement, as documented in comparable regulatory settings. The pre-policy coefficients (2005–2007) remain close to zero—the break occurs precisely at the transition to the enforcement year, not gradually over the pre-period. Appendix F provides further evidence, showing that the cross-sectional ranking of cells by shipping intensity is highly stable in years unaffected by the 2011 boundary modification (Spearman $\rho \geq 0.94$ for 2009, 2012–2014), with a temporary dip in 2010–2011 (ρ between 0.63 and 0.76) driven by the traffic reshuffling around the Channel Islands. Replacing the baseline treatment with 2012 AIS data yields similar coefficients. These patterns confirm that the treatment variable captures stable geographic variation in shipping intensity rather than transient noise.

Spillovers. If the policy redistributed vessel traffic from high-exposure to low-exposure cells, kelp improvements in treated cells could partly reflect worsened conditions in nearby untreated cells, violating SUTVA. Two features of our design mitigate this concern. First, the 10 km hexagonal aggregation internalizes local physical spillovers (e.g., pollution or turbidity drifting between adjacent pixels) within each unit. Second, for the broader rerouting-induced interference—where traffic displaced from one corridor reappears in another—the complementary exercise described next, which uses an external control group, is immune to this concern by construction.

External validation. We subject the main findings to a complementary exercise that relies on an entirely different identification strategy. Following Arkhangelsky et al. (2021), we construct a synthetic counterpart for treated California hexagons using a donor pool of Mexican coastal hexagons. The estimated post-2009 effect is positive and significant, and the result holds whether we include or exclude southern California hexagons from the treated sample. Full details—sample construction, donor weights, and the synthetic-versus-treated trajectory—are reported in Appendix G. Because the continuous DiD and

the synthetic DiD differ in spatial unit, control group, and identifying variation, their convergence strengthens confidence that the pattern reflects a genuine policy effect rather than an artifact of a single design choice.

4 Boundary Modification of the Regulated Zone in 2011

This section exploits the 2011 boundary modification around the Channel Islands to provide a second, mirror-image test of the relationship between maritime exposure and kelp health. While Section 3 showed that kelp improved after 2009 in cells more exposed to pre-policy vessel traffic, the 2011 episode allows us to test the converse: whether kelp *deteriorated* in cells more exposed to shipping when the boundary change re-concentrated traffic in specific corridors. Applying the same continuous difference-in-differences logic to a policy change that *increased* rather than decreased local maritime pressure provides a partially independent test that strengthens the causal interpretation of both results.

4.1 Spatial redistribution of traffic after 2011

This section relies on the same AIS-derived shipping measures (vessel count and distance traveled) described in Section 3 and assigned to kelp hexagons via the buffer procedure defined above.

The 2011 boundary modification around the Channel Islands altered where vessels concentrate as they approach Southern California ports and contributed to a reallocation of traffic toward specific corridors (Klotz and Berazneva, 2022). Figures 2 and 3 illustrate the baseline spatial distribution of traffic exposure in 2009 and 2011, respectively; Appendix H documents the post-modification redistribution directly.

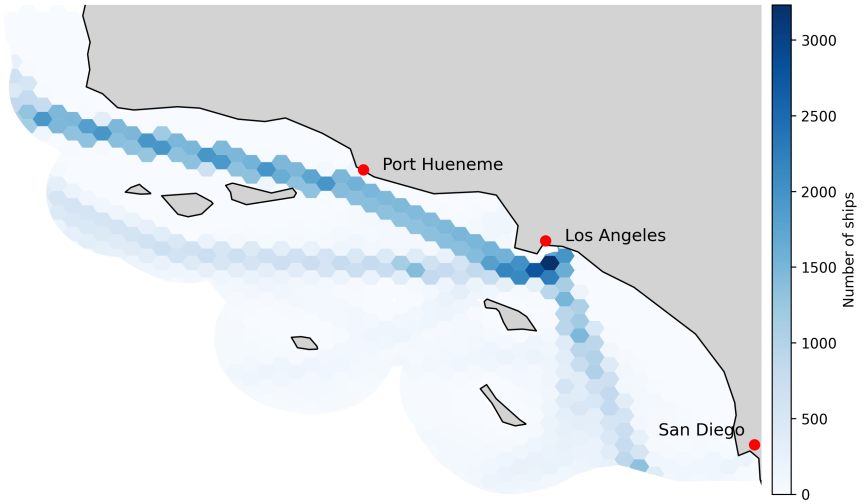

Notes: Number of vessel passages through each 10 km hexagonal cell in Southern California during 2009. Darker shading indicates higher traffic intensity. Data source: AIS.

Figure 2: Exposure of kelp to maritime traffic in 2009

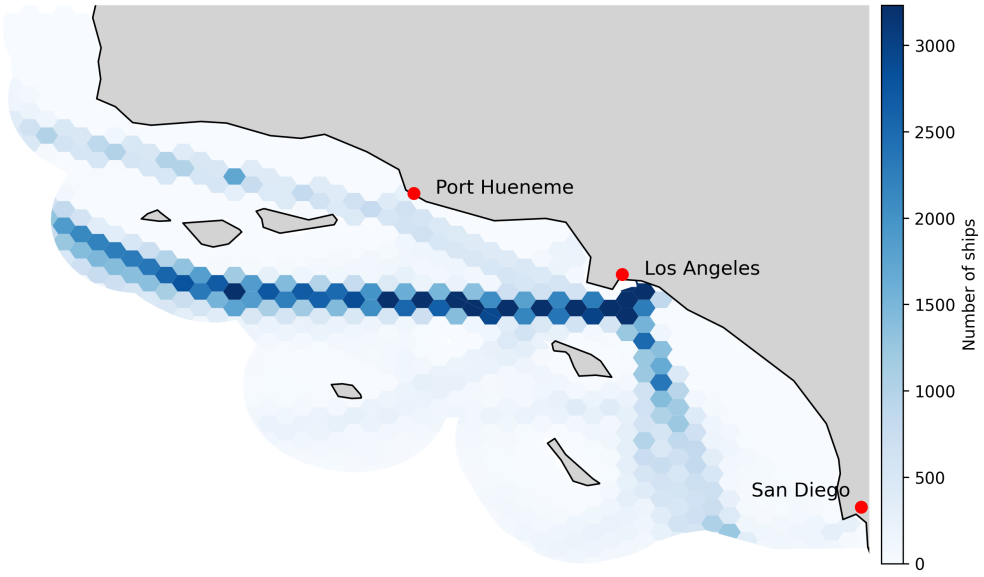

Notes: Number of vessel passages through each 10 km hexagonal cell in Southern California during 2011. The boundary modification took effect in December 2011, so the map overwhelmingly reflects pre-modification routing. Data source: AIS.

Figure 3: Exposure of kelp to maritime traffic in 2011

Appendix H provides direct evidence of this redistribution. AIS cell-level data show a shift in total vessel distance between the 2009 zone and the boundary modification area (Figure 14), consistent with the route adjustments documented in Klotz and Berazneva (2022).

4.2 Empirical strategy

To exploit the 2011 boundary modification, we implement a continuous difference-in-differences design that parallels Eq. (1). As in the 2009 specification, the treatment intensity is measured by shipping activity—now from the pre-modification period. The logic is analogous: kelp cells that were more exposed to vessel traffic before the 2011 boundary change should experience worse outcomes afterward if the reallocation of traffic into those corridors increased nearshore stressors.

Specifically, we measure shipping exposure over the full calendar year 2011 using total distance traveled and vessel count within a 20 km buffer around each kelp hexagon. Because the boundary modification took effect in December 2011—the last month of the year—eleven out of twelve months in our treatment variable reflect pre-modification routing patterns. The remaining month contributes roughly 8% of the annual average: NBIC vessel arrival records for California ports over 2005–2016 show that December ranks sixth out of twelve months, with a traffic share (8.4%) virtually identical to the uniform benchmark (8.3%). Any contamination from post-modification routing in December would blur the spatial contrast between high- and low-exposure cells—mixing pre- and post-modification patterns—thereby attenuating γ toward zero rather than inflating it. The 2011 treatment variable is therefore overwhelmingly predetermined and, to the extent it is not, the resulting bias works against our finding. As a direct check, Appendix I replaces the 2011 treatment with shipping exposure measured during 2010—a year entirely preceding the boundary modification. All coefficients remain negative and significant at the 1% level, confirming that the results are not driven by the inclusion of December 2011 in the treatment variable. Appendix J documents the correlation between the 2009 and 2011 treatment variables, confirming that the boundary modification generated partially independent spatial variation in shipping exposure.

We estimate:

$$Kelp_{it} = \exp\Big(\gamma\, Exposure_i^{2011} \times Post2011_t + \delta\, Z_{it} + f_i + f_t\Big) \cdot \eta_{it}, \tag{3}$$

where $Kelp_{it}$ is the sum of kelp biomass (kg) or canopy area (m^2) in hexagonal cell i during quarter t. $Exposure_i^{2011}$ is the log of pre-modification shipping intensity measured during 2011 (total distance traveled or vessel count within a 20 km buffer zone). $Post2011_t$ equals one from 2012Q1 onward and zero otherwise (the boundary modification took effect in December 2011, so 2011 quarters are coded as pre-treatment). Z_{it} is defined as in Eq. (1). f_i denotes kelp-cell fixed effects (defined as in Eq. (1)) and f_t year–quarter fixed effects.

As in Eq. (1), the cell fixed effects absorb the main effect of $Exposure_i^{2011}$ and any other time-invariant determinant of kelp outcomes; the year–quarter fixed effects absorb common shocks, including the main effect of $Post2011_t$. The equation is estimated by PPML,

as in the main 2009 specification. Standard errors are clustered at the cell level. The sample covers Southern California (latitudes 32°–36°N) from 2005 to 2016, matching the 2009 specification. The year–quarter fixed effects absorb any coast-wide ecological shock, including the 2013–2016 warming episode; such a shock would bias γ only if its spatial footprint were systematically correlated with the 2011 shipping-exposure cross-section, for which there is no ecological basis.

The coefficient of interest is γ. A negative γ indicates that, after the 2011 boundary modification, kelp outcomes deteriorated disproportionately in cells that were more exposed to shipping before the policy change—consistent with the reallocation of traffic into those corridors worsening local maritime pressures.

This specification offers a key advantage over an alternative approach that would interact time-varying shipping exposure with a post-2011 indicator: the predetermined nature of $Exposure_i^{2011}$ eliminates the concern that contemporaneous vessel routing responds to unobserved marine conditions that also affect kelp. The design thus shares the same identification logic as our main 2009 specification, providing internal consistency across the two natural experiments.

4.3 Results

Table 2 reports estimates of Eq. (3). Columns (1)–(2) measure exposure by total distance traveled by vessels ("Distance Traveled (km)"), with biomass and canopy area as dependent variables respectively. Columns (3)–(4) measure exposure by vessel count ("Vessel Count").

	Distance Traveled (km)		Vessel Count	
	(1) Biomass	(2) Area	(3) Biomass	(4) Area
Post 2011 × Exposure	-0.0925***	-0.0914***	-0.210***	-0.208***
	(0.0288)	(0.0282)	(0.0433)	(0.0425)
SST	-7.155***	-7.060***	-7.103***	-7.008***
	(0.904)	(0.909)	(0.903)	(0.908)
Precipitation	-0.0359	-0.0386	-0.0483	-0.0508
	(0.0781)	(0.0773)	(0.0782)	(0.0776)
Significant Wave Height	-3.452***	-3.371***	-3.341***	-3.262***
	(0.958)	(0.951)	(0.956)	(0.949)
Cell FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes
Observations	6000	6000	6000	6000

Notes: Unweighted PPML estimates of Eq. (3). The dependent variable is kelp biomass (kg) or canopy area (m^2). "Distance Traveled (km)" columns (1)–(2) measure exposure by logged total distance traveled by vessels within a 20 km buffer around each kelp hexagon during 2011; "Vessel Count" columns (3)–(4) use the corresponding logged vessel count. All specifications include kelp-cell and year–quarter fixed effects. Standard errors clustered at the cell level (125 clusters) in parentheses. Sample: Southern California (32°–36° N), 2005–2016, 6,000 observations (125 cells × 48 quarters). Weighted estimates, which are qualitatively identical, are reported in Appendix K.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Continuous difference-in-differences: Kelp outcomes and the 2011 boundary modification

Across all specifications, the interaction between post-2011 implementation and pre-modification shipping exposure is negative and statistically significant at the 1% level. The estimates are stable across treatment-intensity measures and dependent variables.

