The Price of an Arc:
Rerouting and the Incidence of Shipping Pollution

Xinming Du*

August 2026

Abstract

Environmental instruments price damage only when the source has a fixed address, and a ship chooses its address voyage by voyage. When the Red Sea closed in December 2023, the same goods sailed 4,133 nautical miles farther around the Cape. I solve both corridors on an ocean grid and read the diversion in satellite columns over 53 countries' coasts. A standard deviation of corridor exposure raises city nitrogen dioxide by 2-3.5 percent and moves exposure off 81 million people onto 37 million earning two fifths as much. A tonne of fuel on the Red Sea arc delivers 3.7 times the exposure it delivers on the Cape arc, yet fuel-scaled instruments price the pair at 0.7. I derive the efficient ratio of arc exposure charges, which does not depend on what exposure is worth, and carry it to twenty-eight passages. The same calculation applies to any emitter that chooses among routes.

JEL Classifications: F18, Q53, R41, Q56, H23

Keywords: shipping emissions, route choice, maritime chokepoints, spatial externalities

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

1 Introduction

What an emission costs depends on where it is released and not only on how much of it there is. A tonne of nitrogen oxides leaves a different burden over a crowded estuary than over open ocean, and an instrument that charges the tonne carries none of that difference unless it is written into the price. The classic remedies write it in through the source's address, by charging the damage each stack does (Pigou, 1920), by trading permits defined on ambient concentration at receptor points (Montgomery, 1972), and by differentiating the charge across zones (Tietenberg, 1978; Fowlie, Holland and Mansur, 2012). Every one of them assumes the source has an address. When the emitter chooses where to be, the address becomes an outcome of the same decisions the instrument is trying to price. Ocean shipping carries the great majority of world merchandise by volume (UNCTAD, 2024), and the particulate matter its exhaust forms carries a mortality burden the literature puts in the tens to hundreds of thousands a year, before the sulfur cap and after it (Corbett et al., 2007; Sofiev et al., 2018). A ship is the polar case of the mobile source, choosing its address voyage by voyage in water no state governs. Yet the instruments that reach shipping charge in proportion to fuel and uniformly in space.[1]

Standard frameworks are not built to record that choice. Emissions inventories attribute bunker fuel to the country that sells it or the flag the ship flies, so the water a voyage crossed enters no ledger. Emission control areas map a boundary and cap what may be burned inside it, so they price the fuel a ship uses inside the box and leave the choice of box unpriced. The damage-differentiated pricing that does condition on location assigns damage to a stack whose coordinates are fixed (Muller and Mendelsohn, 2009; Fowlie and Muller, 2019), and the designs that recover who lives downwind identify a firm's choice of site or a household's choice of neighborhood (Banzhaf, Ma and Timmins, 2019). Two routings of the same cargo between the same ports, one through a narrow sea flanked by cities and one around a continent through open water, are not differentiated by any instrument now in force, or by any account that scores them, according to the populations exposed, even though the people who breathe the exhaust are different populations on different continents earning different incomes. Closing that gap requires a shock that moves the route between the same origins, destinations and products.

This paper prices the missing instrument, the price of an arc, the charge on one stretch of water that would internalize what a transit across it does to the people who live beside it. No road has the scale of this reallocation, since when a corridor moves, the arcs it trades between lie on different continents, so the reallocation runs between populations on different continents, and the counterfactual arc appears in no dataset and has to be solved before it can be priced. The closure of the Red Sea in December 2023 supplies that shock, and I use it to ask two questions, whether

[1] Arc instruments are not new. Routing measures, control areas and speed restrictions attach to stretches of water, and the war-risk premium is a market-set arc price that differentiates by exactly this geography. Nor is a spatially differentiated charge on a mobile emitter new. Road pricing has varied the charge with the time and the place of travel for decades (Vickrey, 1969; Arnott, de Palma and Lindsey, 1993), and the damage a vehicle does varies sharply from one vehicle to the next (Knittel and Sandler, 2018) and with the node that dispatches its power (Holland, Mansur, Muller and Yates, 2016). And an arc-level reallocation of traffic has been read into local air and infant health at toll plazas and along subway lines (Currie and Walker, 2011; Gendron-Carrier, Gonzalez-Navarro, Polloni and Turner, 2022), with the harm depending on who is downwind (Anderson, 2020).

the geometry of an arc measures the incidence of its plume, and what a charge on an arc does to the trade that crosses it. The diversion answers the first question, and a published restriction at a different chokepoint answers the second, the drought that rationed the Panama Canal over the same window. A carrier weighs fuel, time and the tolls and premiums levied on the water it crosses, all of which it pays, while the exposure the same path delivers to a coastline belongs to somebody else and enters no cost the carrier faces.

The efficient instrument built from the two answers fixes the ratio of the two arcs' exposure charges instead of their level, because the value of a unit of exposure is common to both arcs and divides out. After the value divides out, the charge needs only the population beside each arc and the geometry that carries a plume from the water to the shore, read through a dispersion kernel whose single parameter I fit to the satellite column over the water. Both are observable for any arc on the map, so the ordering can be computed for arcs nobody has ever charged. The benchmark I compare against is the set of instruments actually in force, and those are second best twice over, once because a single levy must fold two externalities into one base, and once because a fuel base leaves exposure entirely unpriced. A planner confined to one fuel-scaled instrument can move the level of the charge but never its geography, so no rescaling of the instruments in force reaches the ordering this paper measures, and the estimate I provide reverses the ordering every one of them imposes.

Answering the first question takes variation in the path that satisfies four conditions, a move large enough to be visible in the air, a shock unexpected and set by an agent with no stake in the coastlines on either side, the same origins, destinations and products on both sides of the break, and two sides separated in space. The Red Sea closure satisfies all four. Transits fell 43.6 percent at Bab el-Mandeb and 33.1 percent at Suez and rose 84.9 percent at the Cape of Good Hope, a change in path large enough to be read in a satellite column. The reroute was complete within weeks of the first attacks and it was priced by war-risk underwriters, who bear none of the exposure beneath either arc. The same partners went on exchanging the same goods over a path 4,133 nautical miles longer. The comparison holds origins, destinations, product categories, the technology of the fleet and the location of every factory fixed, while the water the exhaust falls on changes. And the two arcs run past different continents, so the coastline released and the coastline taken on enter the same regression with opposite signs.

A diversion need not be legible on land at all. A lane emits over open ocean hundreds of kilometers from anyone, nitrogen dioxide lives a few hours in the boundary layer, and most of what a corridor emits never reaches a coast, so what survives has to be read against a city's own trend and its country's month, and no one has shown it can be. The direction of the ordering does not follow from the fuel either. The longer detour emits more in total, so the naive ordering makes the Cape routing the dirtier one, exactly the ranking every fuel-scaled instrument implies. It reverses because the two arcs differ in who stands beside them by more than they differ in length. And the incidence could have gone either way. The corridor released a coastline that is both more populous and richer than the one it arrived on, so the aggregate reading and the distributional reading point in opposite directions, and where they cross has to be computed instead of asserted.

The corridor cannot be read off a map, so I solve it. Each lane takes its least-cost path across the ocean twice, once with the canal open and once without it.

A city's exposure is the change in a weight that falls with distance from those paths. The decay length in that weight is measured, not assumed, and the satellite column over the water fits it at 225 kilometers. Which cities sit near a path is fixed by the shape of the coastline and the location of the two ports. Two cities on the same coast, alike in size and industry, can therefore differ in exposure because one sits where the shortest water route rounds a headland.

The solver then predicts data it was never given. It has one free parameter, how far off the coast carriers stand, and that parameter is set on three counting points. At other gates the solver predicts which way the traffic would move, and the counts agree. The satellite column is a second test of the same kind, since nothing in the retrieval knows where the paths are, and the change in the column peaks on the solved corridor and decays away from it.

A corridor delivers cargo as well as exhaust, so what looks like a plume could be the business a port does. Cities the ships merely pass receive exposure and no cargo, so the split between passing and calling cities separates the two. The wind supplies a second separation, since a plume must arrive on the passes when the air moves from the lane to the city, and a harbor's own economy does not depend on which way the air was going.

Because a war and a collapsing currency sit on the emptied corridor's own coasts, anything national is absorbed by construction, and the estimate is identified by one city against another in the same country in the same month. A standard deviation of corridor exposure raises city nitrogen dioxide by 3.5 percent. The two sides of the reallocation carry the same sign and comparable size, at 2.6 percent where the corridor filled and 4.3 where it emptied, the symmetry the incidence account rests on. The response also follows the counted traffic. Read against the traffic gone at each city's own gate, it is 4.7 percent per unit of the gate's shortfall. The most demanding within-country designs put the same response at 2 percent, the two countries with cities on both sides of the reallocation at 1.8 and 2.2, so the headline is a 2-3.5 percent band. Every dollar object below is priced at both ends of it, and every ratio and crossing is free of the coefficient entirely.

The wind separates a plume from a coincidence of geography. On the passes when the air was moving from the lane to the city the response is 3.0 percent, against 1.4 on the passes when it was moving the other way. On the pre-closure passes that premium is 0.5 percent, so an onshore pass carries little for any coastal source until there is a plume to carry.

The composition of the plume identifies what burned. Across four pollutants read on the same city-months, nitrogen dioxide rises 0.099 and carbon monoxide 0.042 while sulfur dioxide falls 0.070 and the absorbing aerosol index falls 0.024. Ship exhaust after the 2020 sulfur cap is high in nitrogen oxides and low in sulfur, and the hulls that took the long route are the large ones whose economics favor a scrubber, which washes sulfur into the sea and does nothing for nitrogen oxides. An index that falls where nitrogen dioxide rises is neither smoke nor dust. Formaldehyde, which burning vegetation makes and ships barely do, holds flat at 0.3 percent.

The extra activity on those coasts came from calls and bunker sales, while cargo throughput barely moved. In the count form the zeros call for, calls at the corridor ports rise 10.7 percent

for a standard deviation of exposure while the cargo those ports handled rises 1.8, a fifth of the pace. The sharpest movements are in the tanker calls at the bunkering ports the detour runs past, 45.3 percent at Cape Town and 28.5 at Las Palmas against 4.3 at Singapore, which sits on neither corridor. Singapore's own bunker sales, a government series with no connection to the transit counts, rose 13.9 percent. Cities the ships merely pass take the same plume as the cities they call at, 3.6 percent against 3.4 with the pooled 3.5 between them, the test that allows the corridor weight to be read as plume incidence instead of harbor activity.

The reallocation moved exposure off 80.9 million people on the Red Sea and Mediterranean corridor onto 36.7 million on the Atlantic and around the Cape, and off a further 38.3 million on the coasts at war, priced at the same coefficient. Of the 36.7 million, 5.4 live in cities with no port of their own, receiving the plume and none of the trade. The coastline that took the plume earns 4,725 dollars a head against 11,748 on the coastline released. At equal weights the diversion is an improvement, and under a weighting that favors the poor strongly enough it is a loss, with the two readings crossing at an inequality aversion of 1.38. Over the empirically standard range of aversion the verdict turns on how the poorest released coasts are valued. Because weighting exposure by inverse relative income makes the exposure-charge ratio of the two arcs a schedule in the same parameter, the aversion at which the verdict flips and the aversion at which the two arcs price equally are one number.

Pricing an arc needs one more input, the trade response to a known change in its cost. The Red Sea closure is a single step at a single date, and the drought over Gatun Lake supplies the graded one. As the lake fell, the Panama Canal Authority cut the maximum authorized draft in the Neopanamax locks from 15.24 meters to 13.41 and restored it over the following year, announcing each cut as a numbered advisory carrying the new limit and the date it took effect. The draft is thus the margin that binds on every ship that crosses and the one published with dates. In a panel of importer by exporter by chapter by month cells with three sets of time effects, a pair loses 7.0 percent of its trade for each thousand nautical miles the canal saves it and each full unit of draft shortfall. At the deepest cut the United States and China trade 3.6 percent less. The episode displaced 162 billion dollars of trade from the routes that cross the isthmus.

A tonne of fuel burned on the Red Sea arc delivers 3.7 times the population exposure of the same tonne burned on the Cape arc. Yet every instrument that scales with fuel charges the Red Sea arc 0.7 times the Cape arc, because a levy on fuel is a levy on distance and the detour is longer. A fuel base puts a weight of zero on where the exhaust lands. Per transit of a given hull the exposure charges stand at 2.6, a third of the premium from reaching more people and two thirds from running closer to them. The two instruments disagree about which arc should cost more, not merely by how much, and the disagreement survives every income weighting up to the crossing. Because the rule needs a population and a geometry and nothing else, I carry it to the twenty-eight passages the transit data covers (International Monetary Fund, 2024), against the throughput rankings of Verschuur, Lumma and Hall (2025). There the gate price of a transit varies by a factor of thirty-six among the heavily used passages and Panama prices at a quarter of Suez. The ordering is not a generic case for taxing chokepoints, since below the indifference value the efficient correction on

Suez is negative and the canal should be cheaper, because the carbon a transit saves outweighs the exposure it delivers.

This paper contributes to four literatures that meet at the leg between origin and destination, and none of them prices that leg. First, damage-differentiated pricing is the closest antecedent. The damage a tonne of emissions does has been computed stack by stack across the United States economy, and pricing each source at its own damage rather than uniformly is welfare improving when costs and damages are known (Muller and Mendelsohn, 2009; Muller, Mendelsohn and Nordhaus, 2011; Fowlie and Muller, 2019). In that whole line a source is a stack whose location is given. A ship chooses its location among arcs, so the charge has to attach to the arc it takes. The valuation literature supplies the frame in which exposure has a price at all, through housing markets, mortality identified on wind direction, mortality downwind of a highway and nowhere else around it, the health toll of taxiing aircraft and the ozone a port pushes over a city (Chay and Greenstone, 2005; Deryugina et al., 2019; Anderson, 2020; Schlenker and Walker, 2016; Moretti and Neidell, 2011). The price it recovers is the parameter that divides out of the ratio I estimate. On the distributional side, pollution and disadvantage come to be correlated through firm siting, household sorting and the politics of enforcement (Banzhaf, Ma and Timmins, 2019), and the reallocation studied here operates through none of them. No household moved and no plant was resited, so the account is short run by construction, and the correlation between exposure and income appears anyway because a carrier minimizing cost chose a different arc.