In columns (1)–(2), which use logged distance traveled by vessels, the coefficient of approximately -0.093 implies that a one-unit increase in log pre-2011 distance traveled is associated with a $1 - \exp(-0.093) \approx 8.9\%$ larger post-2011 decrease in expected kelp biomass and canopy. In columns (3)–(4), the coefficient on logged vessel count is approximately -0.210, implying a $1 - \exp(-0.210) \approx 18.9\%$ relative decline per unit of log vessel count.

These results are the mirror image of the 2009 findings in Table 1. Where the 2009 regulation was associated with disproportionate kelp improvement in high-exposure cells ($\hat{\beta} > 0$), the 2011 boundary modification was associated with disproportionate kelp deterioration in cells exposed to re-concentrated traffic ($\hat{\gamma} < 0$). The symmetry of these two findings—one following a regulation that tightened fuel standards for nearshore shipping, the other following a boundary change that redirected traffic into new corridors—strengthens the causal interpretation: maritime regulation has a quantitatively important and plausibly causal effect on kelp forest health, mediated by changes in the spatial footprint of vessel activity.

Environmental controls enter with expected signs and are consistent with the esti-

mates in Table 1. Sea surface temperature is strongly and negatively associated with kelp outcomes (-7.2 to -7.0), confirming kelp sensitivity to warm-water conditions. Significant wave height is also strongly negative (-3.5 to -3.3), reflecting physical stress from storms. Precipitation is positive but statistically insignificant.

4.4 Dynamic effects

A standard PPML event study for the 2011 boundary modification faces a specific identification challenge: the 2009 ECA regulation, which constitutes the main treatment in the first part of this paper, falls within the pre-treatment window for the 2011 event. Because both treatments rely on cross-sectional variation in shipping intensity, a conventional two-way fixed-effects event study cannot disentangle the 2009 and 2011 effects—the pre-2011 coefficients would reflect a mix of the 2009 treatment dynamics and genuine pre-trends for the 2011 shock.

To address this, we implement the heterogeneity-robust estimator of Chaisemartin and D'Haultfœuille (2024), which accommodates continuous (non-binary) treatments. The treatment variable D_{it}^{2011} equals zero for all cells before 2011 and equals the log of mean vessel distance traveled within a 20 km buffer (measured in 2011) from 2011 onward. The key identifying assumption is that, in the absence of the 2011 shock, the evolution of kelp outcomes would have been linear in the cross-sectional shipping intensity. The estimator fits this baseline trend from the earliest pre-treatment period and uses it to project counterfactual outcomes at each event time; deviations from this projection are the estimated treatment effects. Because the treatment variable is identically zero before 2011, the cross-sectional variation that identifies the baseline polynomial comes from the relationship between shipping intensity and kelp trends during the pre-treatment years, not from variation in the treatment itself.

The panel is collapsed to annual frequency to avoid seasonal heteroscedasticity in quarterly kelp data—winter quarters (Q1, Q4) exhibit roughly ten times smaller variance than spring/summer quarters (Q2, Q3), which would distort the apparent precision of quarterly event-study coefficients. The sample covers 2008–2016, with the first pre-treatment year (2008) used to fit the baseline polynomial, leaving two pre-treatment placebos (2009, 2010) and six post-treatment effects (2011–2016). Standard errors are obtained by bootstrap (999 replications) clustered at the cell level.

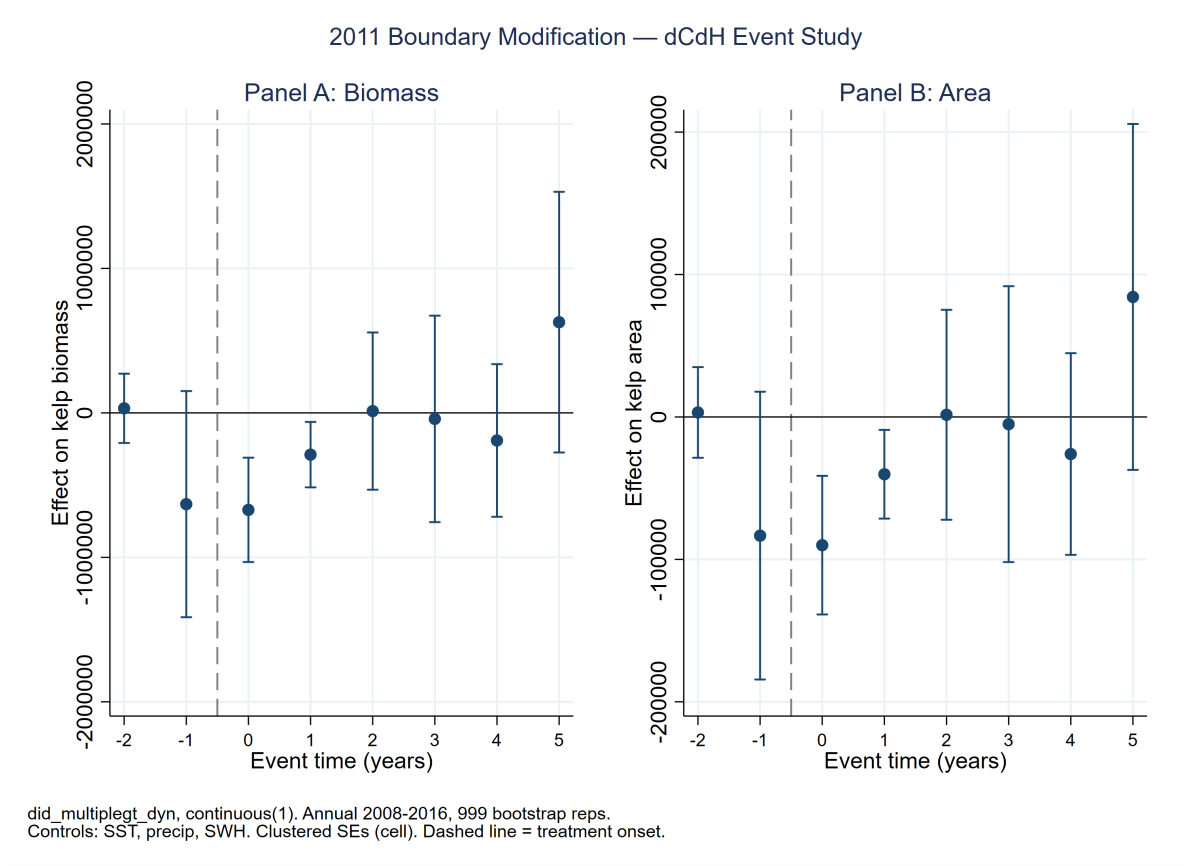

Note: Each panel plots heterogeneity-robust treatment effects from the `did_multiplegt_dyn` estimator of Chaisemartin and D'Haultfœuille (2024) with continuous treatment intensity (linear polynomial). The treatment variable is zero before 2011 and equals log mean vessel distance within a 20 km buffer from 2011 onward. The panel is annual (2008–2016). Controls: log SST, log precipitation, log significant wave height. Error bars are 95% bootstrap confidence intervals (999 replications, clustered at the cell level). Panel A: kelp biomass; Panel B: kelp canopy area. Dashed line marks treatment onset.

Figure 4: Event study: 2011 boundary modification (de Chaisemartin–D'Haultfœuille estimator)

Figure 4 presents the results. Both pre-treatment placebos (event times -2 and -1, corresponding to 2009 and 2010) are statistically insignificant, with confidence intervals comfortably crossing zero, supporting the parallel-trends assumption for the 2011 shock even in the presence of the concurrent 2009 regulation.[13] Following the 2011 boundary modification, the treatment effect turns sharply negative at event times 0 and 1, indicating that cells more exposed to shipping experienced a deterioration in kelp outcomes. By event times 2–5 (2013–2016), the point estimates gradually return toward zero, indicating that the negative effect was transitory. This pattern is consistent with kelp's well-documented capacity for rapid regrowth—giant kelp (*Macrocystis pyrifera*) can grow up

[13]The event time -1 point estimate is negative in the biomass panel, but the wide confidence interval ($p > 0.10$) is consistent with sampling variation rather than a pre-trend.

to 30 cm per day under favorable conditions—consistent with partial ecological recovery, though other time-varying factors may also contribute.

This dynamic pattern complements the static PPML estimates in Table 2. While the static specification pools all post-2011 periods and estimates the average effect under PPML, the event study traces the temporal evolution under OLS with the heterogeneity-robust corrections of Chaisemartin and D'Haultfœuille (2024). The transitory nature of the effect implies that the static estimate, which averages over both the initial damage and subsequent recovery, understates the short-run impact of the boundary modification. Nevertheless, the convergence of both approaches—different estimators, different functional forms—reinforces the conclusion that the 2011 boundary modification worsened kelp outcomes in high-shipping cells, even if the ecosystem partially recovered.

The transitory nature of the 2011 effect contrasts with the persistent 2009 effect, which shows no attenuation through 2016. This asymmetry has a natural biological explanation. The transitory 2011 effect reflects kelp's capacity for rapid regrowth: once the initial damage from re-concentrated traffic is absorbed, giant kelp can recolonize, closing the gap between high- and low-exposure cells within a few years. The persistent 2009 effect, by contrast, is consistent with a sustained change in nearshore conditions: as long as the regulation remains in force, the conditions favoring kelp recovery in formerly high-exposure cells persist, so the positive differential does not attenuate.

5 Conclusion

> "If in any country a forest was destroyed, I do not believe as many species of animals would perish as would here from the destruction of kelp." — Darwin (1839)

Human reliance on kelp dates back to prehistoric times; Erlandson et al. (2007) even propose a "Kelp Highway Hypothesis" explaining that human migration 20,000 years ago depended on this resource. Yet these same kelp corridors are now traversed by giant vessels that leave maritime pollution in their wake.

In this paper, we have exploited two policy episodes—the 2009 Ocean-Going Vessel Fuel Rule and the 2011 boundary modification around the Channel Islands—to estimate the causal effect of maritime regulation on kelp forests along the California coast, operating through differential shipping exposure. Using a continuous difference-in-differences design with shipping-specific exposure measures derived from AIS vessel trajectories, we show that (i) kelp canopy and biomass improved disproportionately in cells more exposed to vessel traffic after the 2009 regulation tightened fuel standards for nearshore vessels, and (ii) kelp deteriorated disproportionately in cells more exposed to shipping before the 2011 boundary change re-concentrated traffic in specific corridors. Event-study diagnostics show approximately flat pre-trends for shipping-specific measures in the 2009 specification (Figure 1), while broader cumulative impact measures exhibit significant pre-trends driven by non-shipping confounders. For the 2011 episode, a heterogeneity-robust event study (Chaisemartin and D'Haultfœuille, 2024) likewise shows no pre-trend despite the concurrent 2009 regulation (Figure 4), and reveals that the negative effect was

transitory—kelp partially recovered within two to three years, consistent with the rapid regrowth capacity of giant kelp. The symmetry of the 2009 and 2011 results—one following a tightening of fuel standards, the other following a spatial reallocation of traffic—strengthens the causal interpretation that maritime activity has a quantitatively important effect on kelp forest health, operating through channels that covary with the spatial footprint of maritime activity, which the regulation altered through both fuel switching and route adjustments.