Second, trade and the environment has measured the emissions embodied in what crosses borders and the welfare arithmetic of taxing them. The question is organized around what a country produces once it opens, through the scale, composition and technique margins (Antweiler, Copeland and Taylor, 2001; Copeland and Taylor, 2004). Yet the carbon burned in shipping what is traded is small beside the gains from that trade (Shapiro, 2016), and trade policy is itself tilted toward dirtier goods, an implicit subsidy to traded carbon (Shapiro, 2021). The freight leg's own emissions have been priced by mode and by the distance a country pair implies (Cristea et al., 2013), the closest the literature comes to the path itself. In all of it the path is a technical coefficient attached to a pair, fixed once the pair is named. A diversion moves it.

Third, the economics of shipping has made the carrier's problem explicit. Ocean freight costs have not fallen the way the usual account supposes, trade costs are endogenous to where ships find their next cargo, and the hub structure of the network sets what a route costs (Hummels, 2007; Brancaccio, Kalouptsidi and Papageorgiou, 2020; Ganapati, Wong and Ziv, 2024). The canals have carried natural experiments before. Feyrer (2021) reads the 1967-1975 closure of this same canal as a distance shock, with trade falling in the pairs the closure lengthened and most of the adjustment complete within a few years. The Panama design is the complement rather than the same number, a published rationing over eighteen months instead of an eight-year change in distance. Days at sea carry a price of their own, the private cost the detour burns alongside the fuel (Hummels and Schaur, 2013). The closure studied here has its own emerging trade-side literature, tracking the collapse in Suez transits and the pressure on capacity and freight rates (Dunn and Leibovici, 2024).

That literature prices the arc to the carrier. What a transit costs the people beside it does not enter.

Fourth, transport networks in spatial equilibrium have made the network itself the object. The literature solves trade and welfare over a continuous topography, prices improvements to the network's own links, recovers the gains a railroad network delivered, and solves for the network a planner would build (Allen and Arkolakis, 2014, 2022; Donaldson, 2018; Donaldson and Hornbeck, 2016; Fajgelbaum and Schaal, 2020). The arcs there are given by history or bought out of a building budget. Here they are opened or closed by agents who are not the planner and priced by tolls, insurance and war, and the choice moves not output but the coastline that receives the exhaust. The toll schedule I trace is the network counterfactual for this pair, lane by lane, and the released weight matrix with the arc condition at every gate is the full-network version, ready for whoever wants to run it.

That the diversion is visible from orbit is not itself new. The shift in the tropospheric column has been read over the Red Sea, Gibraltar and the ocean south of Africa (Pseftogkas et al., 2024), and the same rerouting bounds the cloud response to the sulfur cap over the southeast Atlantic (Diamond and Boss, 2025). Air quality has been documented at African port cities along the route that filled (Zhao et al., 2026), and the Panama restriction has been carried through traffic to pollution on the United States west coast (Lee, 2025). Where the samples overlap the readings agree, since those papers map the same water this paper's counting box and corridor means read. And the port-city movements they report sit inside the coastal response estimated here, so the atmospheric record corroborates the inputs. None of them asks what a transit on one arc instead of the other is worth charging, which needs the exposure a corridor delivers to people rather than the change in a column. It needs the trade response to a price on an arc instead of the traffic response to a closure of one.

Beyond the estimates, I contribute the measurement itself, since what a coastline is exposed to has never been observable, the corridor being a datum in no distributed product. The schematic route layer runs single segments of thirty thousand kilometers, while a count at a gate says nothing about the water between two gates. Solving the corridor on an ocean grid produces the missing object, a weight from every stretch of water to every coastal city, in both states of the world, the counterfactual one included, validated against data that never enters it. Against the closest antecedent, the receptor-level pricing of a fixed source, the kernel replaces a full chemical transfer matrix with one measured parameter. The kernel can be built for an arc that does not exist yet, which is what pricing a routing requires. The solver and the weight matrix are released with the paper and run for any arc on the map in either state.

The policy options the ordering opens are specific and legally constrained. Transit passage through an international strait cannot be charged under the law of the sea,[2] and the Convention of Constantinople of 1888 that keeps Suez open to every flag constrains how far even a canal's own tolls may discriminate (Convention of Constantinople, 1888). The dock remains chargeable, and

[2]The one carve-out is a charge for a specific service rendered to the ship, which is how the Turkish Straits collect under the Montreux Convention of 1936 (Montreux Convention, 1936). An exposure charge cannot be written as one, since what this paper prices varies with where the plume lands instead of with any service the ship receives.

conditioning at the dock on the character of a voyage is settled practice. The European scheme already prices the voyage itself, with the sign reversed, since a Cape routing of the representative crossing pays 201,000 euros against 139,000 for the same voyage through Suez. Thus the one instrument that already conditions on a voyage charges the low-exposure arc the premium in proportion to its length. Reaching past a voyage's endpoints to the arcs it traversed has been tried once and reversed.[3] The lesson is that a charge reaching water outside the charging state's jurisdiction is a trade instrument as well as an environmental one, and the constraint that binds is the willingness of other states to accept it.

The remainder of the paper proceeds as follows. Section 2 describes the Red Sea closure and the instruments now in force. Section 3 sets out the conceptual framework and derives the efficient ratio of arc exposure charges. Section 4 describes how I solve the corridors and build the three panels. Section 5 presents the empirical strategy. Section 6 reports the effect of the diversion on coastal air quality. Section 7 separates the plume from port activity. Section 8 computes the incidence of the reallocation across coastlines. Section 9 assembles the arc price and carries it to every monitored passage. Section 10 concludes.

2 Background

2.1 The Red Sea Closure

Traffic through Bab el-Mandeb carried 8.7 percent of seaborne trade by volume (UNCTAD, 2024) and the Suez Canal moved 72 transits a day on its own count (Suez Canal Authority, 2025) before December 2023, across the shortest water between Asia and Europe. The strait narrows to a few miles at its tightest, and every crossing lane passes through that one water, so a threat to shipping there raises the cost of the whole arc at once. The closure came in steps, beginning with attacks on merchant shipping in the southern Red Sea in November 2023. Within weeks the war-risk premium on that water moved from a twentieth of one percent of hull value to between seven tenths and one percent.[4] Thus a hull worth 150 million dollars paid between one and one and a half million to make a single transit. The charge was triggered by entering listed water and did not depend on what the ship carried, so the price of one stretch of sea moved while nothing about the goods on it did. By January 2024 the major container lines had suspended transits and were routing Asia-Europe services around the Cape of Good Hope. A price set by a party that bears none of the pollution exposure on either coastline can move the traffic without being moved by it.

Transits from December 2023 fell by two fifths at the two Red Sea gates and rose by four fifths at the Cape against the five years before, in the daily counts of Figure S1 at the four passages the design is built on. The ten-point gap between the two gates is traffic that enters the Red Sea

[3] The European Union covered the whole length of every flight touching its airports in 2012 and met retaliation from the states whose carriers flew the covered legs. It suspended enforcement within the year by the stop-the-clock decision and confined the scheme to its own airspace by Regulation 421/2014 (European Parliament and Council, 2014).

[4] The steps are dated in the underwriting press, a nominal 0.05-0.1 percent in the quiet years and often waived, 0.3-0.5 percent by late December 2023, and quotes of 0.75 rising to 1 percent in the days after the January 2024 airstrikes (Lloyd's List, 2024).

without crossing it, a margin I use to read the exposure dose at each gate separately. On the canal authority's own annual accounts, the net tonnage the Suez toll is levied on fell by three fifths. One competing shock lands in the same month the container lines suspended their transits. The authority raised transit dues by 15 percent for tankers and container ships effective 15 January 2024, by Circular 7/2023 issued on 16 October 2023 (Suez Canal Authority, 2023a). Because the circular came two months before the diversion, the rise was decided in a world where the corridor was open. The companion Circular 8/2023 exempted container ships sailing directly from north-west Europe to the Far East, the trade the crossing lanes carry, through June 2024 (Suez Canal Authority, 2023b). While that exemption stood the treated lanes paid no part of the rise, so canal pricing could not have moved them. On the lanes the rise did reach, a dearer canal pushes traffic toward the Cape exactly as the attacks do, and of the two shocks the toll is if anything the smaller.

The reroute changes the path the same pairs sail. After the diversion the traffic-weighted crossing lane runs 4,133 nautical miles farther, around the Cape instead of through the canal. A carrier diverts only when the effective charge on the canal arc exceeds the cost of the extra distance, so the observed switch puts a lower bound on what the Red Sea arc came to cost. That bound, the fuel and the days of the detour the diverted hull actually sailed, is an outlay of the same order as the toll the canal itself charges. Since almost every crossing lane switched, the charge exceeded the detour cost on each of them separately, so the tightest bound is the longest detour instead of the mean one. That lane runs from Jebel Ali to Ambarli, 3,934 nautical miles through the canal against 12,321 around the Cape, and its hull burned 1,241 tonnes to avoid the Red Sea, the top of the 140 lane-level bounds ranked in Figure S2. To keep each bound a floor, I price the detours as fuel alone, dropping the time and the capital a full accounting would add. Even the longest detour's floor sits below the lower edge of the premium band, so on every lane the quoted premium exceeds the cost the carrier revealed by diverting, exactly the pattern a market-set arc price implies.

The Red Sea held half its old transits through 2024 and had won back only part of that loss by the middle of 2026, a stop-start path that enters the panel through the transit count. I count the recovery at the chokepoint rather than model it, so the response is measured against the traffic that actually returned instead of against an assumed end date.[5]

A second conflict crossed the window's final months. The Strait of Hormuz closed to commercial transit after strikes on Iran in late February 2026, in the strait's own count a fall from 77 transits a day to between 3 and 5 across March, April and May of 2026 with a tenth returning by June. Gulf services were suspended, tankers anchored in the roads, and Jebel Ali, one of the 14 endpoint ports and the origin of the longest-detour lane, was itself struck. Cape traffic from March to June 2026 therefore includes ships avoiding Hormuz alongside ships avoiding the Red Sea, and Appendix S1 re-estimates every reading on the months that close before it.

[5]The return timeline is externally traceable service by service. Attacks ended in the autumn of 2025, and a first liner crossed on the MECL string that December (Maersk, 2025). CMA CGM returned its INDAMEX service with a sailing departing Nhava Sheva on 18 January 2026 (CMA CGM, 2025), and the ME11 service followed with first departures in the first week of February (Maersk and Hapag-Lloyd, 2026). The strikes of late February then sent the strings around the Cape again. A memorandum of 14 June 2026 consolidated the United States and Iran ceasefire, and no new attack on shipping had been reported by the June 2026 freeze this paper's institutional statements are dated to (Security Council Report, 2026).

The canal itself closed once before, for six days, when a grounded container ship shut it in the spring of 2021 with no war, no insurance premium and no geopolitics. At the charter rates then prevailing a day at anchor cost more than a day of detour steaming, so carriers began rounding the Cape within days of the grounding. Six days are invisible to a monthly panel and plain at the resolution of the satellite pass, which makes that episode a rehearsal with a different cause, the same geometry and the same wind. Over the months of the Red Sea closure a second canal lost capacity for an unrelated reason, a drought over Gatun Lake that cut draft but closed no passage, the second priced shock of the paper.

2.2 Existing Instruments

No environmental instrument in force prices the choice among arcs. The net-zero framework the International Maritime Organization approved at the committee stage in April 2025, and has not yet formally adopted, would price greenhouse gas intensity per unit of fuel energy on a well-to-wake basis, uniformly in space.[6] In January 2024 the European emissions trading system extended to maritime transport, and it charges a ship on the character of the voyage that brings it to a European dock, in proportion to the fuel that voyage burns.

Seven emission control areas are in force as of the June 2026 freeze, each capping what a ship burns inside a mapped boundary. The Mediterranean has been under a sulfur limit alone since May 2025, and the Canadian Arctic and Norwegian Sea both entered force on 1 March 2026 with their fuel-sulfur limits enforceable twelve months later. An eighth, the North-East Atlantic, is adopted with force from September 2027, covering nitrogen and particulate matter as well as sulfur. The North American zone measurably improved the air and the health of the coasts inside it (Hansen-Lewis and Marcus, 2025). A mapped boundary already moves ships, not only fuel, since hulls lengthen their paths to skirt the California zone instead of burning the cleaner fuel inside it (Klotz and Berazneva, 2022). That is the routing margin this paper prices.

The framework, the trading system and the control areas all scale with fuel or with what fuel contains, so each prices an arc in proportion to how far it is sailed. Canal tolls are the one instrument that charges the arc itself, a genuine arc price set by a single identifiable agent, levied on net tonnage and vessel type without reference to anything downwind. Arc pricing therefore has a working precedent, and the gates on this map go unpriced because treaties made their passage free, not because a gate cannot carry a toll.[7]

3 Conceptual Framework

An *arc* is a stretch of water a carrier may use or avoid and someone may price. A *lane* is a pair of ports. A *path* is the sequence of water a lane takes in one state of the world, and it is what the

[6] The vote was adjourned in October 2025 and rescheduled for October 2026, with the earliest entry into force in March 2028, and its design charges fuel wherever it is burned, the feature that matters here.

[7] Denmark charged most ships entering the Baltic at the Sound for over four centuries, the Copenhagen convention of 1857 bought the dues out (Copenhagen Convention, 1857), and the modern free-transit norm begins there.

solver returns. A *route* is a lane together with the path it takes. The bundle of paths that run through one arc is a *corridor*, and a city is near a corridor, not a path.