These findings have implications for both economic analysis and maritime policy. First, as the international trade literature now examines the "hidden costs" of globalization, we bring kelp forests into this discussion: they play an important role in fisheries, alginate production, and carbon sequestration in coastal ecosystems. An illustrative back-of-the-envelope calculation suggests that the ECA-induced kelp improvement in our study area is worth on the order of $59–67 million per year in ecosystem services—a figure not currently reflected in cost-benefit assessments of maritime emission regulation.[14]

Second, this study supports the objectives of the International Maritime Organization (IMO), which has been pivotal in enforcing emission control areas over the past decade. In settings where regulated zones overlap with nearshore ecosystems, ECAs may have effects beyond their impact on air pollution, and could play a role in promoting sustainable international trade that limits damage to coastal habitats. At the same time, the 2011 episode serves as a cautionary finding: boundary modifications that re-concentrate traffic can undo some of the environmental benefits of the original regulation, suggesting that spatial design matters as much as the stringency of emission standards.

References

Abman, Ryan and Clark Lundberg (2020). "Does Free Trade Increase Deforestation? The Effects of Regional Trade Agreements". In: *Journal of the Association of Environmental and Resource Economists* 7.1, pp. 35–72. ISSN: 2333-5963. DOI: 10.1086/705787.

Ambrose, RF and Bobette V Nelson (1982). "Inhibition of giant kelp recruitment by an introduced brown alga". In.

Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager (2021). "Synthetic Difference in Differences". In: *American Economic Review* 111.12, pp. 4088–4118. DOI: 10.1257/aer.20190159.

Bell, Tom W., Kyle C. Cavanaugh, Vienna R. Saccomanno, Katherine C. Cavanaugh, Henry F. Houskeeper, Norah Eddy, Falk Schuetzenmeister, Nathaniel Rindlaub, and Mary Gleason (2023). "Kelpwatch: A new visualization and analysis tool to explore kelp

[14] We apply the ecosystem service valuations of Eger et al. (2023), who estimate that *Macrocystis* kelp forests provide approximately $72,000/ha/year through fisheries production, nutrient cycling, and carbon sequestration. Our study area contains roughly 5,500 hectares of kelp canopy (pre-period annual mean). Multiplying by the approximately 15–17% improvement implied by the main estimates yields approximately $59–67 million per year. This calculation is approximate: it assumes that the average treatment intensity across cells maps one-to-one onto the semi-elasticity, and it relies on global average per-hectare valuations rather than California-specific estimates.

canopy dynamics reveals variable response to and recovery from marine heatwaves". In: *PLOS ONE* 18.3. Ed. by Alejandro Pérez-Matus, e0271477. DOI: 10.1371/journal.pone.0271477.

Berman, Nicolas, Mathieu Couttenier, Antoine Leblois, and Raphael Soubeyran (2023). "Crop prices and deforestation in the tropics". In: *Journal of Environmental Economics and Management* 119, p. 102819. ISSN: 0095-0696. DOI: 10.1016/j.jeem.2023.102819.

Birch, C.P.D., S.P. Oom, and J.A. Beecham (2007). "Rectangular and hexagonal grids used for observation, experiment and simulation in ecology". In: *Ecological Modelling* 206.3–4, pp. 347–359. DOI: 10.1016/j.ecolmodel.2007.03.041.

Britton-Simmons, KH (2004). "Direct and indirect effects of the introduced alga Sargassum muticum on benthic, subtidal communities of Washington State, USA". In: *Marine Ecology Progress Series* 277, pp. 61–78. ISSN: 1616-1599. DOI: 10.3354/meps277061.

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

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller (2008). "Bootstrap-Based Improvements for Inference with Clustered Errors". In: *Review of Economics and Statistics* 90.3, pp. 414–427. DOI: 10.1162/rest.90.3.414.

Cameron, A. Colin and Douglas L. Miller (2015). "A Practitioner's Guide to Cluster-Robust Inference". In: *Journal of Human Resources* 50.2, pp. 317–372. DOI: 10.3368/jhr.50.2.317.

Carr, Mark H. and Daniel C. Reed (2019). "Ecosystems of California". In: *Ecosystems of California*. Ed. by Harold Mooney and Erika Zavaleta. Berkeley: University of California Press. Chap. Shallow Rocky Reefs and Kelp Forests, pp. 311–336. DOI: 10.1525/9780520962170-021.

Carreira, Igor, Francisco Costa, and João Paulo Pessoa (2024). "The deforestation effects of trade and agricultural productivity in Brazil". In: *Journal of Development Economics* 167, p. 103217. ISSN: 0304-3878. DOI: 10.1016/j.jdeveco.2023.103217.

Castorani, Max C. N., Tom W. Bell, Jonathan A. Walter, Daniel C. Reuman, Kyle C. Cavanaugh, and Lawrence W. Sheppard (2022). "Disturbance and nutrients synchronise kelp forests across scales through interacting Moran effects". In: *Ecology Letters* 25.8, pp. 1854–1868. ISSN: 1461-0248. DOI: 10.1111/ele.14066.

Chaisemartin, Clément de and Xavier D'Haultfœuille (2024). "Difference-in-Differences Estimators of Intertemporal Treatment Effects". In: *Review of Economic Studies*. DOI: 10.1093/restud/rdae008.

Christie, H, KM Norderhaug, and S Fredriksen (2009). "Macrophytes as habitat for fauna". In: *Marine Ecology Progress Series* 396, pp. 221–233. DOI: 10.3354/meps08351.

Cima, Francesca and Roberta Varello (2022). "Potential disruptive effects of copper-based antifouling paints on the biodiversity of coastal macrofouling communities". In: *Environmental Science and Pollution Research* 30.4, pp. 8633–8646. DOI: 10.1007/s11356-021-17940-2.

Conley, Timothy G. (1999). "GMM Estimation with Cross Sectional Dependence". In: *Journal of Econometrics* 92.1, pp. 1–45.

Costello, Christopher et al. (2020). "The Future of Food from the Sea". In: *Nature* 588, pp. 95–100. DOI: 10.1038/s41586-020-2616-y.

Dafforn, Katherine A., John A. Lewis, and Emma L. Johnston (2011). "Antifouling strategies: History and regulation, ecological impacts and mitigation". In: *Marine Pollution Bulletin* 62.3, pp. 453–465. DOI: 10.1016/j.marpolbul.2011.01.012.

Darwin, Charles (1839). "Voyages of the Adventure and Beagle". In: *London: Henry Colburn*.

Dayton, Paul K (1985). "Ecology of kelp communities". In: *Annual review of ecology and systematics*, pp. 215–245.

Dunstan, Piers K., Scott D. Foster, Edward King, James Risbey, Terence J. O'Kane, Didier Monselesan, Alistair J. Hobday, Jason R. Hartog, and Peter A. Thompson (2018). "Global patterns of change and variation in sea surface temperature and chlorophyll a". In: *Scientific Reports* 8.1. ISSN: 2045-2322. DOI: 10.1038/s41598-018-33057-y.

Eger, Aaron M. et al. (2023). "The value of ecosystem services in global marine kelp forests". In: *Nature Communications* 14.1. ISSN: 2041-1723. DOI: 10.1038/s41467-023-37385-0.

Erlandson, Jon M., Michael H. Graham, Bruce J. Bourque, Debra Corbett, James A. Estes, and Robert S. Steneck (2007). "The Kelp Highway Hypothesis: Marine Ecology, the Coastal Migration Theory, and the Peopling of the Americas". In: *The Journal of Island and Coastal Archaeology* 2.2, pp. 161–174. ISSN: 1556-1828. DOI: 10.1080/15564890701628612.

Filbee-Dexter, Karen and Thomas Wernberg (2018). "Rise of Turfs: A New Battlefront for Globally Declining Kelp Forests". In: *BioScience* 68.2, pp. 64–76. DOI: 10.1093/biosci/bix147.

Halpern, Benjamin S. et al. (2009). "Mapping cumulative human impacts to California Current marine ecosystems". In: *Conservation Letters* 2.3, pp. 138–148. ISSN: 1755-263X. DOI: 10.1111/j.1755-263x.2009.00058.x.

Hansen-Lewis, Jamie and Michelle Marcus (2022). *Uncharted Waters: Effects of Maritime Emission Regulation*. Tech. rep. DOI: 10.3386/w30181.

Hersbach et al. (2023). "ERA5 hourly data on single levels from 1940 to present". In.

Jacox, Michael G., Christopher A. Edwards, Elliott L. Hazen, and Steven J. Bograd (2018). "Coastal Upwelling Revisited: Ekman, Bakun, and Updated Upwelling Indices for the U.S. West Coast". In: *Journal of Geophysical Research: Oceans* 123.10, pp. 7332–7350. DOI: 10.1029/2018JC014187.

Johansson, Per, Karl Martin Eriksson, Lennart Axelsson, and Hans Blanck (2012). "Effects of Seven Antifouling Compounds on Photosynthesis and Inorganic Carbon Use in Sugar Kelp Saccharina latissima (Linnaeus)". In: *Archives of Environmental Contamination and Toxicology* 63.3, pp. 365–377. DOI: 10.1007/s00244-012-9778-z.

Keiser, David A. and Joseph S. Shapiro (2019). "Consequences of the Clean Water Act and the Demand for Water Quality". In: *Quarterly Journal of Economics* 134.1, pp. 349–396. DOI: 10.1093/qje/qjy019.

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.

Krumhansl, Kira A. et al. (2016). "Global patterns of kelp forest change over the past half-century". In: *Proceedings of the National Academy of Sciences* 113.48, pp. 13785–13790. DOI: 10.1073/pnas.1606102113.

Lindgren, J. Fredrik, Magda Wilewska-Bien, Lena Granhag, Karin Andersson, and K. Martin Eriksson (2016). "Discharges to the Sea". In: *Shipping and the Environment*. Springer Berlin Heidelberg, pp. 125–168. ISBN: 9783662490457. DOI: 10.1007/978-3-662-49045-7_4.

Maxwell, Sara M. et al. (2013). "Cumulative human impacts on marine predators". In: *Nature Communications* 4.1. ISSN: 2041-1723. DOI: 10.1038/ncomms3688.

McHaskell, Danielle Alexis (2024). "Half a Century of Global Invasion: How Global Trends Can Inform the Spread and Phenology of the Non-Native Kelp, Undaria Pinnatifida, in California, United States". In: *Integrative And Comparative Biology* 64.4, pp. 1087–1101. ISSN: 1557-7023. DOI: 10.1093/icb/icae152.

McPherson, M. L., D. J. I. Finger, H. F. Houskeeper, T. W. Bell, M. H. Carr, L. Rogers-Bennett, and R. M. Kudela (2021). "Large-scale shift in the structure of a kelp forest ecosystem co-occurs with an epizootic and marine heatwave". In: *Communications Biology* 4.1, p. 298. DOI: 10.1038/s42003-021-01827-6.

Nedoncelle, Clément, Philippe Delacote, and Léa Crepin (2024). "Agricultural Exports, Market Power, and Deforestation". In.

Openshaw, Stan (1984). *The Modifiable Areal Unit Problem*. Concepts and Techniques in Modern Geography 38. Norwich: GeoBooks.

Phinney, Jonathan T., Frank Muller-Karger, Phil Dustan, and Jack Sobel (2001). "Using Remote Sensing to Reassess the Mass Mortality of Diadema antillarum 1983–1984". In: *Conservation Biology* 15.4, pp. 885–891. ISSN: 1523-1739. DOI: 10.1046/j.1523-1739.2001.015004885.x.

Rambachan, Ashesh and Jonathan Roth (2023). "A More Credible Approach to Parallel Trends". In: *Review of Economic Studies* 90.5, pp. 2555–2591. DOI: 10.1093/restud/rdad018.