3.1 The Carrier's Route Choice

A carrier moving a unit of trade from origin o to destination d selects a path p through a network of arcs $\mathcal{A}$ to minimize

$$c(p) \;=\; \sum_{a \in p} \big(f \ell_a + \tau_a \big) \tag{1}$$

where ℓ_a is the length of arc a, f the cost per unit distance covering fuel, time and capital, and τ_a an arc-specific charge covering canal tolls, insurance and port fees. The model strips the choice to a carrier picking among paths for one voyage, the margin the shock moved, and leaves out the network, speed and frequency, margins that move on longer horizons than a closure announced in a month. Each omitted margin returns in the empirics as a check on the estimate and not as part of the model, capacity absorbed by the longer voyage in the control corridor, speed in the convoy variant, and the network in the transshipment hubs. Let p^* be the least-cost path with the canal arc available and p° the least-cost path without it. Diversion occurs when

$$\tau_{\text{canal arc}} \;>\; f\big(\ell(p^\circ) - \ell(p^*) \big) \tag{2}$$

so the observed switch bounds from below the increase in the effective cost of the Red Sea arc, the quantity a canal authority would need to price it and does not hold. I recover the distance difference by solving both paths.

3.2 The Dispersion Kernel

Emissions occur along the path and reach the coast on the wind. City i receives

$$P_{it} \;=\; \sum_{a} \omega_{ia}\, q_{at}, \qquad \omega_{ia} \;=\; \frac{1}{D} \int_a \exp\!\big(-\delta_i(s)/D \big)\, ds \tag{3}$$

where q_{at} is traffic on arc a in month t and $\delta_i(s)$ the distance from the city to the point of the path at arc length s, so a lane that runs beside a coast for days counts for more than one that crosses at a point. Only the geometry of the path enters the kernel. The wind enters where it can be checked, in the response I read against it on each satellite pass. Because the reroute moves the traffic term and leaves the weight untouched, the weight serves as a fixed intensity and the traffic as the shifter. The weight is a dimensionless intensity, not a share, since it is not normalized to sum to anything over cities or over arcs. That convention cancels out of every ratio the paper reports and enters dollar levels only through the indifference chain that prices them.

The kernel carries chemical loss and no geometric spreading, so a concentration that in a physical plume falls with dispersion as well as with decay is made to fall with decay alone. That overweights long lanes at moderate distance against short lanes close in, and the Red Sea arc is the one with short lanes close in, so the exponential understates the Red Sea arc's advantage. And the kernel

is isotropic while the two coastlines have different prevailing flow. The pass-level test reads the response against the wind directly, and a repriced ratio scales each city's weight by the share of hours the air actually carries the plume ashore.

The kernel weights the tracer, while the damage the mortality literature attributes to ship exhaust is secondary particulate matter, which forms over a day or so as nitrogen oxides oxidize and find ammonia, downwind of where the column is read. The transport-chemistry accounting prices ship-attributable mortality through the particulate pathway (Sofiev et al., 2018). That chemistry has a geography, and the geography runs against the headline, not for it. Nitrate formation is ammonia-limited over most coasts. The IASI record catalogs industrial ammonia hotspots on the Nile delta and the Gulf of Suez (Van Damme et al., 2018), while the upwelling coasts the diversion moved onto carry none of its cataloged hotspots. As a result, a map from tracer to damage that respected the ammonia would weigh the exposure removed more heavily than the exposure delivered, and the equal-treatment reading the ratio uses is the smaller of the two.

From the tracer the ratio needs only an ordering, because it survives any monotone map from tracer to damage that treats the two coastlines alike. The distance profile shows the tracer decaying within the fitted length of the water, so the kernel weights the near field, where the population also is. Running both corridors through a plume model in the real wind and boundary layer fields recovers the kernel's ratio from the dirtier side, and I check the particulate's displacement in the same plume model, since the monotone-map condition covers the chemistry of the conversion and not the geography of where the damage lands.

The fleets themselves also differ across the two arcs, and that margin can be signed, not left inside the monotone-map condition. The largest classes round the Cape, carrying the scrubber-fitted share, and the sulfur split shows their chemistry directly, sulfur dioxide falling where the ships arrived and not where they left. A scrubber removes the sulfur and not the nitrogen, so the exhaust delivered to the gained coastline carries less of the sulfate precursor per tonne of fuel than the exhaust removed from the released one, and a map from tracer to particulate damage that respected the fuel chemistry would run the same direction as the ammonia geography, weighing the removed exposure more heavily. The per-tonne unit itself is unmoved by hull composition, since nitrogen oxides emitted per unit of engine work are a property of the slow-speed engine class both fleets run, recorded in the tier emission factors of the IMO greenhouse gas study (International Maritime Organization, 2020). Hull size divides out of the per-tonne ratio and enters only the per-transit readings, where the observed-fleet variant prices it.

Damage also depends on what is breathed at the surface while the column is a vertical integral, and the map between the two is the depth of the layer the air mixes through. That depth is not the same on the two coastlines. A narrow sea flanked by desert mixes deep in the warm season, while the cool upwelling coasts the diversion moved onto sit under a shallow capped marine layer, so a given column change is a larger surface change on the coast that gained. The depth itself is measurable and enters the estimation as a direct check.

The decay length D is a physical length, measured from the column over the water itself, not assumed. I build the corridor weights at 150 kilometers, below the fitted length, so that one kernel

serves the estimation, the mechanism tests and the pricing. Weights built short of the fitted length make the coefficient read conservatively. And the common-kernel condition behind equation (5) is testable on the water itself. Where it fails, it fails in the direction that makes every common-kernel reading in this paper, in the estimation and in the pricing alike, the conservative one.

3.3 The Efficient Arc Charge

A planner choosing arc charges internalizes what the traffic on each arc does to the people it reaches. Writing N_i for the population of city i and θ for the loss a unit of exposure imposes on one person, the damage arc a does per unit of traffic, written $\mathcal{D}_a$ to keep it apart from the decay length D, is

$$\mathcal{D}_a \;=\; \theta \sum_i N_i \, \omega_{ia} \tag{4}$$

and the efficient exposure charge on that arc is proportional to it. The level of $\mathcal{D}_a$ needs θ, while the ratio between two arcs does not, because θ is common to both and divides out

$$\frac{\mathcal{D}_a}{\mathcal{D}_{a'}} \;=\; \frac{\sum_i N_i \, \omega_{ia}}{\sum_i N_i \, \omega_{ia'}} \tag{5}$$

Everything on the right is observed once the corridors are solved.

Remark 1. *If the loss a unit of exposure imposes is common across the people the two arcs reach and constant in the level of exposure, if each arc carries the plume to the coast by its own kernel measured off its own water, and if carbon is priced by its own instrument, the efficient relative exposure charge on those arcs is identified by population and geometry alone, without any estimate of that loss. If instead the loss varies with income at elasticity η, the relative exposure charge is identified by population, geometry and the income distribution, and the level of the loss still cancels.*

The remark is stated in its general form, one kernel per arc, and the kernels are the object I measure, not assume. Fitted off its own water each corridor returns its own length, so the common kernel at one length is an implementation choice and the conservative one, and every headline is the smaller of the two readings. The inverse-column variant relaxes the level condition, and the second instrument of equation (6) below meets the carbon condition.

Entering equation (4) multiplicatively and common to every arc, θ cancels in equation (5). Let arc a carry emissions e_a per unit of traffic and exposure $\mathcal{D}_a/\theta$ per unit of traffic, and let p_c be the price of carbon. The planner who may set both a uniform charge on fuel and an arc-differentiated charge on exposure solves for

$$\tau_a^* \;=\; p_c \, e_a \;+\; \theta \sum_i N_i \, \omega_{ia} \tag{6}$$

on every arc, and the second term is the one no instrument now in force contains.

With one instrument the planner must choose a weight between the two objectives and the arcs are ordered by whichever wins. Given both instruments the ordering problem disappears, because

the fuel charge prices distance and the arc charge prices who lives beside it. The missing instrument is not a different carbon price. It is a second charge whose base is the population a transit passes. The ratio orders the two arcs, and it shows when an instrument orders them the other way, as every fuel-scaled charge does here. Because routing responds to the difference between charges rather than to their ratio, implementing the ordering still requires a value for θ, and everything this paper reports about relative prices is free of that value while everything it reports in dollars carries the indifference calculation that supplies one. Neither object requires a shock. The diversion's role is to establish that ω_{ia} is plume incidence, the common-loss condition relaxes to the income form the remark states, and the rule carries to passages where nothing has happened.

Remark 2. *An instrument whose charge is proportional to fuel prices two arcs in the ratio of their lengths, so it orders them correctly only if the arc that delivers more exposure is also the longer one.*

Fuel burned on a voyage follows from deadweight, distance and the efficiency ratio the fourth greenhouse gas study reports (International Maritime Organization, 2020). The same hull sails both arcs when a voyage is diverted, so a fuel-scaled charge sets the relative price at the ratio of route lengths. Observed fleets differ across the two arcs, the larger hulls round the Cape, and pricing the instruments at the sizes actually sailing corrects for the difference. A chokepoint is by construction the shorter way and the more confined water. The two remarks therefore point in opposite directions whenever the confined arc is short, and how far apart they point is an empirical magnitude, not a matter the theory settles.

3.4 The Planner's Problem

The total efficient charge comes from one first-order condition. Let q_a be the traffic that routes on arc a. Carriers are price takers who route by equation (2), so a charge τ_a enters their delivered cost one for one and the traffic on the arc responds with elasticity ε to that cost. Shippers earn surplus $S(q_a)$ with $S'(q_a)$ equal to the delivered cost they face. Each transit delivers $\sum_i N_i \omega_{ia}$ of exposure at a loss of θ per unit and burns e_a of carbon at a price p_c, and both fall on people who are not party to the transaction. The planner chooses τ_a on every arc to maximize

$$\sum_a \Big[S(q_a) - c_a q_a - \theta q_a \sum_i N_i \, \omega_{ia} - p_c \, e_a q_a \Big] \tag{7}$$

where c_a is the resource cost of a transit.

The charge carries no cross-arc term, since the exhaust a transit lays down reaches only the water that transit crosses. That separation holds on each arc's exclusive water, since a coastline like Iberia's takes exhaust from both routings and its exposure does not move when a charge moves traffic between them. A lane delivers its common-water exposure under either routing, so the shared water drops out of the margin, canceling from the difference $\tau^*_{suez} - \tau^*_{cape}$ that routing responds to. The theory-consistent ratio for the routing margin is therefore the one computed on exclusive water, while the full charge levels keep the shared water in. Both carry the ordering, because

common levels cancel from the difference. The robustness variants relax the assumptions below one at a time.

Assumption 1. *Carriers of a given hull take charges as given and route by equation (2), so routing responds to the difference between arc charges.*

Assumption 2. *Damage is linear in traffic at a loss per unit of exposure common across the people an arc reaches, with no other distortion on the routing margin.*

Assumption 3. *A transit's exposure lands on the coast it passes, so water shared by both routings cancels from the difference between the two arcs' charges.*

A single price per unit of emission implements the optimum only when the damage per unit is the same everywhere it is released (Baumol and Oates, 1988). That is the condition two arcs with different coastlines beside them fail. When carbon is priced by its own instrument, the remaining wedge between two arcs is the exposure term alone, and equation (5) identifies its ratio without θ.

The carrier, the planner and the canal authority each put a charge on the same arc. The carrier faces τ_a, the sum of tolls, premiums and port fees actually levied, and routes by equation (2). A planner would set τ_a^* from equation (6). The canal authority is the one agent on this map that prices an arc and can adjust it, setting τ_a^m to maximize its own receipts, an inverse-elasticity markup on the tonnage that crosses. The three differ in more than level.

Proposition 1. *For carriers of a given hull, routing responds to the difference between arc charges, so a toll levied at one gate and set to the difference between the two efficient charges implements the efficient routing margin exactly. It cannot implement the efficient level, because it falls only on voyages that use the gate and so distorts whether a voyage is made and where it terminates, and because it never reaches a lane that passes no gate at all.*

Differentiating equation (7) under Assumptions 1 and 2 and using $S'(q_a) = c_a + \tau_a$ gives equation (6) arc by arc. A toll at the gate enters delivered cost only on the gated arc, so a toll equal to the difference between the two total efficient charges reproduces that difference on the routing margin, which under Assumption 3 is computed on exclusive water. The level margin needs each arc charged its own damage, which one gate cannot do, and a lane that crosses no gate faces no charge at all.

A canal authority prices the arc through its own water and no agent prices the arc around the continent, so τ_a^m is positive on one arc and zero on the other by construction. Since the priced arc is the one that delivers more exposure, raising the toll does move exposure in the direction efficiency wants. It cannot separate the two margins. A planner setting τ_a^* from equation (6) charges each arc its own damage and leaves the traffic where the charges put it. In contrast, a gate toll has one number to set and buys every unit of exposure it removes with traffic pushed onto the arc it cannot reach.[8] With carbon priced separately, efficiency requires the exposure component of the charge

[8]The Suez authority in practice posts a banded schedule by tonnage and vessel type, rebates particular trades, and was discounting through this window to draw traffic back. Every one of those numbers is levied at the same gate, so none can reach a lane that never passes it.

on the two arcs to stand in the ratio of equation (5), which is above one against the Red Sea, so the instrument that would implement it has to reach an arc that passes no gate.[9]

Equation (4) treats the kernel as proportional to the exposure a transit on arc a delivers to city i, a reading that fails if the data pick up the port and not the plume. A corridor delivers exhaust to every city it passes and cargo only to cities with a port, so write Π for the cities the ships pass without stopping and Σ for the cities they call at. Cities in Σ receive both the exposure and the port access that comes with a call, while cities in Π receive only the exposure, so if the estimated response to the kernel is the same in both sets, the kernel measures the plume and not the traffic in the harbor. In a spatial equilibrium the two sets would eventually part ways in prices, port access capitalizing into wages and rents (Roback, 1982) and a permanent arc charge capitalizing the removed plume the same way. Yet a diversion measured over a few years prices neither margin, so the split between passing and calling separates the plume from the port over a window in which neither capitalization arrives.

4 Data

4.1 Solving the Corridors

The ocean grid comes from the Natural Earth physical vectors (Natural Earth, 2022), rasterized to a tenth of a degree, with each cell joined to its eight neighbors by great-circle edge weights. The grid spans 35 degrees west to 130 east and 45 south to 62 north, which holds every lane either state can take. Cells near land take a multiplicative penalty because carriers stand off the coast, an edge multiplier that decays exponentially in the cell's distance to shore with a 30 kilometer scale and a unit penalty at the waterline, the one free functional form the calibration sweep varies. One Dijkstra solve from an origin returns the distance and predecessor tree for every ocean cell, and 24 origins are solved in each state of the world. Four waterways are narrower than one cell and enter as explicit arcs, Suez, the Bosporus and the Dardanelles, without which the Black Sea is a separate ocean, and Panama.[10] Every solved lane steps between neighboring cells, so the two corridors carry the same overhead in open water and it cancels from the comparison of their lengths.

Endpoints are the 14 busiest container ports east of the canal and the 10 busiest west of it, ranked by daily container calls over the five years before the closure, and pairing every eastern port with every western counterpart gives the 140 crossing lanes, with the ports listed in Table S1. A lane carries the product of the calls at its two ends, because the transit data counts hulls at fixed points and never follows a ship from one port to another, and a length-normalized variant of the weight is computed beside it. Shanghai and Yangshan enter separately in the baseline and merge in the agglomeration variant. Another 30 lanes join the busiest ports on each side to each other, read

[9] A congested gate puts a second wedge between the revenue-maximizing and the efficient toll, and the anchorage hours it generates emit beside the gate city, the margin the canal-bank variant prices. The statements here are therefore about geography, and the queue is a further reason gate instruments underperform arc ones.

[10] Kiel is not one, so Baltic traffic routes through the Danish straits. The choice does not matter, since a Baltic lane runs the same water in both states.

off the same predecessor trees the crossing lanes use, and their distance does not change between the two states.

The one free parameter is how far off the coast carriers stand. Because the transits that later multiply the corridor weights are counted at three fixed points, the solved paths have to run through those points, and the standoff is set by a rule written on the gates alone, applied once, before any exposure is computed. The rule reads each setting's distance to the counting points in Table S2, computes the miss at the two confined gates, Bab el-Mandeb and Suez, keeps the settings whose maximum confined-gate miss sits within one grid cell of the best, the three settings at four kilometers off Bab el-Mandeb and ten off Suez, and among the survivors selects the smallest Cape miss, 47 kilometers against 70 and 158, the row marked used. Among the admitted settings the traffic-weighted extra distance moves within 87 nautical miles, two percent of the 4,133 the headline carries, and the lanes Figure 1 draws in each state sit on the counting points by construction, with every coastal city beside them at its standardized exposure in Figure 2.

4.2 Coastal Cities, Population and Income

Cities come from the global human settlement urban center database (European Commission, 2024) and the port database. The box each city's column is read over is thinned so no two overlap, merging neighbors closer than 78 kilometers, a radius that comes from the retrieval box itself and not from the 40 kilometers inside which a city is counted as owning a port, and the merge bites hardest on the dense Mediterranean coast.

The estimation sample spans 53 countries with a median of four coastal cities each, and every panel the paper estimates on is set out in Table S3. Ten further countries are counted in the incidence account and excluded from every regression, holding a seventh of the account's population, under one rule fixed before any estimation, that the coast already carried armed conflict, a blockade or a collapsing currency when the ships diverted, each of which moves emissions on its own and none of which the corridor explains. Those coasts hold 74 cities, marked in the city counts Table S4 gives country by country. The full register is printed at the end of the appendix and released with the data, every city with its coordinates, shoreline distance, population and exposure, so the account can be audited city by city.

4.3 Satellite Columns

Nitrogen dioxide, sulfur dioxide, carbon monoxide and the absorbing aerosol index come from the offline product of the tropospheric monitoring instrument (Veefkind et al., 2012) on Sentinel-5P, as gridded and served on Google Earth Engine. The quality screen is the mission's own. Nitrogen dioxide retrievals are kept above a quality value of three quarters and each other product at its own documented cut. I leave non-positive sulfur retrievals as the retrieval reports them, and they exit the sample only where the logarithm is undefined, 5,985 city-months. Each scene also carries the processor version that produced it, and a version change moves polluted scenes and clean ones differently. The closure sits inside one version, in force from August 2022 to November 2024, so the scenes on the two sides of the break are identically processed, in Table S5.

Each gas is read over a box 0.7 degrees on a side centered on the city and averaged to the month, the tropospheric column for nitrogen dioxide and the total column for the other gases. The instrument's ground pixel is 5.5 by 3.5 kilometers at nadir and the composite the archive serves regrids it to 1.1 kilometers. A city-month with no valid retrieval is missing rather than zero, and that is why sulfur dioxide has the smallest sample of the four. The column is the same quantity measured the same way on every coast in the sample, read at the same local time, half past one in the afternoon, so the hour is common to the comparison. A coefficient on the log column is a proportional change, and the median coastal city sits at 20.9 micromoles per square meter, so each percent is a fifth of a micromole on that city.

4.4 Shipping, Trade and Weather

Transits at the chokepoints, port calls, tanker calls and cargo volumes come from the daily panels of IMF PortWatch (International Monetary Fund, 2024), which builds them from vessel positions by the algorithms of Cerdeiro et al. (2020). A call is a berthing instead of an arrival in the roads, so calls at a port and transits at a gate are not the same ships counted twice.[11] The transit series is a raw daily count of whole ships, not a moving average, so the December 2023 break is the four-day step the data show and not the week-long ramp that smoothing would draw. The cargo field is tonnage inferred from vessel positions, not a customs declaration.

The trade panels come from customs records instead, the Panama panel monthly imports by 39 reporters from 162 partners across 96 chapters of the harmonized system and one residual code, assembled from the UN Comtrade monthly service (United Nations Statistics Division, 2026). A second customs panel covers 27 reporters and 103 partners over a window a year and a half longer, and a third, the United States Census port-level series by customs district, counts imports by coastal customs district on each seaboard. The reporter sets hold four fifths of world imports, with a canal-dependence distribution as wide as the world's because the reporters sit on both sides of both canals. Every dollar aggregate is summed over reporters alone, so the displaced-trade figures are floors on the world numbers, and what each panel and the passage counts hold is set out in Table 1.

The Panama Canal Authority's numbered advisories give the maximum authorized draft and its effective date, and Singapore's Maritime and Port Authority publishes monthly bunker sales in tonnes. Wind at the synoptic hour nearest each satellite pass comes from the hourly ERA5 reanalysis (Hersbach et al., 2023), and the temperature and rain in the weather controls come from the ERA5-Land monthly means (Muñoz Sabater, 2019), defined over land only. A coastal box takes the land cells it contains, while the wind, needed over the water, comes from the full reanalysis. I freeze every series a regression or an account reads at the end of June 2026, and every statement about institutional timing is written as of that date. Four exhibits draw their sources past the freeze and say so in their notes, the daily transit figure, the authority's published tonnage, and the carbon table and ledger whose 969 transit-series days run to late July.

[11] The canal blockage of March 2021 shows it directly, since hundreds of ships lay at anchor off Port Said and Suez for six days and the call counts at those ports never left their ordinary one to three a day.

5 Empirical Strategy

5.1 The City Panel

Exposure is the change in corridor weight the closure implies, $\Delta\omega_i = \omega_i^{\text{Cape}} - \omega_i^{\text{Suez}}$, standardized by its standard deviation among coastal cities. At the length the weights are built at, one standard deviation of exposure is the gap between a city 40 kilometers from the corridor and one 200 kilometers away. Thus an estimate is the gap in the column between the near town and the far one, attributable to the route. The estimating equation is

$$y_{it} \;=\; \beta\left(\Delta\omega_i \times \text{Post}_t\right) \;+\; \alpha_i \;+\; \lambda_{c(i)t} \;+\; \gamma_i \cdot t \;+\; \varepsilon_{it} \tag{8}$$

where y_{it} is the log column in city i and month t, Post_t turns on in December 2023, α_i are city effects, $\lambda_{c(i)t}$ country by month effects and $\gamma_i \cdot t$ city-specific linear trends. Inference is clustered on the city, with randomization on the design's own assignment as the primary reading and the parametric, the spatial and the selection-adjusted readings as corroboration. Estimates are in log points throughout and the text writes them as percent. Where a magnitude compounds several inputs, as at the deepest Panama cut, I use the level conversion and the table note states it.

The corridor that emptied runs along a coast where a war was fought over the same window and through a country whose currency collapsed in March 2024. Country by month effects absorb anything national, leaving one city compared against another in the same country in the same month. City trends are required as well, because nitrogen dioxide carries strong city-level secular trends that correlate with where the corridors run. The diversion also absorbed capacity from the world fleet (Notteboom, Haralambides and Cullinane, 2024) and raised freight rates everywhere, which moves every coast together and is absorbed by the month effects every specification carries. The comparison requires only that a city's exposure moved with the corridor and not with its neighbor's. A handful of coasts, Spain and Morocco, sit beside both corridors at once, so their exposure is the difference between two weights that both changed, and I construct the same exposure two independent ways.

The design rests on three conditions. The first is that, once each city carries its own linear trend, exposure is as good as randomly assigned within country and month. A city's corridor weight is a distance to a solved least-cost path, so it is set by the shape of the coastline and the location of the two endpoint ports instead of by anything the city or its country chose. The balance test regresses predetermined outcomes on the assignment, and the permutation and rotation exercises put the estimate against what rearranged coastlines produce.

The second is that, absent the diversion, cities differing in their distance to the solved corridors would have moved in parallel within country. The placebo closures and the control corridor, the 30 within-side lanes whose solved distance the closure leaves unchanged, test both halves of the claim,

the date and the geography. An event study enters the same exposure interacted with half-years of event time,

$$y_{it} \;=\; \sum_{k\neq -2} \gamma_k \left(\Delta\omega_i \times \mathbf{1}[\,\text{half-year} = k\,]\right) \;+\; \mu_{m(t)}\Delta\omega_i \;+\; \alpha_i \;+\; \lambda_{c(i)t} \;+\; \varepsilon_{it} \tag{9}$$

where $\mu_{m(t)}$ lets each calendar month carry its own loading on exposure, and the regression is weighted by city population. Here k counts half-years from the closure, so $k = 0$ runs from December 2023 to May 2024 and $k = -1$, which contains the November attacks, is estimated, not omitted. The base is $k = -2$, the half-year ending in May 2023, since the months just before the closure are already treated. Under the identifying assumption the coefficients before the closure are zero. One date treats every exposed city at once, so there is a single cohort and the staggered-adoption corrections of Sun and Abraham (2021) and Borusyak, Jaravel and Spiess (2024) have nothing to repair here.

The third is that the closure reaches the column only through the traffic on the corridor. The wind supplies its direct test, since a plume must arrive when the air moves from the lane to the city and be absent when it moves the other way. The pollutant signature and the split between cities the ships pass and cities they call at rule out the two channels that remain, biomass burning and the harbor's own business. The calendar, the weather and the carbon scheme that began one month after the closure enter directly.

5.2 The Satellite-Pass Panel

The wind separates a plume from anything else that happens to sit beside a shipping lane. Exhaust reaches a city only when the air is moving from the lane to the city, so a plume has to be larger on the passes when it is, and anything not carried by the air has no reason to move with the direction of the air. That test cannot be run on monthly means. Nitrogen dioxide lives a few hours in the boundary layer, the satellite crosses once a day near local noon, and a coast whose sea breeze reverses twice a day averages over a month into a coast with no wind at all. Those are the coasts the corridor runs along.

The unit therefore becomes one city on one satellite pass with any valid retrieval in the box. At this box size essentially every overpass qualifies, so the screen drops pixels inside a pass and almost never a whole pass. For the 120 coastal cities the corridor reaches most strongly, reading at the pass gives 32 usable passes per city each month after the quality screen, 28-36 for four cities in five, against the one a monthly mean allows. Intermittent monitoring carries a bias of its own (Zou, 2021), a second reason the test runs at the pass instead of at the month.

I read the wind against the bearing from the nearest solved lane of either state, not of one, so the bearing is a fixed feature of a city's geometry and does not move when the traffic does. The wind comes from the nearest of the four six-hourly reanalysis fields, within three hours of the pass.

The dated-wind falsification puts the wind at the pass against air a day old, which carries no plume. The pass-level equation is

$$\begin{aligned} y_{ip} \;=\; & \beta_1\big(\Delta\omega_i \times \mathrm{Post}_p\big) \;+\; \beta_2\big(\Delta\omega_i \times \mathrm{Post}_p \times \mathrm{Down}_{ip}\big) \\ & +\; \beta_3\big(\Delta\omega_i \times \mathrm{Down}_{ip}\big) \;+\; \beta_4\big(\mathrm{Post}_p \times \mathrm{Down}_{ip}\big) \;+\; \psi\,\mathrm{Down}_{ip} \\ & +\; \alpha_i \;+\; \lambda_{c(i)t(p)} \;+\; \gamma_i \cdot t(p) \;+\; \varepsilon_{ip} \end{aligned} \tag{10}$$

where p indexes one satellite overpass of city i and $t(p)$ is the calendar month the pass falls in, so the trend and the country-month effects are the monthly panel's own. Here Down_{ip} is an indicator for the wind at the synoptic hour nearest that pass projecting positively onto the bearing from the nearest solved lane to the city.