Rogers-Bennett, L. and C. A. Catton (2019). "Marine heat wave and multiple stressors tip bull kelp forest to sea urchin barrens". In: *Scientific Reports* 9.1. DOI: 10.1038/s41598-019-51114-y.

Santos Silva, João and Silvana Tenreyro (2006). "The Log of Gravity". In: *The Review of Economics and Statistics* 88.4, pp. 641–658. URL: https://EconPapers.repec.org/RePEc:tpr:restat:v:88:y:2006:i:4:p:641-658.

Schiff, Kenneth, Jennifer Gosselin, and David Diehl (2011). *Southern California Bight 2008 Regional Monitoring Program: Volume IV Sediment Chemistry*. Tech. rep. 628. Costa Mesa, CA: Southern California Coastal Water Research Project. URL: https://www.sccwrp.org.

Shapiro, Joseph S. and Reed Walker (2018). "Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade". In: *American Economic Review* 108.12, pp. 3814–3854. DOI: 10.1257/aer.20151272.

Smith (2015). "The putative impacts of the non-native seaweed Sargassum muticum on native communities in tidepools of Southern California and investigation into the feasibility of local eradication". In: *Marine Ecology* 37.3, pp. 645–667. ISSN: 1439-0485. DOI: 10.1111/maec.12335.

Solon, Gary, Steven J. Haider, and Jeffrey M. Wooldridge (2015). "What Are We Weighting For?" In: *Journal of Human Resources* 50.2, pp. 301–316. DOI: 10.3368/jhr.50.2.301.

Steneck, Robert S., Michael H. Graham, Bruce J. Bourque, Debbie Corbett, Jon M. Erlandson, James A. Estes, and Mia J. Tegner (2002). "Kelp Forest Ecosystems: Biodiversity, Stability, Resilience and Future". In: *Environmental Conservation* 29.4, pp. 436–459. DOI: 10.1017/S0376892902000322.

Williams, Susan L. and Smith (2007). "A Global Review of the Distribution, Taxonomy, and Impacts of Introduced Seaweeds". In: *Annual Review of Ecology, Evolution, and Systematics* 38.1, pp. 327–359. ISSN: 1545-2069. DOI: 10.1146/annurev.ecolsys.38.091206.095543.

Yang, Kai, Amelie Meyer, Peter G. Strutton, and Andrew M. Fischer (2023). "Global trends of fronts and chlorophyll in a warming ocean". In: *Communications Earth amp; Environment* 4.1. ISSN: 2662-4435. DOI: 10.1038/s43247-023-01160-2.

Yunker, Mark B., Robie W. Macdonald, Roxanne Vingarzan, R. Harvey Mitchell, Darcy Goyette, and Stephanie Sylvestre (2002). "PAHs in the Fraser River Basin: A Critical Appraisal of PAH Ratios as Indicators of PAH Source and Composition". In: *Organic Geochemistry* 33.4, pp. 489–515. DOI: 10.1016/S0146-6380(02)00002-5.

A Pollution along shipping maritime routes

Section 2 discusses several channels through which maritime activity may affect kelp, including vessel discharges, antifouling contaminants, ballast water, and the deposition of combustion byproducts from heavy fuel oil (HFO) burning. A common feature of these channels is that their effects scale in intensity with the concentration of vessel traffic. Heavily trafficked corridors accumulate stressors through repeated passages, sediment disturbance, and the continuous release of discharge and exhaust byproducts. This appendix provides suggestive evidence that our AIS-based treatment variable captures meaningful variation in nearshore stressor exposure across kelp cells. Specifically, we show that the spatial footprint of vessel traffic in our study area is consistent with the geographic distribution of a relevant pollution indicator prior to ECA enforcement.

We match our AIS-derived shipping exposure to sediment chemistry stations from the Southern California Bight 2008 Regional Monitoring Program (Schiff et al., 2011), before ECA enforcement. This survey, coordinated by the Southern California Coastal Water Research Project (SCCWRP), sampled sediment chemistry across several stations spanning the nearshore and offshore waters of the Southern California Bight. At each station, sediment samples were analyzed for a suite of contaminants including trace metals and polycyclic aromatic hydrocarbons (PAHs), providing spatially extensive baseline measurements of seafloor pollution. We concentrate our analysis on high molecular weight pyrogenic PAHs — the chemical signature of petroleum combustion (Yunker et al., 2002) — providing a measure of HFO-related contamination that is partly distinguishable from other anthropogenic sources.

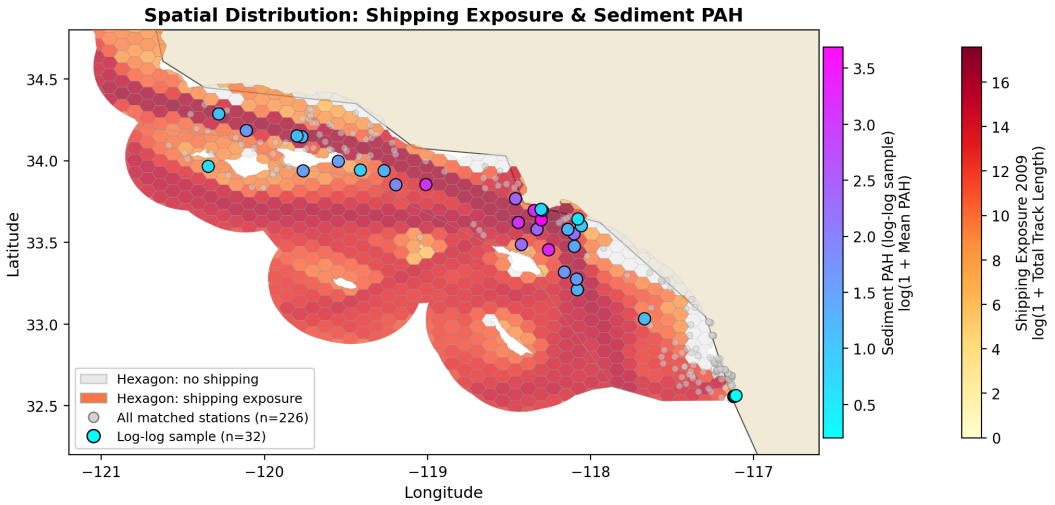

This figure presents the spatial distribution of shipping exposure and sediment PAH concentrations in the Southern California Bight. Hexagons (ECA grid, 10 km resolution) depict shipping intensity in 2009, measured as total ship track length (log-transformed, YlOrRd color scale). Gray hexagons indicate no recorded shipping activity. Light gray dots represent all sediment monitoring stations (Bight 2008 survey) matched to the nearest shipping grid cell using a nearest-neighbor assignment with a 10 km distance cutoff. Colored dots (cool color scale) highlight the subset of stations with strictly positive PAH concentrations and shipping exposure, used in the log-log regression analysis. Coastline from Natural Earth.

Figure 5: Shipping exposure and sediment PAH concentrations in the Southern California Bight

Figure 5 shows that stations with the highest pyrogenic PAH concentrations are systematically located in areas of high pre-policy vessel traffic. This pattern holds in a simple cross-sectional regression: log pyrogenic PAH concentrations are positively and significantly correlated with log mean vessel track length across the 32 Bight stations within our study area (Pearson $r = 0.41$, $p < 0.05$; Spearman $\rho = 0.55$, $p < 0.01$). The stronger Spearman rank correlation suggests that the relationship is not driven by a small number of outlier stations, but reflects a broad spatial gradient in contamination that tracks shipping intensity across the study area.

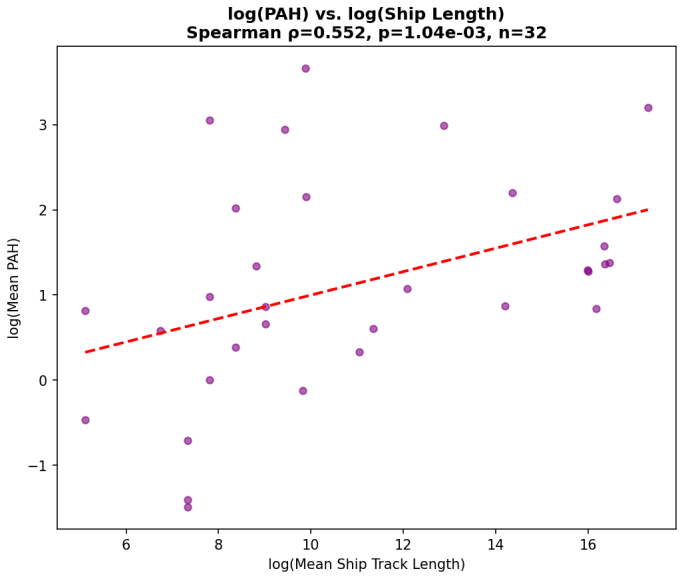

Each point represents a sediment monitoring station with strictly positive mean PAH concentration (Bight 2008) and shipping exposure (total ship track length, 2009). Both variables are expressed in natural logarithm. The red dashed line shows the ordinary least squares (OLS) fit, with the estimated slope () reported in the legend. Pearson's correlation coefficient (r) and Spearman's rank correlation (), along with their associated p-values, are displayed in the panel title. Station-to-grid assignment is based on nearest-neighbor matching with a 10 km distance threshold.

Figure 6: Log-log relationship between shipping exposure and sediment PAH concentrations

We caution that this exercise is illustrative rather than causal — sediment PAH concentrations reflect multiple anthropogenic sources, and Bight stations do not map perfectly onto our kelp hexagons — but the consistency between our treatment variable and an independent pollution indicator strengthens confidence that the DiD estimates capture a genuine ecological response to changes in nearshore maritime conditions.

B Weighting Robustness and Cumulative Impact Index: the 2009 ECA

The main specification uses unweighted PPML on the restricted sample of 81 cells recording at least two vessel passages per month on average. This appendix examines robustness to alternative weighting schemes and to an alternative treatment measure.

Table 3 reports weighted PPML estimates on the same restricted sample, using log vessel count as observation weight (offset trick; see Solon et al., 2015 for a discussion of weighting motives in regression). The coefficients are virtually identical to the unweighted estimates in Table 1: 0.135 versus 0.136 for vessel distance and 0.152 versus 0.155 for vessel count, all significant at the 1% level. This stability confirms that the results

are not sensitive to the weighting scheme once cells with negligible shipping exposure are excluded.

	Distance Traveled (km)		Vessel Count	
	(1) Biomass	(2) Area	(3) Biomass	(4) Area
Post 2009 × Exposure	0.1352***	0.1308***	0.1517***	0.1463***
	(0.0372)	(0.0366)	(0.0446)	(0.0439)
SST	-5.364***	-5.194***	-5.387***	-5.216***
	(1.176)	(1.161)	(1.181)	(1.167)
Precipitation	0.0842	0.0899	0.0911	0.0965
	(0.1093)	(0.1065)	(0.1100)	(0.1071)
Significant Wave Height	-3.426**	-3.278**	-3.495**	-3.345**
	(1.538)	(1.512)	(1.541)	(1.513)
Cell FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes
Observations	3880	3880	3880	3880

Notes: Weighted PPML estimates of Eq. (1). The dependent variable is kelp biomass (kg) or canopy area (m^2). "Distance Traveled (km)" columns (1)–(2) measure exposure by logged mean distance traveled by vessels within a 20 km buffer during January–June 2009; "Vessel Count" columns (3)–(4) use the corresponding logged mean vessel count. Observations weighted by log vessel count (offset trick). All specifications include kelp-cell and year–quarter fixed effects. Sample: Southern California (32°–36° N), 2005–2016, restricted to cells with at least two vessel passages per month on average (81 cells, 3,880 observations). Standard errors clustered at the cell level (81 clusters) in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Weighted PPML: Kelp outcomes and the 2009 ECA

As a complementary exercise, we replace the AIS-based treatment with the cumulative human impact index of Halpern et al. (2009), which aggregates 25 anthropogenic stressors into a single measure. The coefficient is positive and significant, but event-study diagnostics reveal severe violations of the parallel-trends assumption (pre-policy coefficients of -0.18 to -0.14, all significant at the 1% level; joint Wald test $p < 0.001$). These pre-trends are driven by non-shipping stressors bundled in the index, motivating our preference for shipping-specific AIS measures.