I carry the two lower-order interactions because neither is absorbed. City effects take out $\Delta\omega_i$ and country by month effects take out Post_p, but $\Delta\omega_i \times \mathrm{Down}_{ip}$ would otherwise load whatever makes exposed cities more often downwind, and $\mathrm{Post}_p \times \mathrm{Down}_{ip}$ whatever the closure did to the wind itself. The coefficient of interest is β_2, the increment on a downwind pass relative to an upwind pass at the same city, the signature a plume implies and a coincidence of geography does not. A continuous version replaces the indicator with the wind projection divided by its standard deviation, and a placebo version estimates the same split on the pre-closure passes alone, where there is no plume for the wind to carry.

Throughout the paper a *person log point* is one person bearing a change of one in the log of the nitrogen dioxide column for one year, written person-log-point-year where a price attaches to it. The *detectable effect* beside a null is the smallest effect the test had the power to find. The *window rule* is that anything read at the half-year or finer excludes the two pandemic years while anything read as a single pre-and-post contrast keeps them, since the month effects absorb a common shift. Each family that uses the rule is also shown the other way. Accounts that *enumerate* a coastline rather than sample it carry no standard errors, since the uncertainty that remains comes from the assignment of exposure and not from a sample draw. That is why the permutation and rotation readings sit beside every headline. The stars some tables print are a display convention.

6 The Effect of the Diversion on Coastal Air Quality

6.1 The Shock over the Water

The reallocation can be measured over the water before any city enters. The transits PortWatch counts record where hulls went, and the column over the water measures what their exhaust did on the way, the two drawn on one map with a counting box at the Cape in Figure 3. Only one of the eight counting points pairs a box of open ocean with traffic that moved, 70 percent water and a quarter of it more than 75 kilometers from any land, which is why the Cape carries the question. In contrast, the Suez box is the Nile delta, the Bab el-Mandeb box the Yemeni coast, and the Panama box has a city at each end of the canal. The three Caribbean passages are open water whose traffic

moved by a tenth at most against the Cape's four fifths, which makes them a placebo instead of a second estimate.

Two fifths of the tropospheric column over the Cape box belongs to the ships. The column in the counting box at the Cape rises with the transits counted there, in Figure 4 and Table 2. No port sits inside the box, so nothing in the count is harbor activity, and the far-field trend is netted, not assumed, in Table S6. Because the column over the water is a background plus a ship term proportional to the traffic, the elasticity to the traffic reads directly as the ship term's share of the total.

The estimate is 0.413, and the trend reading of Table S6, which puts it at a third, is the other end of what the same identity supports. Errors read with a Newey-West kernel at six lags, set by the serial correlation of monthly traffic and retrievals, leave the estimate where it is. With the two pandemic years out, which a common collapse of shipping and air would otherwise inflate, it is 0.381 in Table 2's own third column, the column I report, and the trend table's own pandemic-out column reads 0.389.

A box drawn around a headland could be picking up the coast rather than the sea. A quarter of the box lies more than 75 kilometers from any land. Keeping only that water moves the elasticity from 0.413 to 0.372 in Table S7, so the box measures the water and not the land beside it. Panel A of Table S8 splits the traffic by hull class to read which class the relationship runs through, one relationship divided by the fleet mix. Panel B of Table S8 replaces hulls by capacity, where deadweight tonnage returns 0.396 against the count's 0.413, so the fuel burned moves the column, not the number of hulls.

The passages whose traffic barely moved check the 0.413. Run at the three Caribbean passages in Table S9, the same regression gives -0.131, 0.099 and -0.180. None can be told from zero. The column fell over the corridor the ships left and barely moved over the corridor they took, drawn against the solved lanes in Figure 3 and averaged within 150 kilometers of each corridor in Table S10, with the error taken over 125 kilometer spatial blocks.

The solver was never fitted to the remaining gates, and their counts check the corridor construction itself. Traffic through Sunda, the Indonesian strait the solved paths substitute toward, rose 19.2 percent, and the 20 monitored passages the diversion never touches moved 5.0 percent at the median in absolute value, with a signed median of -2.6, in Table S11. Gibraltar is the subtlest of the predictions, since in the closed state the Mediterranean lanes come in through the strait while the North Europe lanes leave it altogether, and the count shows the near-cancellation the solved paths imply, a move of 3.4 percent inside the spread of the untouched passages. The European route-density rasters are built from vessel positions this construction never touches, and inside the Mediterranean basin the cargo-class density falls monotonically away from the solved trunk, 3.9 times the basin's background within ten kilometers of the paths, 2.5 times at ten to twenty-five, 1.4 at fifty to one hundred. The water the solver draws is the water the transponders fill.[12]

[12]The basin carries ferries, feeders and coastal trades no crossing lane claims, which is why the right reading is the gradient onto the trunk, not a share of everything afloat.

6.2 The Coastal Response

The three columns of Table 3 are one design read three ways. As a switch the shock raises nitrogen dioxide by 3.5 percent per standard deviation of exposure, an interval of 2.4-4.5. In the dose column I match each city to its own gate, since a Mediterranean city sees what comes through Suez and a Red Sea city what comes through Bab el-Mandeb and the two counts fell by different amounts. The response is 4.7 percent per unit of the own gate's traffic gone. At the shortfalls the gates averaged over the post period, the dose and the switch are one reading. Split by side, the corridor that filled carries 0.026 and the corridor that emptied 0.043. Exposure is negative on the side the ships left, so a positive coefficient there is a fall in the column and not a rise. The two sides read as one, with an observed gap of 0.017 against the 0.037 the design could detect. The incidence account needs that symmetry, and the crossing variants price what each side's own coefficient would do instead.

The design's own contrast holds without pooling, drawn city by city in Figure 5. Each city's column is demeaned within its country-month and read against its own line fitted on the pre-closure months, and the cities the corridor reached hardest on each side are averaged. Because the two coasts sit at different latitudes and their residual cycles run in antiphase, each side's series loses its own pre-period seasonal means. In the lower panel the difference between the two fluctuates around zero for five years, steps up in the month the corridor moved, by what the switch estimates, and stays up to the end of the window. The identifying variation sits beside each reading in Table S12, separating the rows that are one design at a different resolution from the rows that are a different design.

The coastline was not already moving apart before the ships were. Entered by half-year around the diversion in Figure 6, the same exposure puts the two bins before it at 0.014 and 0.015, neither statistically distinguishable from zero and each a third of the first treated bin. Estimated instead of omitted, the bin holding the November attacks puts the whole of the anticipation the data show at 0.015, since the bin holds at most two attack-weeks diluted over six months. I keep the pandemic years out of this window, since they sit in the middle of the pre-period and moved shipping and urban nitrogen dioxide by amounts that differ across cities. On the same window the switch itself is 0.033 against the full sample's 0.035 in Table S13, so the exclusion buys a readable pre-period and costs almost nothing.

The response appears in the half-year of the suspension at 0.041, falls back to 0.026 in the half-year after it, and reaches 0.070 a year out, the dip sitting where the traffic sat, the same path the dose specification prices. Because the column holds no stock, the late rise is the traffic path deepening, not anything cumulating. The bins also do not decay, and the incidence account needs that, since the account is a flow per year and holds only while the diversion does. A year out the response is larger than at the suspension, and the corridor is still 45 percent empty at the end of the window, two and a half years on. The permanence the account assumes is the permanence the data show.

As written, the switch imposes that the traffic left in December 2023 and never came back. Yet it came back in part, and Table S14 puts that path in place of the switch. Entered beside the switch, the continuous path takes 0.076 and the switch takes -0.022. Because the corridor emptied

at the date the switch turns on, the pair does not decompose into a date effect and a traffic effect. The column establishes that the traffic path fits at least as well as the date, and it reconciles the dose with the switch, since 0.076 times 0.436 is 0.033 against the switch's 0.035, where 0.436 is the share of the pre-closure transit norm missing at Bab el-Mandeb averaged over the post period, the same computation that yields the gate's average fall, so the jointly estimated path implies the switch to within its error. Inside the post period, where the traffic is the only thing that moves, the response is 0.090, read as the joint response to the traffic and the canal authority's discounting since both moved over that window. The three readings climb as the identifying variation shrinks. They are one response at three resolutions.

The response falls off with distance from the coast in the shape a plume has. Each band's coefficient is scaled by the spread of exposure that band carries. Rescaled that way, the profile in Figure S3 runs from 0.041 within 25 kilometers to 0.007 past two hundred, and Table S15 reads 0.026 in the nearest band. The 200-400 kilometer band is the exception, at 0.067 with Egypt in and 0.025 with Egypt out, since greater Cairo and the delta move with something correlated with exposure over these months and no species that lives hours reaches that far inland. City size does not divide the response either, at 0.031 in the larger half of the cities and 0.031 in the smaller, in Table S16.

The port, Egypt, capacity arriving where the ships did, the calendar, the European trading scheme and the air itself all enter the estimate directly without moving it, and the placebo closures at the two clean Decembers find nothing. Appendix S1 runs each of these checks and names what each one rules out. It also carries the specification and the inference the estimate rests on.

6.3 Wind Direction Tests

The plume arrives on the wind, in proportion to the wind, and it did not arrive before there was a plume to carry. Estimated five ways in Table 4, Equation (10) with no wind in it returns 0.022 against the monthly panel's 0.035, and the gap is the sample, not the estimator, since the same overpasses averaged to city-months return 0.021 on the same cities. The pass panel is the 120 cities the corridor reaches most, and that subsample's own response sits at two percent, the same level the most demanding within-country designs read, so the wind contrast is estimated inside the subsample it describes.

The response on a downwind pass is 0.030 against 0.014 on an upwind pass, and the difference of 0.017 is statistically significant. The upwind reading has a physics rather than a confound behind it, since nitrogen dioxide lives hours and the satellite crosses once, in the early afternoon, so a morning of onshore flow leaves a loading the afternoon pass can still read after the sea breeze reverses. Scaled by how hard the air was moving from the lane, not merely by which way, the response rises 0.016 for each standard deviation of that wind. On the passes before the closure the downwind premium is 0.005 against the treated premium's 0.017, so no premium of the treated size was there before the closure. The pass with onshore flow gained its premium only when the ships arrived.

One reading of where the plume comes from is the nearest point of the corridor, and the kernel-weighted mean bearing over the whole corridor is the reading the exposure measure implies. On the second the split sharpens in Table S17, 0.036 on the passes the air carries against 0.009 on the rest. The same table dates the wind back, and matched to the wind of one or two days earlier the premium is gone, as a species that lives hours requires, since a plume yesterday's wind carried is no longer there to read. Because the source is a line instead of a point, alongshore flow arrives loaded. The falsifications the geometry does license, the dated wind and the pre-closure passes, both come back empty.

Rehearsed on the six days of the 2021 grounding and read on the same exposure gradient as the 2023 closure, the wind test gives 0.080 on all passes and 0.138 on the downwind ones, in Table S18. In contrast, the same six calendar days of the springs of 2019 and 2022 carry 0.012 and 0.013, with confidence intervals that include zero. The 2021 regressor is the share of a month the canal stood shut, so its coefficient is a full closure. Scaled by the share of the corridor actually gone after December 2023 it implies 0.035 against the 0.035 the later episode is estimated at, in Table S19. That is a ratio of 1.00.

The same plume is read at two dilutions, over the water and at the city, and Table S20 works out the arithmetic that connects them, bounding the estimate from each side. Starting from the counting box, where the ships are two fifths of the column, a whole corridor's plume is 2.8 percent of what the median city breathes. From the fitted enhancement over the open water instead, decayed to the median city's distance, the same corridor is worth 0.6 percent. The estimate, rescaled from one standard deviation of the reallocation to a whole corridor, is 1.6 percent, between the two anchors. That is where it should sit, since the box concentrates every lane at one point and overstates a coastal plume, while the open-water far field lets the plume dilute over sea room a coast does not have and understates it.

6.4 The Composition of the Plume

The four pollutants move the way ship exhaust moves them. The shock is read on all four at once in Table S21, each divided by its own within-city standard deviation, on the city-months where all four retrievals are usable. Nitrogen dioxide rises 0.099 and carbon monoxide 0.042, sulfur dioxide falls 0.070 and the absorbing aerosol index falls 0.024, drawn side by side in Figure S4. Resampled as one family in the stepdown manner of Romano and Wolf (2005), the four survive together, in Table S22.

Carbon monoxide is the long-lived comparator the distance profile needs, since a species that survives for weeks is mixed well enough that a country by month effect takes most of it. The rise in nitrogen dioxide and the fall in the aerosol index are the two facts that discriminate, since biomass smoke and Saharan dust both raise that index and both bracket these corridors, so an index that falls where nitrogen dioxide rises is not smoke and is not dust.

The sulfur reading is a differential and not a level, since, split by corridor in Table S23, sulfur dioxide falls 0.178 where the ships arrived and moves 0.029 where they left, with a confidence interval that includes zero. In the Mediterranean, the sulfur rule that binds from May 2025 leaves

the lost-corridor reading where it was. The level's own sign carries a censoring component, since a scrubbed fleet cannot subtract from a baseline it never raised and the retrieval loses exactly the months the column falls through the detection floor, so the fuel chemistry is identified by the between-corridor contrast and its timing, not by the fall alone.[13]

Nitrogen dioxide lasts hours in the boundary layer and carbon monoxide lasts weeks, so a fresh plume has to be steep in distance in the first and flat in the second, while persistent spatial structure would look the same in both. In Table S25 nitrogen dioxide falls from 0.041 within fifty kilometers of the coast to 0.008 in the band from two to four hundred, a fifth of what it was, while carbon monoxide holds three fifths of its coastal value out to the same distance. The two profiles separate exactly as the two lifetimes imply.

Two gases separate exhaust from burning on the ground. Formaldehyde is a gas that burning vegetation makes and ships barely do, and under each city's own calendar-month effects the corridor holds it flat at 0.003 while nitrogen dioxide carries 0.026, in Table S26. The null bounds formaldehyde below 0.014, half the nitrogen dioxide response. Ozone is read on the total column, nine tenths of which sits in the stratosphere, and with twenty-two degrees of latitude between the two sides that column carries season beside titration by fresh nitric oxide. The sign runs in the direction the chemistry implies. The argument rests on the signature itself, since a signal rising in nitrogen dioxide and leaving formaldehyde alone is exhaust in the air instead of burning on the ground.