Sensitivity to sample restriction. Figure 7 traces the coefficient as a function of an increasingly restrictive sample cutoff. The estimate stabilizes once cells with negligible shipping exposure are excluded, confirming that the baseline restriction removes noise from cells where the treatment variable captures geographic features rather than policy-relevant variation.

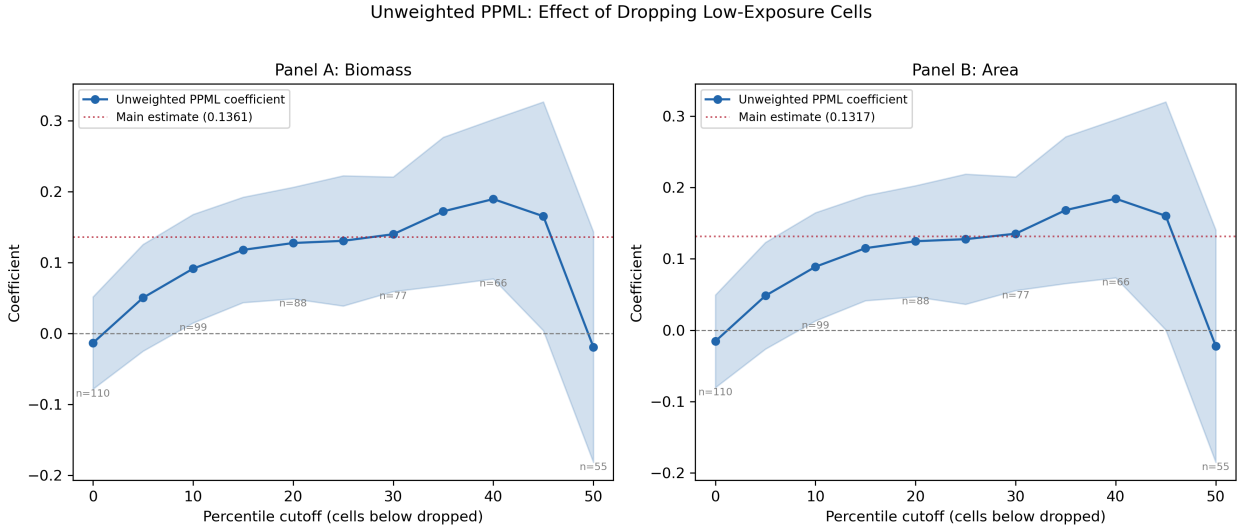

Notes: Each point reports the unweighted PPML coefficient from Eq. (1) estimated on the subsample of cells above the indicated percentile of pre-policy shipping exposure (logged mean vessel distance). Cutoffs range from the 5th to the 50th percentile. The shaded area shows 95% confidence intervals based on standard errors clustered at the cell level. The horizontal dashed line marks the main estimate for each dependent variable. The number of cells in each subsample is indicated below the confidence band. Sample: Southern California, 2005–2016.

Figure 7: PPML coefficient as a function of sample restriction

Table 4 confirms this pattern across alternative specifications. Panel A reproduces the main unweighted estimate. Panel B applies vessel count weights (offset trick) on the same restricted sample; the coefficients are nearly identical. Panel C uses pre-period kelp biomass (2005–2008 average) as an alternative weight, capturing the ecological size of each cell rather than its shipping exposure; all four coefficients remain positive and significant at the 1% level.

	Distance Traveled (km)		Vessel Count	
	(1) Biomass	(2) Area	(3) Biomass	(4) Area
Panel A: Main specification (unweighted PPML)				
Post 2009 × Exposure	0.1361***	0.1317***	0.1548***	0.1493***
	(0.0359)	(0.0352)	(0.0432)	(0.0424)
Observations	3880	3880	3880	3880
Panel B: Weighted by vessel count (offset trick)				
Post 2009 × Exposure	0.1352***	0.1308***	0.1517***	0.1463***
	(0.0372)	(0.0366)	(0.0446)	(0.0439)
Observations	3880	3880	3880	3880
Panel C: Weighted by pre-period kelp biomass				
Post 2009 × Exposure	0.1029***	0.0995***	0.1242***	0.1196***
	(0.0296)	(0.0286)	(0.0382)	(0.0365)
Observations	3736	3736	3736	3736
Controls	Yes	Yes	Yes	Yes
Cell FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes

Notes: Panel A reproduces the main specification from Table 1 (unweighted PPML on cells with ≥ 2 vessel passages/month). Panel B applies vessel count weights (offset trick) on the same restricted sample. Panel C uses pre-period kelp biomass (2005–2008 cell-level average) as observation weight; 3 cells with zero pre-period kelp are dropped. All specifications include kelp-cell and year–quarter fixed effects. Controls: log sea surface temperature, log precipitation, log significant wave height. Standard errors clustered at the cell level in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Weighting robustness: alternative specifications

C Placebo Test for the 2009 ECA

A natural concern with the main results is that the positive post-2009 coefficient might reflect a pre-existing differential trend rather than a causal effect of the ECA regulation. To address this, we conduct a placebo test by assigning a fictitious treatment date within the pre-policy period and verifying that no significant effect is detected.

Specifically, we restrict the sample to the pre-policy years 2005–2008 and define a fake "post" indicator equal to one from 2007Q3 onward (and zero before). We then re-estimate Eq. (1) using the same PPML specification as in the main analysis: the same treatment variables (vessel distance and vessel count), the same controls (SST, precipitation, significant wave height), and the same cell and year–quarter fixed effects. The only difference is that the genuine post-2009 indicator is replaced by the fake 2007Q3 indicator, and the sample contains no post-regulation observations.

Table 5 reports the results. All four placebo coefficients are small, negative, and statistically insignificant, with p-values ranging from 0.26 to 0.30. This stands in sharp contrast to the main estimates in Table 1, where the coefficients are positive and significant ($+0.136$ for vessel distance, $+0.155$ for vessel count, both $p < 0.01$). The placebo coefficients are also opposite in sign, ruling out the possibility that any pre-existing positive differential trend drives the main result.

	Distance Traveled (km)		Vessel Count	
	(1) Biomass	(2) Area	(3) Biomass	(4) Area
Fake Post 2007Q3 × Exposure	-0.0753	-0.0694	-0.0859	-0.0804
	(0.0673)	(0.0669)	(0.0755)	(0.0749)
Observations	1248	1248	1248	1248
Controls	Yes	Yes	Yes	Yes
Cell FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes

Notes: PPML estimates of Eq. (1) with a fictitious treatment date set at 2007Q3 on the pre-policy sample (2005–2008 only). The dependent variable is kelp biomass (kg) or canopy area (m^2). "Distance Traveled (km)" columns (1)–(2) measure exposure by logged mean distance traveled by vessels within a 20 km buffer during January–June 2009; "Vessel Count" columns (3)–(4) use the corresponding logged mean vessel count. All specifications include kelp-cell and year–quarter fixed effects. Sample restricted to cells with at least two vessel passages per month on average. Standard errors clustered at the cell level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Placebo test: Fake treatment date (2007Q3) on pre-policy sample

These results are consistent with the positive effect identified in the main analysis being specific to the actual 2009 ECA enforcement date rather than arising from pre-existing differential trends between high- and low-shipping cells.

D Balance Table

Table 6 reports pre-treatment (2005–2008) means of key variables across quartiles of the treatment variable (log vessel distance traveled, January–June 2009). The last column reports the p-value from a Kruskal–Wallis test for differences across quartiles.

The treatment variable exhibits strong variation by construction: mean log vessel distance ranges from 9.3 in the lowest quartile to 13.3 in the highest. Vessel count shows a similar gradient. SST, precipitation, and significant wave height all differ significantly across quartiles ($p < 0.01$), reflecting geographic sorting—high-shipping cells are concentrated near ports, which tend to be in calmer, warmer waters. These environmental variables are controlled for directly in the regression.

Pre-treatment kelp biomass and canopy area are highest in Q1 (lowest exposure) and decline monotonically with shipping intensity, indicating that high-exposure cells do not have systematically better kelp conditions before the policy. This pattern is reassuring: it rules out the concern that the positive post-2009 coefficient reflects a pre-existing advantage of high-shipping cells.

	Q1 (low)	Q2	Q3	Q4 (high)	KW p
Log vessel distance (Jan-Jun 2009)	9.311	10.491	12.470	13.321	0.000
	[0.491]	[0.290]	[0.564]	[0.182]	
Kelp biomass (kg)	1655116.8	801416.6	409706.0	321770.5	0.000
	[3976548.5]	[2459511.5]	[929510.4]	[904773.3]	
Kelp canopy area (m^2)	220694.5	107794.3	54520.6	43562.9	0.000
	[521297.4]	[333245.2]	[121129.7]	[123756.5]	
Log SST	2.842	2.883	2.798	2.844	0.000
	[0.134]	[0.124]	[0.125]	[0.124]	
Log precipitation	2.611	2.506	2.779	2.729	0.009
	[1.052]	[1.116]	[1.221]	[1.164]	
Log significant wave height	0.8859	0.7991	0.8732	0.7286	0.000
	[0.1626]	[0.1313]	[0.1748]	[0.1465]	
Log vessel count (Jan-Jun 2009)	0.9722	1.516	3.593	4.300	0.000
	[0.1562]	[0.310]	[0.468]	[0.134]	
Cells	21	20	20	20	

Notes: Cell-quarter means during the pre-treatment period (2005–2008), by quartile of log vessel distance traveled (January–June 2009, 20 km buffer). Standard deviations in brackets. The last column reports the p-value from a Kruskal–Wallis test of equality of distributions across quartiles. Sample: 81 cells (20–21 per quartile), Southern California (32°–36° N), restricted to cells with at least two vessel passages per month on average.

Table 6: Pre-treatment balance across treatment-intensity quartiles

E Financial Crisis Robustness

The 2009 ECA enforcement coincides with the 2008–09 global financial crisis and the Great Trade Collapse, which sharply reduced port throughput worldwide. This coincidence raises a natural concern: if the crisis differentially improved ecological conditions in high-shipping cells—for example, through reduced industrial discharge, ancillary port activity, or vessel traffic unrelated to the sulfur regulation—then our estimates could partly reflect the recession rather than the ECA.

We address this confound using vessel arrival records from the National Ballast Information Clearinghouse (NBIC), a federal database jointly administered by the Smithsonian Environmental Research Center and the U.S. Coast Guard that records every commercial vessel arrival at U.S. ports.[15] We filter to overseas arrivals at the four Southern California ports relevant to our study area (Los Angeles, Long Beach, San Diego, and Port Hueneme) and aggregate to port-by-quarter frequency over the 2005–2016 study period. We assign each kelp cell a measure of port activity by combining vessel arrivals from the four ports, giving more weight to ports that are geographically closer to that cell [16]

Figure 8 documents the raw port-activity data at quarterly frequency. Panel (a) shows total quarterly overseas vessel arrivals at Southern California ports, which dropped sharply after the Lehman collapse (September 2008) and remained depressed through 2016. Panel (b) decomposes this by port (indexed to 2007 = 100), revealing that the decline was broad-based—all four ports experienced substantial and persistent reductions, with San Diego and Port Hueneme experiencing the largest proportional declines.

[15] Data available at `https://nbic.si.edu/`. The NBIC database covers all vessel arrivals from 2004 onward.

[16] Formally, for each kelp hexagon we compute the inverse-distance-weighted sum of quarterly arrivals across the four ports, using haversine distances from cell centroids to port coordinates. This serves as a proxy for the structural economic forces driving local maritime traffic, and controls for the possibility that kelp cells near high-activity ports follow distinct pre-existing trends independent of the policy change. Restricting to overseas arrivals ensures the measure captures shipping traffic patterns.