Something could have moved every retrieval at once, since three of the four products fall on the gained side. In Table S27 the corridor shows nothing on the retrieval's own cloud fraction and moves the surface albedo by two thousandths of a unit, well below the albedo sensitivity the air-mass-factor error budget already carries (Lorente et al., 2017). Entering the absorbing aerosol index as a control, a bound rather than a preferred reading since the index is itself an outcome, moves the estimate from 0.035 to 0.037 in Table S28. The species separate by distance the way their lifetimes predict, and the response is present on the downwind passes and half that on the upwind ones at the same cities in the same months. That joint pattern is hard to reconcile with a fixed feature of the coastline or with an artifact of the retrieval.

7 Distinguishing the Plume from Port Activity

The usual proxy for onshore activity is night radiance, and radiance cannot say whether anything besides exhaust moved with the diverted traffic. The sensor sees a lit vessel at anchor, and lit vessels arrived in whichever corridor gained traffic, so the radiance signal would repeat the traffic shock instead of measuring anything the ships brought ashore. In contrast, a port call is counted

[13]The fitted fleet is the largest hulls, and the largest hulls rounded the Cape. Scrubber-fitted hulls carried 36.4 percent of global container capacity at the end of January 2024, up from 32.4 a year earlier, on a far smaller share of vessels, on Alphaliner's own fleet count (Alphaliner, 2024). Sulfur is also the one species the retrieval loses, and losing it is not random. A missing sulfur month becomes 1.3 percentage points more likely where the fleet arrived, in Table S24, the pattern a genuine fall through the detection floor produces, and it censors the fall toward zero, not away from it.

only when a ship enters a harbor, so the series cannot be mistaken for the traffic that merely passes offshore.

In the log form Table S29 finds calls at the corridor ports rise 2.9 percent for a standard deviation of exposure and tanker calls the same 2.9. To keep the pandemic out of the comparison, I drop 2020 and 2021 from the window. Coastal activity over those two years measures the pandemic more than it measures shipping, and the pandemic did not hit the corridor cities in the same proportion as the rest. Varying the fixed effects runs the estimate between 0.012 and 0.030 across the four specifications of Table S30. Placebos that move the shock to December 2019 and December 2022 find nothing at either date on either outcome, in Table S31.

Because a logged count with its zeros patched is exactly the object recent econometric work warns against (Chen and Roth, 2024), Table S32 estimates calls, cargo, and tanker calls by Poisson, keeping the zeros and needing no transformation, and I rest the claims of this section on that reading. Pooled calls come in at 0.102, which exponentiates to 10.7 percent, the one conversion used everywhere, and the cargo those calls brought comes in at 0.018, where the logged cargo shows nothing at all. Tanker calls enter at 0.075 under the same estimator, and the administrative fuel record below gives an independent check where one exists.[14] The log form was losing the months a port recorded no tonnage, exactly the months the arrival of a new route should change.

Because fuel is sold and cargo handled where a ship stops, a real arrival effect also sits on the cities that own a port and is absent from the cities ships only pass. Splitting both outcomes that way in Table S33, calls rise 0.057 where ships stop against zero where they pass, and the equality of the two is firmly rejected. Tanker calls run 0.022 against -0.002, so the split is decisive for the calls and points the same way for the fuel. The columns move on both kinds of city alike, the calls move only where a port sits, and that contrast separates the exhaust from the business.

One port on the gained coast ran the reverse experiment on its own. Ship-to-ship bunkering in Algoa Bay stopped in September 2023, when the revenue service detained the bunker tankers over a customs assessment, a policy shock with no connection to the diversion and the opposite sign to it. If the column measured the harbor economy, Algoa Bay's column should have fallen with its business. Instead its tanker calls fell by seven tenths and its column rose 0.067 against 0.030 at the other South African coastal cities, on an exposure of one and a half standard deviations, in Figure S5. The harbor's business collapsed while the diverted corridor kept running offshore, and the column still rose with the corridor. The port channel and the plume channel move in opposite directions here, at one port and under a shock of its own.

Tanker calls are the sharper of the two outcomes because a tanker call is bunkering, and bunkering is the one service on this coast that has to move when the route does. A ship buys fuel along the track it sails, so a rerouted track moves the purchases with it. Across the ten ports the detour runs past, calls rise by a quarter or more at the four bunkering ports, most at Cape Town at 45.3 percent, in Figure S6. Singapore and Houston, which sit on neither corridor and serve trades the closure never touched, move 4.3 and 7.9 percent, so the rise belongs to the detour and not to a

[14] South Africa publishes no comparable monthly bunker-sales series, so the gained coast has no administrative mirror.

freight cycle. The mirror is at Fujairah, the bunkering port that fueled the traffic the Red Sea lost, where calls fall 2.3 percent. The diversion shifts bunker purchases toward the coasts the ships now run past.

Singapore is the one port on this map that reports what it sold, publishing monthly bunker sales in tonnes, a government series with no connection to the transit counts or to anything in orbit. Sales run 4.11 million tonnes a month over the five years to November 2023 and 4.68 million since, a rise of 13.9 percent, in Figure S7. The two Singapore readings agree, the calls nearly still because its own trades were untouched, the sales up because the voyages that still pass it have grown longer and take on more fuel. Fuel arrived where cargo did not, and at the 26 cities with no port of their own no cargo could have come ashore with the calls. Across the corridor as a whole the logged cargo comes in at -0.033 in Table S34, against the 0.018 the count form gives.

Sorting cities by proximity to a harbor rather than by title to one leaves the coefficient where it stands. In Table 5 a city counts as call-at when any port sits within 40 kilometers of it. Cities the ships merely pass show 0.036 and cities they call at show 0.034 against the pooled 0.035, and the data cannot tell the two apart. The same equality holds with the radius drawn at 20 or at 10 kilometers instead. The equality shows the kernel measures the exhaust and not the business in the harbor, so the weight enters equation (5) as written.

The side that lost the corridor anchors the comparison, and on the side that gained it the same split gives 0.053 without a port and 0.034 with one, the no-port cities responding more and not less, which points away from the harbor. The Algoa Bay reversal already ran the claim at a single port (Ship & Bunker, 2025). The column over the ocean beside that coast rises most where no harbor sits to explain it. At the level of the individual overpass, with thirty-fold the observations, the call increment on the response is -0.004, an equality an order tighter than the monthly test could state. The filled coast reads the same way as the emptied one.

8 Incidence

I count exposure in person log points, the flow the estimate implies while the diversion holds. The chain from a transit to a dollar runs transit to column through the kernel, column to log through the outcome, and log to person log points through the population, and person log points become money through whatever value a reader gives θ, a price this paper never fixes. Figure 7 draws the account on the coastline itself, before any table reduces it to totals. The released side is long and dense. The side that took the plume is short and sparse. Banjul, for scale, carries 1.3 million people at an exposure near two standard deviations, and contributes 85 thousand person log points a year to the taken side. I hold one account canonical throughout, the estimation-sample coasts under the exclusive population assignment at the pooled coefficient, and every other reading in this section is a named departure from it. Table S35 puts the whole family on one page. Each row carries its population, its net and its ratio, with the departure that produced it named alongside.

The one coefficient carried across the coastline and weighted by who lives there adds the population without adding identification, in Table 6. On the sample the coefficient comes from, the

reallocation takes exposure off 80.9 million people and puts it on 36.7 million. With the ten coasts at war added, where the coefficient is carried in rather than estimated, the count rises to 119.2 million released against the same 36.7. Because the account scales linearly in the coefficient, a person-weighted coefficient would leave it where it stands, the grid's population-weighted cells running 0.033-0.038 around the unweighted 0.035. Of the 36.7 million the exposure lands on, 5.4 million live in 26 cities with no port of their own. The people the plume arrived on captured none of the trade, the mismatch the pass-call split identifies and Table 6 measures.

The Gulf coastline sits in the account's Neither row with exposure at zero, its final-months columns pick up the Hormuz conflict of 2026 instead of anything this closure did, and the zero weight keeps that conflict out of every total. Yet the Neither row still carries content, since the diversion absorbed fleet capacity, carriers canceled sailings on lanes whose distance never moved, and the column fell beside those lanes too, an improvement I leave uncounted. Priced at the control corridor's own coefficient, that withdrawal removed a further 1.2 million person log points across 58 million people, a row the account carries outside the totals, and its sign runs in the diversion's favor.

I extend the account to the 100-200 kilometer band at that band's own coefficient, 0.035 estimated on a sample that keeps Egypt, and the asymmetry widens, in Table S36. The extension adds more to the released side than to the taken one, since Cairo and the Nile delta enter at that distance on the corridor that emptied, and the person-weighted ratio of released to taken rises from 3.6 to 5.1. Priced at the coefficient estimated with Egypt out of the band sample the ratio is 4.8, and taking every Egyptian city out of the band entirely leaves it at 3.8, above where it stood inside 100 kilometers. Thus the delta amplifies the band's widening and does not create it. The ten largest cities on each side are named with their populations, exposures and incomes in Table S37. Plotting each city's realized change in the column on the change the account's arithmetic predicts for it, Figure S8 puts the cloud on the diagonal at a slope of 1.10 on the estimation sample alone. On the war coasts, where the coefficient is carried in and not estimated, the slope is 0.30, so those coasts are reported apart from the sample.

Age is the damage heterogeneity that matters most, since mortality risk per unit of the pollutant climbs steeply past sixty-five, and this margin works in the diversion's favor. In Table S38 the coastline released is the older one, 8.2 percent of its people over sixty-five against the taken side's 5.3, so the diversion took the plume off 9.7 million people past that age and put it on 1.9 million. Weighting each person's exposure by the over-65 share widens the asymmetry to 4.4 times. Age and income therefore pull against each other, and I report the income weighting as the headline because it is the one the distributional literature asks for and the less favorable of the two. Weighted by both at once, each person carrying their country's over-65 share and the inverse of their relative income, the arc ratio stands at 1.62 and the account still removes more than it delivers, so the two margins do not cancel the ordering even jointly.

The city panel stops at 65 degrees east. Because the solver runs to 130 and the account needs no satellite data, Table S39 carries the same kernel, population and coefficient to every eastern urban center either state's corridor reaches. The eastern coast releases more than it takes, 179.2

million people against 94.6, because the Indian Ocean leg toward Bab el-Mandeb empties while only part of the Indonesian water fills. Adding it removes a further 1.6 million person log points a year, so the account inside the window is the conservative one.

The account's size depends on how the coastline is counted, and Table S35 carries every counting. The box account on the cities with an income figure gives 2.8 million person log points a year net, the same account over every city with a population count gives 3.6, and retiring the box rule altogether gives 11.7. Yet the argument turns on the person-weighted ratio between the two sides, and the ratio runs 3.0, 3.6 and 4.9 across the three, widening as the account reaches further from the city centers. The unthinned account of Table S40, rebuilt over every urban center separately with no thinning at all, takes exposure off 302 million people and puts it on 137 million, and it does not belong in this progression, because its boxes overlap, so the same air is counted toward more than one city and both sides inflate together. Its level moves further still, and the pricing section uses the ratio.

The account survives without any city definition at all. With the box rule retired and the gridded population mosaic (Bondarenko et al., 2020) summed onto the solver's own grid in Table S41, the diversion releases 14.8 million person log points across 422.3 million people and takes on 3.0 million across 89.2, an asymmetry wider than the city account's, because the rural coast the city rule leaves out sits mostly along the corridor that emptied. The calendar is a further choice, since the account counts people on a 2020 grid over a window that runs to 2026 and the two coastlines grow at different rates. Every account and ratio so far is built on the WorldPop aggregated grid, and reading the constrained mosaic at 2020 and at 2025 under the same exclusive assignment, Table S42 finds the person-weighted ratio unmoved while the net removed grows with the population, so five years of differential growth add to the net removed instead of eroding it.

Whichever coast took the exhaust also took the calls and the fuel sales, and whichever coast lost it lost them, so the harm and the commerce arrived as one bundle and I price only the first. Beside the harm the commerce is not large. In the count form that keeps the empty months, the calls outran the cargo they carried five to one. What the coastline received was passage and the fuel it bought, the object an exposure instrument prices, far ahead of the trade that came with it.

Incidence has two halves. The first is how much exposure each coastline takes, which Table 6 settles. The second is what a unit of it is worth to the people who take it, the half where income enters. The question is the one the environmental-justice literature asks of instruments on land, where a market can reallocate pollution across places that differ in income (Hernández-Cortés and Meng, 2023) and where the exposure gaps are read from the same satellites (Hsiang, Oliva and Walker, 2019; Colmer, Hardman, Shimshack and Voorheis, 2020; Currie, Voorheis and Walker, 2023). Here nobody moved and no permit traded. Yet the reallocation between incomes arrived anyway, on a carrier's routing.

The corridor that emptied ran past an unusually affluent coastline, at 11,748 dollars a head against the 7,689 of the coasts the shock never touched, while the corridor that filled ran past one at 4,725, poorer than either, in Table 7, incomes being GDP per capita in current dollars for 2022 (World Bank, 2024). Read on the African cities alone the side that took the plume earns four fifths

of what the side that released it earns, so the gradient survives the comparison that puts countries of a kind against each other. It survives with the Gulf and Israel removed and with the largest cities on either side taken out, in Table S43. Both sides span the income range, so the gradient is not one continent against another. The side that gained the corridor holds Spain along with Sierra Leone and Liberia, and the side that lost it holds Egypt and Syria along with Saudi Arabia and Israel.

I apply one coefficient to every city and then weight by income, a construction that compounds the response and the welfare weight if the coefficient itself moves with income, and Table S44 splits the response at the median. The richer half carries 0.039 and the poorer half 0.021, a gap the data cannot separate from zero, so the weighting varies the welfare weight alone. Letting each city carry its own half's coefficient moves the crossing, the point at which an income weighting turns the account's verdict, because the plume arrives where the estimated response is the weaker one. The income figure is a national one, the 2022 vintage fixed before treatment so the closure cannot move its own welfare weights. A downscaled grid that carries income at each city's own administrative cell (Kummu, Kosonen and Masoumzadeh Sayyar, 2025) reads the crossing at 1.11, below the 1.38 the fixed vintage gives, while holding the taken side's income near two fifths of the released side's. Neither the gradient nor the crossing is an artifact of every city carrying its country's mean.

The same account supports two readings, one at equal weights and one under a weighting that favors the poor. Across the coastline the corridor released more exposure than it delivered. At equal weights the diversion is an improvement. Yet the people it arrived on earn two fifths of what the people it released earn, and four times as large a share of them live below the median income of this coastline. Against an absolute line of 2,000 dollars a head the shares are 41 percent and 13 percent. Under a weighting that favors the poor strongly enough it is a loss. The corridor comes off Egypt, Yemen, Algeria and Saudi Arabia and lands on South Africa, Senegal, Morocco, The Gambia and Sierra Leone, and the income gap between those two lists is what the distributional reading turns on. Table 8 locates the point where the two readings balance.

The unit the account is kept in encodes a judgment about damages. In log points a swing on a clean coast weighs as much as the same proportional swing on a dirty one, the reading a proportional damage function implies, while in level units, person micromoles instead of person log points, the dirty coast weighs more, the reading an absolute one implies. In levels the diversion removes more than it delivers by an even wider margin, and the verdict does not turn until an inequality aversion of 2.34, both readings in Table 8. As a result, the two metrics bracket the crossing between 1.38 and 2.34, read at named weights in Table S45, and the improvement at equal weights holds under either.

Chained onto the measured pre-closure surface concentration through the meta-analytic all-cause risk for nitrogen dioxide (Huangfu and Atkinson, 2020) at each country's own death rate, the account's per-city change becomes the quantity the mortality literature prices, in Table S46. Applied mechanically, that function implies 107 deaths a year added on the coast the plume arrived on and 394 avoided on the coast it left, a net of 287 avoided, with the risk function's own interval spanning 140-570. Reading the column at the tightest of the measured station slopes would scale

both coasts down together without moving their ratio. Age-standardizing the crude death rates would raise the avoided count relative to the added one, the coast the ships left being the older one. The translation's two unpriced choices run in opposite directions and neither changes the sign.

The crossing point and the exposure-charge schedule come out of the same calculation. I weight a person's exposure by their income relative to the mean raised to minus η, the aversion parameter of Atkinson (1970) read as a social marginal welfare weight (Saez and Stantcheva, 2016). The income-weighted exposure charge on the two arcs, read by a global utilitarian planner rather than by any toll-setting agent, then becomes a schedule in η, drawn in the left panel of Figure 8.[15] Because the diversion moves exposure from one arc to the other, its welfare effect is proportional to the difference of the two weighted loads, and it is zero exactly where the two arcs price equally. The 1.38 at which the log-point verdict flips and the 1.38 at which the two income-weighted exposure charges equalize are the same number because they are the same equation, and the level metric moves both together to 2.34.

Valuation enters the same equation as one more exponent. If a unit of exposure is valued in money and the valuation rises with income with elasticity η_v, the weighted load of arc a is

$$L_a(\eta) \;=\; \sum_i N_i \, \omega_{ia} \left(\frac{y_i}{\bar{y}} \right)^{\eta_v - \eta} \tag{11}$$

so every crossing in this section depends on the pair only through $\eta - \eta_v$. A reader who prices damage in money at the revealed-preference elasticity near one (Ito and Zhang, 2020) reads the whole schedule shifted right by one. The crossing sits near 2.4, outside the 0.5-2 range conventional aversion spans, so under the cross-country valuation convention the diversion is an improvement at any conventional aversion. Read the other way, reversing the verdict requires both the physical-unit convention and an aversion near the top of that range.

The physical unit is the baseline on substantive grounds. Within a jurisdiction regulatory practice values a unit of physical risk identically across the people it covers. In contrast, across countries the benefit-transfer convention income-scales the value near unit elasticity (Hammitt and Robinson, 2011; Viscusi and Masterman, 2017), exactly the money-metric reading at η_v near one. Thus the uniform baseline is the domestic standard and the environmental-justice position, the money-metric variant is the cross-country convention, and I print both and elect neither. At equal weights the ratio is 2.59 on the full coastline and 2.26 on the cities with an income figure, which is where the curve starts. An aversion of one half prices the arcs at 1.80 and log utility prices them at 1.31, so the diversion is near a wash at log utility, at 1.38 the two arcs price equally, and above that the Cape arc is the dearer one and the diversion is a loss.

No single number produces the crossing on its own, and the reading kept as the headline is not the one most favorable to the diversion. Read under each side's own coefficient in Table S47, the crossing is 2.11, the more favorable to the diversion because the side that took the plume is the poorer one and takes the smaller coefficient. Letting the coefficient follow the city's own income

[15] The reference income the weights are read against is the coastline's own mean, a choice with no content for any crossing, since rescaling every income by one number multiplies both weighted loads together and moves no equality.

puts it at 1.84. I keep the pooled reading and the lowest crossing of the three. Because the whole account scales linearly in the coefficient, it can be priced at the most demanding reading the paper produces. At the 0.02 of the within-country designs the net removed is 2.1 million person log points, the attributed net 1.3, the indifference value 2,245 dollars and the mortality translation 164 net avoided deaths. Each is the baseline scaled by four sevenths or its inverse, and the crossings do not move.

Converting incomes at market rates reads the account as money, the unit a levy collects, while converting at purchasing power reads it as consumption, the unit a welfare weight attaches to, so the baseline and the parity variant answer two different questions. Winsorizing incomes at the fifth and ninety-fifth percentiles puts the crossing at 1.38, purchasing power parity at 1.20, and removing the single largest weighted contributor puts it at 1.67, all in Table S47.

The heaviest of those contributors is Banjul, the urban-center agglomeration of the settlement layer, which is wider than the municipality, which earns 836 dollars a head, so the one place to inspect when the crossing moves is Greater Banjul.[16] Dropping each of the ten heaviest weighted cities in turn moves the crossing between 1.41 and 1.67, always upward. Capping the weight at its ninety-fifth or ninety-ninth percentile leaves it unchanged, so the headline crossing is the low one. The one income that moved during the window is Egypt's, and repriced at its 2024 figure the crossing is 1.58, in Table S47, because the released coastline is poorer at those prices and the improvement correspondingly larger.

The crossing is therefore a reading on the coastline the income data covers, and Table S48 prices the coastline that data leaves out. Poor coasts on the corridor that emptied are the ones most likely to lack a 2022 income figure, Yemen's nineteen cities and nine million people first among them, so completing them moves the crossing in one direction only. With the war coasts out of the account entirely the crossing sits at 1.02, in Table S48, and with Egypt out of the account it sits at 0.60, in Table S49, the largest single instance of the same mechanism. I treat that row as a diagnostic, since an account of who breathes the corridor cannot delete 33 million people who do. With the dropped cities returned at half the last income their country ever reported or at that income itself, the weighted released side grows so poor that the net never turns. At two and four times the last report the crossing returns at 3.25 and 1.76.

The candidate readings run from 1.02 to 3.25 against the standard range of aversion shaded in Figure S9, so a reversal inside that range needs an aversion above one, and the only reading near the band's floor is the diagnostic that deletes Egypt's 33 million from the account. Every crossing variant in this section moves through this one mechanism, the income assigned to the released coastline the data does not price. Yet the improvement at equal weights, the income gradient of 0.37-0.43 and the efficient-ratio ordering stand in every row.

The weighted sums are concentrated, with half of the taken side in five cities at an aversion of one and a half, in Table S50. Over the fitted decay length's own sampling distribution the ratio stays between 2.2 and 2.6 against the baseline 2.59 and the crossing between 1.27 and 1.41

[16]The 1.3 million is consistent with the 2024 census, which counts 2.42 million nationally, with Brikama and Kanifing, the two areas the coastal belt runs through, holding 63 percent of it (Gambia Bureau of Statistics, 2024).

against 1.38, in Table S51, which in money keeps the indifference value between 1,000 and 1,900 dollars. The coastline is enumerated, not sampled, in the design-based sense of Abadie, Athey, Imbens and Wooldridge (2023), the same view of uncertainty the paper's permutation inference takes. Because the response cancels out of the ratio, the decay length is the only estimate the two headline ratios inherit within one construction, while Table S52 collects what the construction itself moves. Enumeration removes sampling error but leaves measurement error, and the measurement error here has a known sign. Undercounted informal settlements sit on the released side (Wardrop et al., 2018), so a corrected count would add people to the side the diversion relieved and move both the ratio and the crossing in the direction the paper already reports as conservative.

9 The Arc Price

9.1 The Efficient Ratio against the Instruments in Force

Every instrument that scales with fuel reverses the order, since a levy on fuel is a levy on distance, so the price it puts on an arc is proportional to the length of that arc. The Cape arc is longer and it reaches fewer people, so the emissions trading system, the fuel standard and the framework the organization has approved all charge more for the arc that delivers less exposure.

The total efficient charge on a transit is the value of exposure times what the transit delivers plus the carbon price times what it burns. The ratio of those total charges across the two arcs is therefore a curve in what exposure is worth. At a value of zero the curve starts at 0.69 in the right panel of Figure 8, drawn alone at full size in Figure S10, the fuel ratio and the relative price every fuel-scaled instrument charges at any value, the Cape dearer because it is longer. As the value of exposure rises the curve climbs through one at 1,279 dollars for a person-log-point-year, and approaches 2.59, the ratio of the two exposure charges, as exposure dominates the ledger.

The crossing falls at the same 1,279 the carbon table reports, and the match is a consistency check instead of a construction, since I build the curve from the arc price account and the table from the counted fuel chain, and the two meet because both price the same two arcs. A dollar figure here means whatever value a reader gives the exposure, and this paper never fixes one. Every fuel-scaled instrument in force or proposed sits on the horizontal line, and the mapped control areas change what burns inside a box without moving a ship between arcs. The gap between that line and the curve is what this paper measures.

The reversal is not an argument against pricing carbon, since the longer voyage does burn more fuel and the extra fuel is the reason a carbon instrument charges it more, in Table 9. Two externalities and one instrument is a Tinbergen problem (Tinbergen, 1952), and equation (6) solves it, a uniform charge on fuel for the carbon and an arc-differentiated charge on exposure for the plume. The rerouting onto the Cape adds 24.2 million tonnes of carbon dioxide a year. An independent accounting of the same detour reaches the same conclusion about the sign, that a

charge on carbon prices the diversion in a direction its designers did not intend (Peng, Wang and An, 2024).[17]

Of the 24.2 million tonnes, 15.8 are attributable to this closure, with at most a transit a day of the remainder pushed the same way by Panama and none of it fleet growth, since the passages the diversion never touched shrank over the same window, in Table S55. The hull classes that carry the corridor's exhaust match one for one across the two counts in Table S56, tankers within a transit a day, which rules out a sanctioned-oil fleet running the Cape. The remainder is dry bulk, ten transits a day that the table rules out one at a time as diverted traffic, as Black Sea war traffic and as fleet growth. The attribution rule the carbon account uses follows from those counts, charging the closure the share of the Cape's rise that Bab el-Mandeb lost, so 24.2 becomes 15.8 by one scalar, computed once and inherited by every attributed row.

The exposure side takes the same scalar and would carry a larger number without it. The diversion removes 3.6 million person log points of exposure a year, of which the scalar makes the same 65 percent attributable, 2.4 million. Yet nothing but the closure emptied the strait, so the removal needs no discount of its own and the 5.0 million removed on the corridor the ships left is attributable in full. Read that way the attributable removal is 4.1 million and the indifference value below falls to 734 dollars. I carry the uniform scalar in the text because it is the more demanding of the two.

A fuel levy reaches the carbon and leaves the exposure unpriced, while a mapped control area does the reverse. The last rows of Table 9 put the two on one line. At the central regulatory price of carbon, 190 dollars a tonne at the 2 percent discount rate the schedule is published for (U.S. Environmental Protection Agency, 2023a), the attributable emissions are worth 3.0 billion dollars a year against the 2.4 million attributable person log points. The diversion is therefore welfare neutral when a log point of nitrogen dioxide for one person for one year is worth 1,279 dollars. That is 45 dollars a person a year at the 3.5 percent a standard deviation of exposure delivers, and it sits on an exhaust stream whose particulate matter carried tens of thousands of deaths a year worldwide before the sulfur cap (Corbett et al., 2007). External per-ton damage schedules do not transfer into this bracket without relabeling a density difference as a valuation, since each is calibrated to its own country's receptor population, so the anchors here are the paper's own chains.

The underwriting market supplies the one outside number worth holding beside these. A hull that kept the canal paid between one and one and a half million dollars a transit for the listed water, while the exposure a diverted transit took off the coastline, 247 person log points at the account's attribution, prices at 316,000 dollars at the indifference value. Thus the war charge the

[17] The account is computed at a common service speed, and the speed margin runs one way. Diverted ships sped up to hold their schedules, the largest container classes by 5 percent in the first quarter (UNCTAD, 2024), and fuel per mile rises with the square of speed. The carbon the diversion adds therefore sits above the printed 24.2, the per-tonne ratio above its printed 3.73, and the per-transit ratio repriced with the observed speed-up reads 2.35 in Table S53. The speed margin stretches the band at its dear end and moves no reading toward one, and the convoy speed the canal mandates works the same margin from the other side, in Table S54.

market actually levied ran three to five times the efficient exposure charge, and the diversion it produced is the one this paper prices.[18]

The mortality translation bounds whether plausible valuations sit above or below that point, pricing a person log point at eight deaths in a hundred thousand. At an American regulatory value of a life of 11 million dollars, the level of the federal guidance in current dollars (U.S. Environmental Protection Agency, 2024) read under the income-invariant convention, that rate prices a person log point at 880 dollars, and the exposure differential on a representative transit then stands at two thirds of its carbon differential. Income-scaling the valuation at the unit elasticity the benefit-transfer literature uses across countries (Hammitt and Robinson, 2011; Viscusi and Masterman, 2017) pushes the same ratio toward a tenth. Letting the nitrate particulate the exhaust forms downwind into the damage, the pathway the transport-chemistry accounting prices ship mortality through (Sofiev et al., 2018), pushes it past one by an order of magnitude the tracer chain alone cannot see. The range therefore straddles one, so under the convention regulators use the ordering error of the instruments in force is first-order in level as well as in geography. At the range's bottom the geography correction still stands whole, since a fuel base puts a weight of zero on exposure at any level. The chain delivers an indifference point. It does not fix what exposure is worth, so a reader holding any value can see which side of that point it falls on.