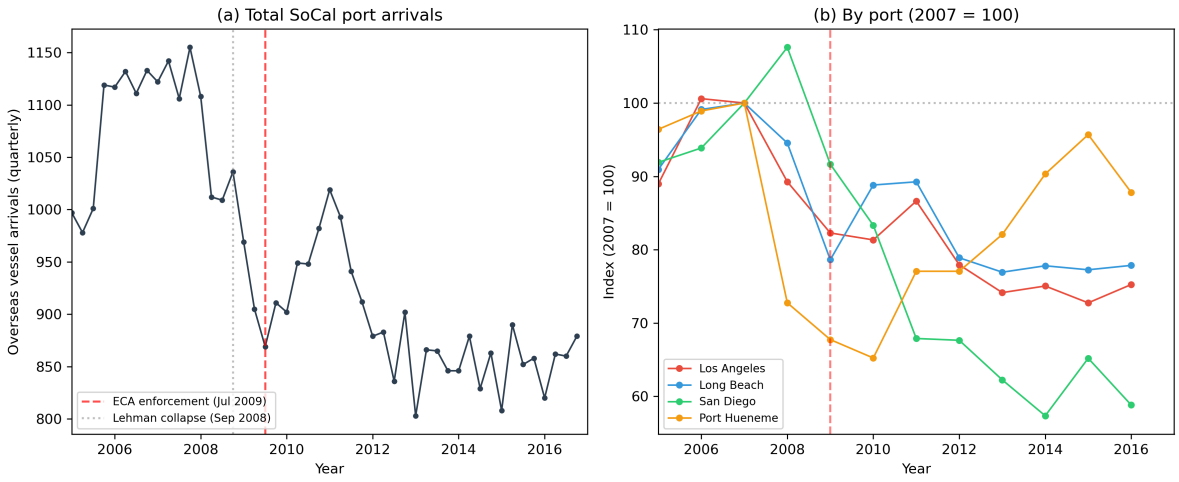

Note: Panel (a): total quarterly overseas vessel arrivals at four Southern California ports (Los Angeles, Long Beach, San Diego, Port Hueneme). Panel (b): arrivals by port, indexed to 2007 = 100. Dashed red line marks ECA enforcement (July 2009); dotted gray line marks the Lehman collapse (September 2008). Data source: National Ballast Information Clearinghouse (NBIC).

Figure 8: Overseas vessel arrivals at Southern California ports

Figure 9 presents the core identification argument by juxtaposing port activity with the kelp event study. Panel (a) shows the same data aggregated to annual frequency (2007 = 100). Port activity fell to approximately 81% of its 2007 level in 2009 and never recovered: by the end of the sample, arrivals stood at only 76% of the pre-crisis baseline. If reduced economic activity at ports—rather than the ECA—were driving kelp improvements, the effect should fade as ports remain depressed or should intensify as port activity declines further. Neither pattern is consistent with the event study in Panel (b), which shows positive coefficients for the interaction of shipping exposure and post-2009 indicators despite the continued depression in port activity.

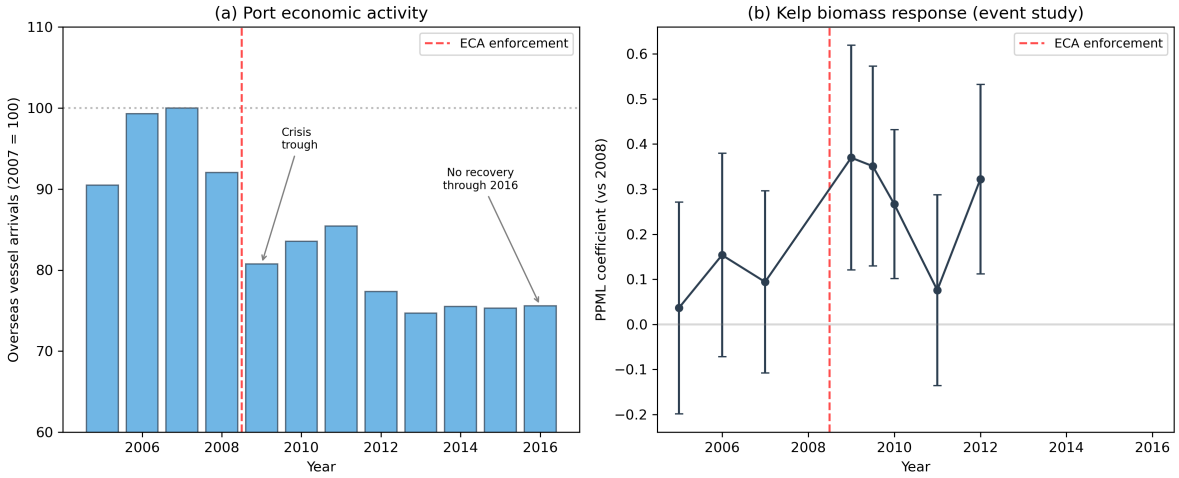

Note: Panel (a): annual overseas vessel arrivals at four Southern California ports (Los Angeles, Long Beach, San Diego, Port Hueneme), indexed to 2007 = 100. Data source: National Ballast Information Clearinghouse (NBIC). Panel (b): PPML event-study coefficients from Eq. (2), using vessel distance as treatment intensity (reference year: 2008). The dashed vertical line marks ECA enforcement (July 2009). 95% confidence intervals based on standard errors clustered at the cell level.

Figure 9: Port economic activity and kelp biomass response

Table 7 formalizes this argument by adding the log of distance-weighted port arrivals as a time-varying control in the PPML specification. Panel A reproduces the baseline estimates from Table 1. Panel B adds the port-activity control. The shipping-exposure coefficient is virtually unchanged: the attenuation is at most 2% across the four specifications, and all coefficients remain statistically significant at the 1% level. The port-activity variable itself enters positively and significantly (coefficients between 0.58 and 0.68, $p < 0.01$), indicating that it has genuine explanatory power for kelp outcomes—yet it does not absorb the treatment effect. This stability demonstrates that the differential kelp improvement in high-shipping cells is not driven by the differential reduction in port economic activity during the financial crisis.

	Distance Traveled (km)		Vessel Count	
	(1) Biomass	(2) Area	(3) Biomass	(4) Area
Panel A: Baseline (no port activity control)				
Post 2009 × Exposure	0.1361***	0.1317***	0.1548***	0.1493***
	(0.0359)	(0.0352)	(0.0432)	(0.0424)
Panel B: Controlling for port economic activity				
Post 2009 × Exposure	0.1363***	0.1319***	0.1520***	0.1465***
	(0.0360)	(0.0353)	(0.0434)	(0.0426)
Log port arrivals	0.6831***	0.6698***	0.5851***	0.5751***
	(0.2084)	(0.2081)	(0.2011)	(0.2009)
Observations	3880	3880	3880	3880
Controls	Yes	Yes	Yes	Yes
Cell FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes

Notes: PPML estimates of Eq. (1). Panel A: baseline specification (identical to Table 1). Panel B: adds log distance-weighted overseas vessel arrivals at Southern California ports (NBIC data) as a time-varying control. The port-activity variable is constructed as the inverse-distance-weighted sum of quarterly arrivals across four ports (Los Angeles, Long Beach, San Diego, Port Hueneme) for each kelp hexagonal cell. All specifications include kelp-cell and year–quarter fixed effects. Sample: Southern California (32°–36° N), 2005–2016, restricted to cells with at least two vessel passages per month on average (81 clusters). Standard errors clustered at the cell level in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Robustness to controlling for port economic activity

The combination of these two findings—coefficient stability when controlling for port activity, and the divergent timing of port depression versus kelp improvement—provides evidence that the financial crisis does not confound the estimated ECA effect.

F Anticipation and Treatment Timing

A potential concern with our identification strategy is anticipation. Because the regulation was announced before the July 2009 enforcement date, vessel operators may have begun adjusting their routes in the preceding months. If such adjustments altered the *spatial pattern* of shipping exposure, our treatment variable—measured during January–June 2009—could reflect anticipatory routing rather than the pre-enforcement equilibrium, biasing our estimates.

There are strong economic reasons to doubt that anticipation is empirically relevant in this setting. The regulation imposes the use of low-sulfur marine fuel within the ECA, which is substantially more expensive than the heavy fuel oil used outside it (Klotz and Berazneva, 2022). Switching to compliant fuel thus entails a direct marginal cost increase on every voyage entering the zone. Vessel operators therefore have a clear incentive to *delay* compliance until enforcement begins, not to anticipate it: using low-sulfur fuel before it is required would raise operating costs with no offsetting regulatory benefit. Anticipation through *route adjustments*—sailing outside the ECA to avoid the zone entirely—is equally unlikely before enforcement, since the penalty for noncompliance did not apply until July 2009. In short, the cost structure of the regulation creates a strong incentive to maintain pre-policy routing patterns until the enforcement date, making it improbable that the January–June 2009 treatment variable reflects anticipatory behavior. Nonetheless, to substantiate this reasoning empirically, we provide two formal tests below.

We cannot address this concern by constructing a pre-2009 treatment variable from AIS data, because global AIS coverage was limited before 2009 and the raw AIS files in our dataset begin in January 2009. Instead, we provide two pieces of evidence that anticipation did not distort the spatial distribution of shipping across kelp cells.

Temporal stability of cell-level shipping rankings. If anticipation had reshuffled the spatial pattern of vessel traffic during January–June 2009, we would expect that pattern to differ systematically from patterns observed in later years, when any anticipatory rerouting would have subsided. Figure 10 reports the Spearman rank correlation between the January–June 2009 measure and the corresponding full-year measure for 2009 through 2016, computed across the 81 cells in our estimation sample. For both vessel distance and vessel count, rank correlations are high ($\rho \geq 0.94$) in years unaffected by the boundary modification (2009, 2012–2014). The correlations drop in 2010–2011 (ρ between 0.63 and 0.76), consistent with the 2011 boundary modification reshuffling traffic patterns in a subset of cells near the Channel Islands. Importantly, the correlations rebound sharply in 2012–2014 ($\rho \geq 0.94$), confirming that the long-run spatial ordering of cells by shipping intensity is highly persistent. The 2010–2011 dip does not threaten the 2009 identification—it reflects the 2011 policy shock, not anticipation of the 2009 rule—and the high correlations in all other years indicate that anticipatory rerouting did not meaningfully alter which cells were more or less exposed to vessel traffic.

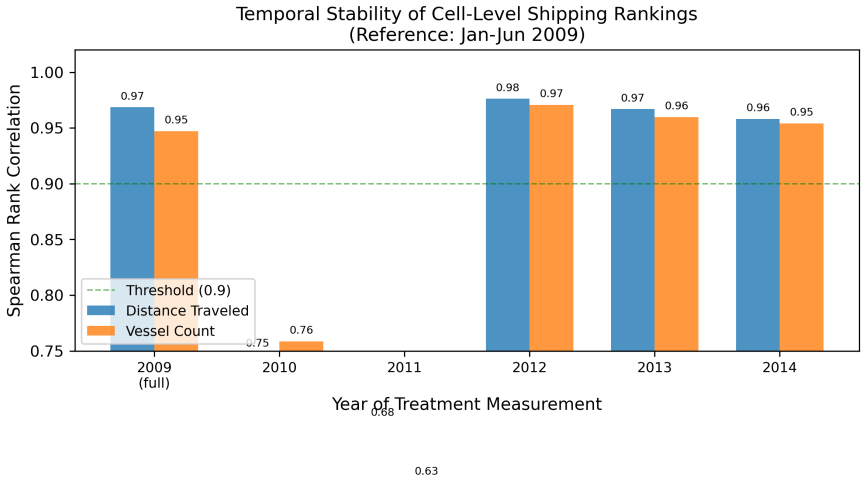

Note: Each bar reports the Spearman rank correlation between cell-level shipping exposure measured during January–June 2009 (the baseline treatment) and the corresponding full-year measure for the indicated year. Computed across the 81 cells in the estimation sample. The dashed line at 0.9 is shown for reference. Correlations are high (≥ 0.94) in all years except 2010–2011, where the boundary modification reshuffled traffic near the Channel Islands.

Figure 10: Temporal stability of cell-level shipping rankings

Later-year treatments. As a stability check, we re-estimate Eq. (1) replacing the January–June 2009 treatment with 2012 AIS data. Because the 2009 rule affected routing, the 2012 cross-section is not strictly predetermined; this exercise therefore tests whether the spatial pattern of the results is robust to using a post-policy proxy rather than providing an independent causal estimate. Table 8 reports the results. The coefficients are larger in magnitude than the baseline (+0.156 for distance, +0.192 for count), consistent with the spatial pattern of shipping becoming more precisely estimated as AIS coverage improves, and are all significant at the 1% level. These results confirm that the choice of treatment timing does not qualitatively affect the main findings.

	Distance Traveled (km)		Vessel Count	
	(1) Biomass	(2) Area	(3) Biomass	(4) Area
Panel A: Baseline (January–June 2009)				
Post 2009 × Exposure	0.1361***	0.1317***	0.1548***	0.1493***
	(0.0359)	(0.0352)	(0.0432)	(0.0424)
Panel B: 2012 (post-adjustment)				
Post 2009 × Exposure	0.1559***	0.1512***	0.1922***	0.1860***
	(0.0389)	(0.0382)	(0.0489)	(0.0481)
Observations	3880	3880	3880	3880
Controls	Yes	Yes	Yes	Yes
Cell FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes

Notes: PPML estimates of Eq. (1). Each panel uses a different measurement period for the treatment variable: Panel A uses January–June 2009 (baseline) and Panel B uses full-year 2012 AIS data. Both panels use the same post-2009 indicator, controls (SST, precipitation, significant wave height), and cell and year–quarter fixed effects. Sample restricted to cells with at least two vessel passages per month on average. Standard errors clustered at the cell level in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Anticipation robustness: Alternative treatment timing

Taken together, the evidence indicates that the spatial distribution of shipping exposure across kelp cells was highly persistent throughout our study period. The January–June 2009 treatment variable is consistent with the long-run spatial pattern of maritime activity rather than transient anticipation-induced rerouting, and the main results are robust to alternative treatment timing.

G Synthetic Differences-in-Differences

A natural complement to the continuous DiD of Section 3 would be a binary comparison of kelp outcomes in regulated California ports versus an unregulated port outside the boundary.[17] However, because the regulation applies uniformly across California, the effective number of treatment clusters is very small (three treated port areas versus one control port). With only one control-group cluster, there is no within-group variation to estimate the control-group residual variance, making cluster-robust inference essentially impossible. Even wild bootstrap corrections become unreliable when the number of clusters falls below roughly five or six (Cameron et al., 2008; Cameron and Miller, 2015). This structural limitation, inherent in any single-jurisdiction policy, prevents credible inference from a binary design regardless of the number of sub-cluster observations within each port area.

The continuous DiD in Section 3 partially addresses this concern by exploiting within-zone variation in exposure intensity, so that inference rests on 81 kelp-cell clusters rather than on a handful of port areas. The synthetic control exercise reported below provides a complementary external comparison that does not rely on within-California exposure variation: rather than selecting a single control port, it constructs a data-driven counterfactual from the full set of Mexican coastal hexagons using the method of Arkhangelsky et al. (2021).

The analysis is conducted on a 10 km hexagonal grid. We aggregate kelp biomass from 30 m $\times$ 30 m pixels to hexagons and construct an annual series for each hexagon. Treated units are coastal hexagons located within the 2009 regulated boundary; we exclude hexagons surrounding the ports already analyzed in the main text and those within the 2011 boundary modification area. The donor pool consists of coastal hexagons along the Mexican coast. Figure 12 displays the treated and donor units used in the synthetic control construction.

The synthetic control is constructed as a weighted combination of donor hexagons selected to replicate the treated pre-2009 biomass trajectory as closely as possible. The resulting weights are reported in Figure 11.

[17] We estimated a standard two-way fixed-effects DiD comparing kelp cells near Los Angeles–Long Beach, San Diego, and Port Hueneme to cells near Ensenada, Mexico—a deep-water port embedded in the same Pacific shipping network but outside the California regulation. Post-2009 interaction terms were positive for all three treated ports; however, the standard errors from this specification are unreliable given the small number of clusters (see discussion in the text). These results are available on request.

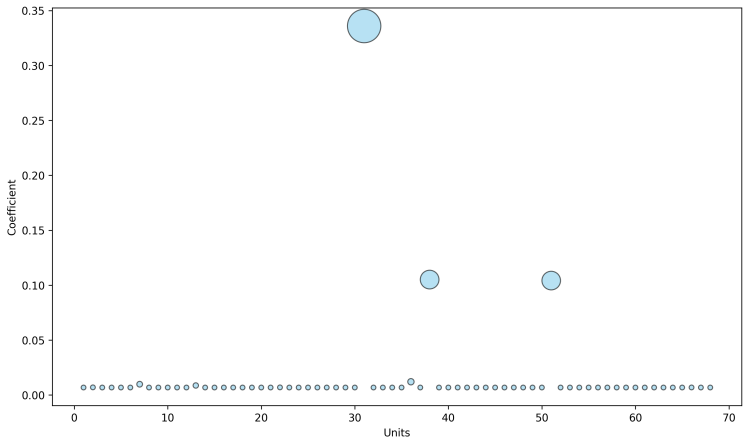

Notes: Weights assigned to each donor hexagon in the synthetic control construction. Donor hexagons are Mexican coastal cells. Larger values indicate greater contribution to the synthetic counterfactual. The outcome variable used for matching is summed kelp biomass (kg) per hexagon-year.

Figure 11: Synthetic control donor weights

A small number of donor hexagons receive most of the weight, which indicates that these locations are the most informative for matching the treated pre-period path. This feature is common in synthetic control applications.

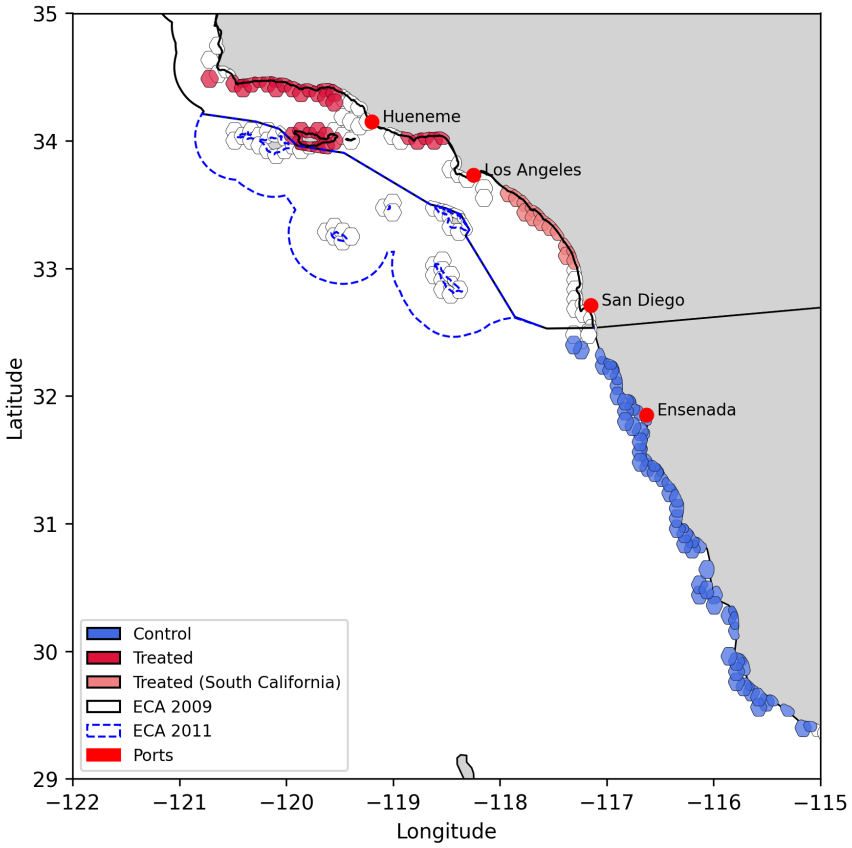

Notes: Red hexagons are treated units within the 2009 ECA boundary; light red hexagons are the southern California subset excluded in column (2) of Table 9. Blue hexagons are the Mexican coastal donor pool. Geographic scope: Southern California to Baja California.

Figure 12: Map for the synthetic control analysis

Table 9 reports the estimated post-2009 effect for two treated samples: (1) all treated hexagons shown in red in Figure 12, and (2) a treated sample excluding the southern California subset (light red in Figure 12).

	(1) All treated	(2) Excl. Southern CA
	Biomass (ln)	
Post 2009 (SDID)	1.391**	1.911**
	(0.642)	(0.753)
N	819	721

Notes: Synthetic differences-in-differences estimates (Arkhangelsky et al., 2021) of the post-2009 ECA effect on log kelp biomass. The dependent variable is the natural log of summed kelp biomass (kg) per hexagon-year. Column (1) uses all treated hexagons within the 2009 ECA boundary (red in Figure 12); column (2) excludes southern California hexagons (light red in Figure 12). The donor pool consists of coastal hexagons along the Mexican coast. Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Synthetic differences-in-differences: Post-2009 ECA effect on kelp biomass

In both cases, the estimated effect is positive and statistically significant, indicating that biomass in treated California hexagons evolves more favorably after 2009 than in their synthetic counterparts constructed from the Mexican donor pool. Figure 13 shows the same result in levels: the synthetic series closely tracks the treated series prior to 2009 and then falls markedly afterwards, while the treated series remains comparatively stable. The post-2009 decline in the synthetic series reflects the realized evolution of the donor locations that best reproduce the treated pre-period trajectory.

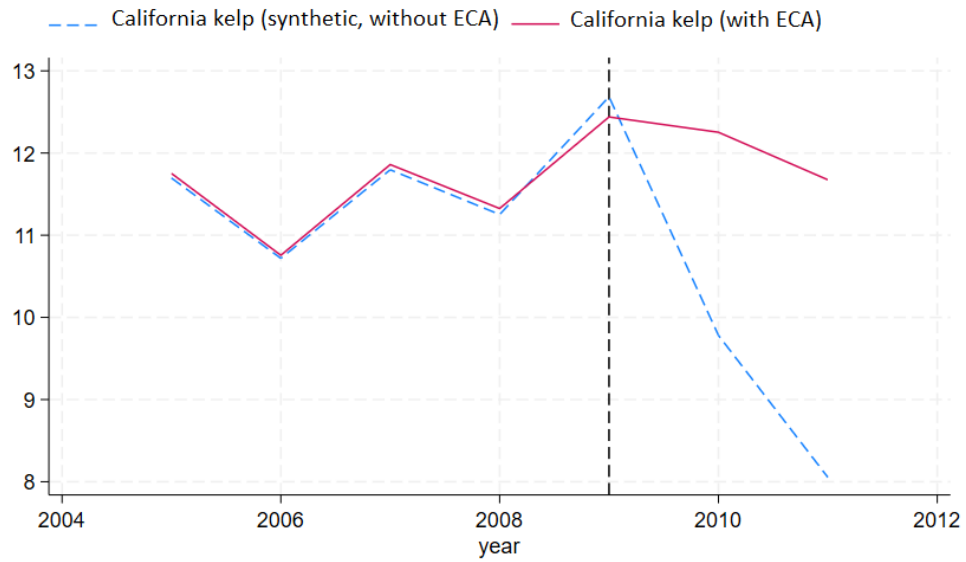

Notes: Annual kelp biomass (summed across hexagons) for treated cells within the 2009 ECA zone—excluding port areas analyzed in the main text—compared to a synthetic counterpart constructed from untreated Mexican coastal hexagons using the method of Arkhangelsky et al. (2021). The vertical dashed line marks the 2009 ECA implementation.

Figure 13: Synthetic control: treated vs. synthetic kelp biomass trends

H After 2011 Boundary Modification

Figure 14 plots the log of total distance traveled by vessels, averaged across hexagonal cells in each zone, comparing the 2009 regulated zone and the area corresponding to the 2011 boundary modification. The series illustrates a redistribution of traffic across these zones over time, consistent with route adjustments and the boundary modification documented in Klotz and Berazneva (2022).

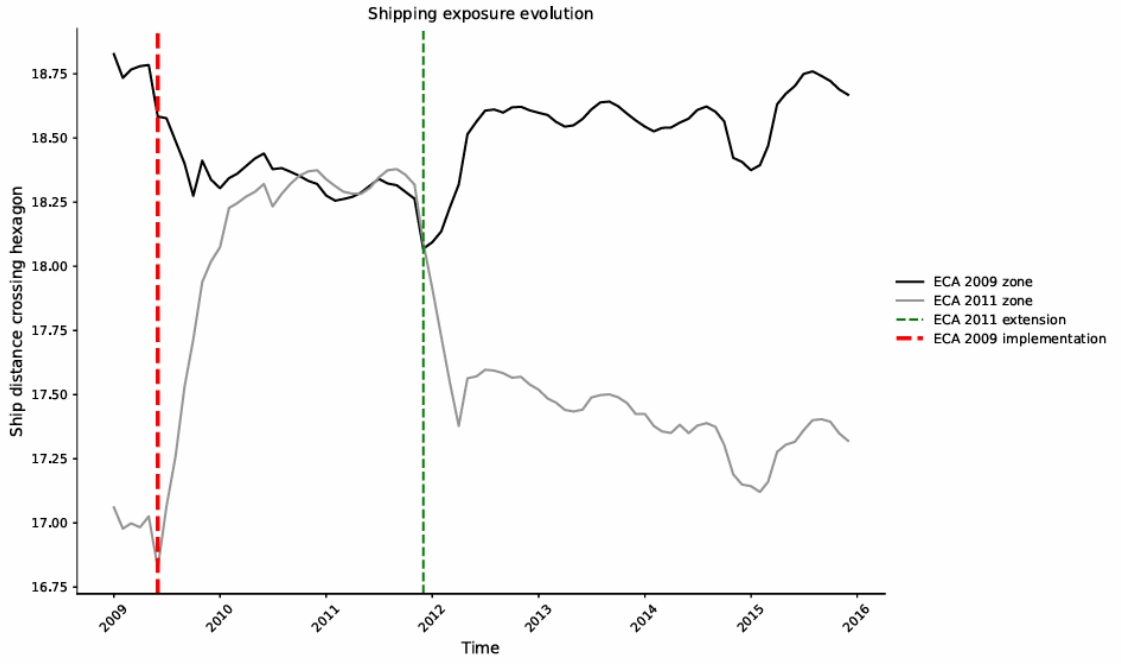

Notes: Log of total distance traveled by vessels (km), averaged across hexagonal cells in each zone, for cells near kelp forests. "ECA 2009 zone" aggregates cells within the original 2009 regulated boundary; "ECA 2011 zone" aggregates cells in the 2011 boundary modification area. Data presented monthly from 2009 to 2016. Data source: AIS.

Figure 14: Evolution of AIS shipping patterns

I Robustness to Pre-Treatment Exposure: 2010 Shipping

Because the 2011 boundary modification took effect in December 2011, the full-year 2011 treatment variable includes one month of post-modification routing. Although the main text argues that any contamination is quantitatively negligible and biased toward attenuation, we provide a direct robustness check by replacing the 2011 treatment with shipping exposure measured during 2010—a year entirely preceding the boundary change.

Table 10 reports PPML estimates that are otherwise identical to Table 2: same sample, same fixed effects, same controls. The only change is the treatment variable, which now uses logged vessel distance and count within a 20 km buffer measured during 2010 instead of 2011.

	Distance Traveled (km)		Vessel Count	
	(1) Biomass	(2) Area	(3) Biomass	(4) Area
Post 2011 × Exposure (2010)	-0.0766***	-0.0757***	-0.213***	-0.211***
	(0.0123)	(0.0122)	(0.0460)	(0.0451)
SST	-7.046***	-6.950***	-7.118***	-7.022***
	(0.914)	(0.920)	(0.901)	(0.907)
Precipitation	-0.0480	-0.0505	-0.0452	-0.0478
	(0.0777)	(0.0771)	(0.0788)	(0.0781)
Significant Wave Height	-3.326***	-3.245***	-3.363***	-3.282***
	(0.962)	(0.955)	(0.959)	(0.952)
Cell FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes
Observations	6000	6000	6000	6000

Notes: PPML estimates of Eq. (3) with shipping exposure measured during 2010 instead of 2011. The dependent variable is kelp biomass (kg) or canopy area (m^2). "Distance Traveled (km)" columns (1)–(2) measure exposure by logged total distance traveled by vessels within a 20 km buffer around each kelp hexagon during 2010; "Vessel Count" columns (3)–(4) use the corresponding logged vessel count. All specifications include kelp-cell and year–quarter fixed effects. Standard errors clustered at the cell level in parentheses. Sample: Southern California (32°–36° N), 2005–2016.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Robustness: 2010 shipping exposure as treatment variable

All four coefficients are negative and statistically significant at the 1% level. The vessel count coefficients are virtually unchanged (-0.213 with 2010 treatment versus -0.210 with 2011 treatment), well within each other's confidence intervals. The distance coefficients are smaller in absolute value (-0.077 versus -0.093), which is expected: the 2010 routing pattern predates the boundary modification by over a year, so it provides a slightly less precise spatial match to the corridors actually affected by the December 2011 change. The attenuation is consistent with classical measurement error in the treatment variable rather than with the 2011 results being inflated by post-modification contamination. In sum, using a purely pre-treatment measure of shipping exposure preserves the sign, significance, and economic interpretation of the main findings.

J Treatment Correlation: 2009 vs. 2011 Exposure

The 2009 and 2011 specifications use different treatment variables measured at different dates: January–June 2009 AIS data for the main ECA analysis and full-year 2011 AIS data for the boundary modification. Because both variables capture the spatial distribution of vessel traffic around the same set of kelp cells, they are necessarily correlated. However, the strength and nature of this correlation matters for interpretation: if the two treatment variables were nearly collinear, the 2011 results would largely replicate the 2009 cross-sectional variation rather than exploiting independent spatial information from the boundary modification.

Figure 15 plots the cell-level 2009 treatment against the 2011 treatment for both distance traveled and vessel count. The Pearson and Spearman correlations are reported in each panel. For vessel distance, the Pearson correlation is $r = 0.67$ and the Spearman rank correlation is $\rho = 0.69$; for vessel count, the correlations are lower at $r = 0.58$ and $\rho = 0.63$.

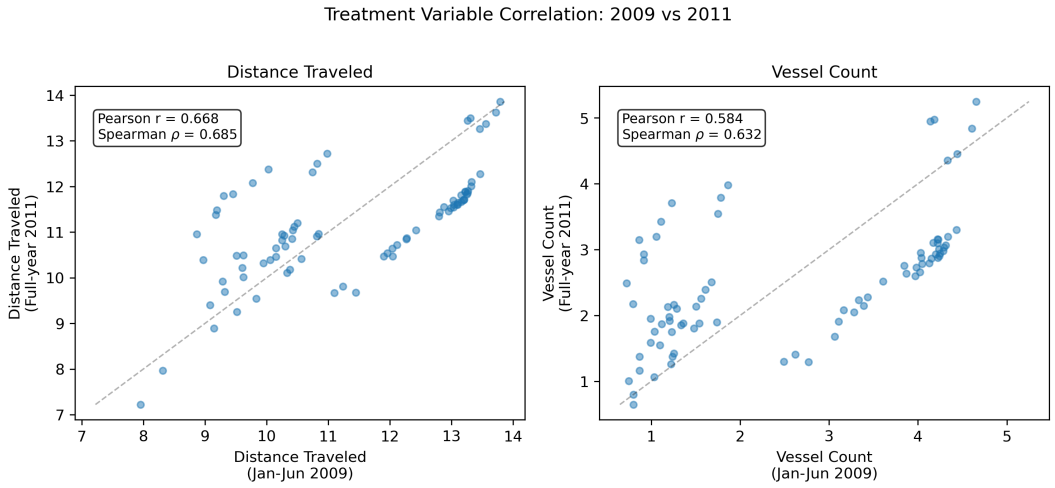

Notes: Each point represents a kelp hexagonal cell. The x-axis measures log shipping exposure during January–June 2009 (the 2009 ECA treatment variable), and the y-axis measures log exposure during full-year 2011 (the boundary modification treatment variable). Left panel: log mean distance traveled (km) within a 20 km buffer. Right panel: log mean vessel count. Both axes are in natural logs. Dashed line is the 45-degree reference. Pearson and Spearman correlations are reported in each panel.

Figure 15: Treatment variable correlation: 2009 vs. 2011 shipping exposure

The correlations are moderate, indicating that the 2011 boundary modification genuinely reshuffled relative shipping intensities across cells. The scatter plots show clear dispersion around the 45-degree line, particularly for vessel count ($r = 0.58$), where several cells that had moderate traffic in 2009 became heavily trafficked by 2011, and vice versa. This supports the argument that the two natural experiments exploit partially independent sources of spatial variation: the 2011 results are not mechanically driven by the same cross-sectional pattern as the 2009 specification, but instead reflect a distinct

reallocation of traffic generated by the boundary modification.

K Weighted PPML: 2011 Boundary Modification

Table 2 in the main text reports unweighted PPML estimates for the 2011 boundary modification. Table 11 below reports weighted estimates, where observations are weighted by log vessel count (20 km buffer, 2011). The results are consistent in sign, significance, and magnitude: all four coefficients are negative and statistically significant at the 1% level. The vessel distance coefficient is -0.102 (weighted) versus -0.093 (unweighted); the vessel count coefficient is -0.134 (weighted) versus -0.210 (unweighted). The two sets of estimates fall within each other's confidence intervals, confirming that the 2011 results are insensitive to the weighting scheme.

	Distance Traveled (km)		Vessel Count	
	(1)	(2)	(3)	(4)
	Biomass	Area	Biomass	Area
Post 2011 × Exposure	-0.102***	-0.101***	-0.134***	-0.133***
	(0.0251)	(0.0248)	(0.0473)	(0.0466)
SST	-6.821***	-6.703***	-6.876***	-6.757***
	(1.003)	(1.004)	(1.002)	(1.003)
Precipitation	0.0432	0.0432	0.0430	0.0429
	(0.0824)	(0.0814)	(0.0820)	(0.0810)
Significant Wave Height	-3.424***	-3.318***	-3.407***	-3.301***
	(0.978)	(0.973)	(0.978)	(0.973)
Cell FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes
Observations	5856	5856	5856	5856

Notes: Weighted PPML estimates of Eq. (3). The dependent variable is kelp biomass (kg) or canopy area (m^2). "Distance Traveled (km)" columns (1)–(2) measure exposure by logged total distance traveled by vessels within a 20 km buffer around each kelp hexagon during 2011; "Vessel Count" columns (3)–(4) use the corresponding logged vessel count. Observations weighted by log vessel count (20 km buffer, 2011). All specifications include kelp-cell and year–quarter fixed effects. Standard errors clustered at the cell level in parentheses. Sample: Southern California (32°–36° N), 2005–2016.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Weighted PPML: Kelp outcomes and the 2011 boundary modification