The efficient exposure charge stands at 2.6 on the Red Sea arc against the Cape arc per transit of a given hull, and at 3.7 per tonne of fuel, in Table 10. The table carries the person-weight one crossing delivers on each arc in its own rows, so a reader willing to defend a value of exposure obtains the level by multiplication. Yet the instruments in force charge it 0.7 times. A levy on fuel prices the arc that reaches 133.6 million people below the arc that reaches 96.0 million, for the mechanical reason that the second corridor is 4,133 nautical miles longer at the crossing lanes' traffic weights. Because a chokepoint pairs the shorter way with the more confined water, any instrument that prices distance underprices chokepoints relative to the routes that avoid them, and the error grows with how much distance the chokepoint saves.

The ratio survives each of these one-at-a-time constructions. Netting the water both routings cross, pricing each lane at its closest point, fitting each corridor's own length, reweighting by the hours the wind blows ashore and rebuilding the exposure at every standoff the calibration allows all leave the Red Sea arc dearer by a wide margin. The family and the margins that run each way are carried in Appendix S3.

9.2 The Ratio at Every Passage

The rule does not stop at this pair, because the efficient exposure charge prices an arc by the exposure a transit on it delivers, and the population and geometry that exposure is made of exist at every passage whether or not a war closed it, so I carry the same kernel around the world in Table 11. The exposure there is a different object from the one the arc tables price. Along an

[18]The 190 is the schedule's value for emissions in 2020 at the 2 percent discount rate. The indifference value scales one for one with it, 337 dollars at a 50 dollar carbon price, 1,279 at 190, 2,019 at 300, with the 2030 vintage raising the 190 reading by a tenth.

arc, exposure is summed over every city the corridor passes, the quantity a charge on using the arc should internalize. At a gate, it is the population beside the passage itself, the quantity a charge collected at that gate can reach.

The Red Sea arc prices at 2.6 times the Cape arc as an arc, while at its two gates the same corridor splits into a Suez at 40 and a Bab el-Mandeb at 8 on the index of Table 11, mapped in Figure S11, and the two measures bracket each other. Against the Cape's own gate at 4.5, Suez prices at 8.9 times and Bab el-Mandeb at 1.7, and the arc-level ratio sits between them, so on this corridor the integral is anchored by the gates instead of by the open water between them. The gate measure depends on how far a plume is taken to travel, and Table S57 moves that scale from a third of the one used to more than three times it. Across that whole range the ordering the rule produces correlates with the one at the scale used at 0.887 or above, so the ranking follows from where people live rather than from where the decay is cut.

Each passage is priced on its own, and the passages substitute, so the index ranks what each passage delivers. It is not a schedule of charges that could be imposed together without recomputing every one. A charge at Sunda that pushed a transit to Lombok would trade a passage priced at 100 for one priced at 17, still removing five sixths of the exposure it targets, while pushing it all the way to Malacca trades 100 for 38. Substitution thins what a single charge buys and reorders nothing, because the three passages sit a factor of two to six apart while the detours between them are a day's sailing.

Ranked by the gate measure, the passages do not sort by traffic, and the difference is the point of the measure. The rankings the literature carries score a chokepoint by the trade that crosses it (Verschuur, Lumma and Hall, 2025), which answers what the world would lose if the passage closed. This table answers what the passage costs the people beside it while it stays open. The Sunda Strait carries an eighth of the transits of the Korea Strait beside it and the rule prices it at four times as much, since Jakarta sits on the passage and nothing of that size sits on the other. Confined water makes these rankings conservative for the passages already ranked first, since the corridor pair showed a confined sea concentrating the same exhaust eleven times over open water, and the index prices geometry alone with no credit for the trapping.

A throughput ranking puts the great funnels first wherever they run, while an exposure ranking puts the inhabited ones first. Among the ten busiest passages the price varies by a factor of nineteen. The twenty-one that carry twenty transits a day or more span a factor of thirty-six once Kerch is set aside under the same named war rule that keeps the ten conflict coasts out of the estimation sample. That exclusion shrinks the statistic instead of inflating it, since with Kerch included the factor is two hundred. Finally, the two passages that actually charge a toll sit far apart on the index, Suez in the top three and Panama at a quarter of its price.

9.3 Behavioral Responses to a Toll

The canal authority prices its own arc and nobody prices the Cape, so a toll raised toward the damage the canal does moves traffic onto the arc that carries no charge at all, and the exposure goes with it. Equation (2) determines which lanes move at which toll and the solver returns the

exposure that results, the whole schedule of Figure 9 and Table S58, against the toll in the units the carrier's comparison is made in, nautical miles of detour. Little happens through a toll worth 3,000 nautical miles, where six lanes and three percent of the traffic have left. Two thirds of the traffic leaves inside the next thousand miles, and the exposure the coastline takes falls to under half of what it was, because the arc the traffic leaves is the one that reaches more people. A lane leaves whole when the toll passes its own detour, so the schedule is a staircase, not a curve, and the grid is drawn at fifty miles through the range where the steps fall.

The schedule holds within-lane volumes fixed, while the paper's own elasticity implies the volumes shrink with delivered cost, and that margin can be priced. Converting the trade-value elasticity into transit thinning at a constant value per transit, a toll deep enough to empty the canal thins the remaining volumes by six percent at the Red Sea's own trade elasticity. That elasticity is what a published restriction at a second chokepoint identifies, and Appendix S2 estimates it. The thinning is an order below the routing steps it would sit beside and proportionally smaller at the weaker readings of that elasticity, so the demand margin bends no feature of the figure. Within-lane hull heterogeneity smooths the staircase the same way, since a lane's hulls leave over a band of tolls instead of at a single one. That rounds the steps and shifts the revenue peak without moving which side of it the exposure floor sits on.

The income-weighted exposure index falls less, stops falling sooner, and then turns. The unweighted index bottoms at 37 and needs a toll worth 6,000 nautical miles to get there, while the income-weighted one bottoms at 62 by 5,250 and rises to 67 by the far end, marked on Figure 9. Rebuilding the whole schedule without Cape Town, the largest single exposure on the gained side at 2.65 standard deviations, moves the floors to 34 and 59. A toll set past the weighted floor makes income-weighted incidence worse, because what is left to move is traffic whose detour runs past the poorest coast on the map, and the income of the people bearing the exposure falls from 11,901 dollars to 9,922 along the way.

Neither floor is the toll a treasury would pick. Revenue peaks at a toll worth 3,300 miles, immediately before the first heavy lane leaves, while the exposure floor needs 6,000. So the revenue-maximizing authority and the exposure-minimizing planner stand on opposite sides of the same step, Proposition 1 in numbers. The proposition's termination margin exists in data, in the coast-switching a customs ledger records. Trade enters that ledger twice, once as canal crossings and once as the lanes' own trade. The volume that still moves through the canal falls to a hundredth of where it started, while the trade of the same lanes holds at 96 percent, since a lane that leaves pays its own detour instead of the toll.

At the far end of the schedule the fall in exposure is bought at 25.9 million tonnes of carbon dioxide a year against the 24.2 the closure itself moved, the limit of what a single gate can reach and the reason the charge belongs on the arc itself. The schedule holds the network fixed, so it prices the traffic a toll moves and leaves aside the lanes a carrier would redraw around it.

The theory picks out one computable point on the schedule, the total efficient extra charge on a canal transit, the carbon price times the difference in what the transit burns plus the value of exposure times the difference in what it delivers. I evaluate it in Table S59 against a toll the

authority already posts, 552,000 dollars for a container transit on the canal's own published dues and tonnage in the year the closure began. The correction is computed per fixed hull, the base the paper's fuel and exposure chains measure, because the observed mix is itself an outcome of the toll, sorting with any charge the schedule sets, and a schedule written on the mix would move its own base.

At 200 dollars for a person-log-point-year the differential is 264,000 dollars below zero, meaning the canal should be cheaper, not dearer. It reaches zero at 1,279, where the carbon a canal transit saves exactly pays for the exposure it delivers, and rises to 177,000 above zero at 2,000. Across that range the correction runs between a half of the posted toll below and a third above it, a rewriting of a price the authority already sets. Below 1,279 no traffic should be pushed out of the strait, and the schedule beyond its origin only describes what heavier tolls would do. Above it the schedule lists the points a toll-setter could choose among. The policy question turns on one value, and the carbon account returns that same value from a different chain. The agreement is why the exercise runs without ever fixing what exposure is worth.

9.4 The Aggregate Ledger and the Instrument

The diversion's whole bill sits in one ledger, in Table S60, each entry labeled by what it is. Exposure is one externality the account prices and the carbon a second, while the fuel and the extra days are the carriers' own resources, spent because war risk priced higher than the detour, and the canal's receipts are a transfer between treasuries and carriers. Pricing the war itself would need the whole ledger, which is outside what this paper prices. Counting the private fuel beside the carbon moves the balance point to 2,575 dollars, a figure to be read against rather than added in, since the fuel is a private cost and not an externality. Pricing the added ship time at the range the published container charter indices spanned moves the private balance point to between 3,764 and 4,556 dollars, raising the bar the war premium had to clear.

The ledger's displacement row uses the Red Sea's own trade elasticity, the -0.010 per thousand nautical miles estimated beside the Panama term of Appendix S2, so nothing is carried across between two shocks of different kinds. Applied to the 1,699 billion dollars a year that Suez-dependent pairs traded before the closure, it moves 65 billion dollars a year to a different origin instead of destroying it, a floor on the world figure since only the reporting importers enter the sum.

The authority's loss is not the fall in transits, because the toll is charged on net tonnage and the ships that left were the large container ships. Container transits collapsed from 583 in October 2023, the last full month before the attacks, to 120 in November 2025 (Xeneta, 2025). The authority's own accounts and satellite AIS are measured differently, neither is derived from the other, and they agree, in Table S61.[19] Two quantities the toll never sees sit outside that agreement, the exposure

[19]The authority's published revenue runs the same way, falling from a record 10.3 billion dollars in 2023 to 4.0 billion in 2024 (Suez Canal Authority, 2025). The war-risk band in the bound is the opening quote, and the realized path of premia sat above it through 2024, so the revealed lower bound on the arc's effective cost is if anything low. The eventual full return is itself a routing shock the released weight matrix prices in advance, the same arcs crossed in the opposite direction.

the incidence account counts and the trade the published restriction at Panama displaced. An authority sets its toll knowing what it collects and knowing neither of these. The diversion reveals a bound on what the arc cost, and Table S62 states it in the two things a carrier weighs. A voyage burns 611 tonnes of extra fuel and takes 10.8 extra days at a service speed of 16 knots. In aggregate the Cape has since carried 161 million ship-nautical-miles more than its pre-closure pace, 105 million of them following from the transits counted out of the Red Sea.

A mapped boundary is what a levy needs to reach an arc, since a canal authority cannot price the Cape and no toll reaches an arc that passes no gate, but a port-state levy conditioned on entering mapped water does. The seven control areas in force and the North-East Atlantic designation adopted for 2027 are already exactly that mapped water. The control areas themselves move no ship between arcs, but the boundaries they run on are the infrastructure an arc levy needs. Even the newest map traces this paper's reallocation, covering the European approach the diverted traffic sails while excluding the exclusive economic zones around the Canaries, Madeira and the Azores, in the boundaries of Appendix VII of revised MARPOL Annex VI. As a result Las Palmas, the bunkering hub whose tanker calls rose second most on this coastline, sits squarely in the water the mapped instrument leaves out. The nitrogen side of the designation, moreover, binds only on newly built engines.

Under the law of the sea no charge may be levied on a foreign ship by reason only of its passage, Article 26 for the territorial sea, and the transit passage regime for straits confers no charging power of its own (United Nations, 1982). That is why the two tolled waterways on the map are both sovereign canals and why most of the passages in Table 11 cannot carry a toll at all. A charge that varies with the water a ship crossed on the way to a dock is therefore a measure applied at the border that differentiates by the route, not by the good, the terrain of the trade agreements rather than the sea conventions. Differentiating among like goods by their carriage raises the most-favored-nation and national-treatment questions the border-adjustment literature has mapped for carbon (Böhringer, Fischer, Rosendahl and Rutherford, 2022), and the answer runs the same way there, an origin-neutral charge on a measured externality of delivery, not on the good or its origin.

An exposure charge prices a passage by who lives beside it, so the countries whose water it prices hardest are not the countries that own the ships. Read as an obstacle, that is one reason a uniform charge has survived politically where a differentiated one has not. As a design, the same fact is the instrument's constituency, since the coastal states of the dense straits bear the damage and would receive the hypothecated revenue. In contrast, the aviation scheme failed precisely because it taxed foreign carriers over foreign airspace with no beneficiary coalition abroad.[20]

The dock is the margin that remains, and port-state conditioning on the voyage is legally settled. Already the European scheme charges a ship on the character of the voyage that brings it to the dock, with a transshipment rule so that a call split at an entrepôt outside the scheme does not take the voyage out of the scheme's reach. Written down, the instrument is one line, a levy at the

[20]The damage-bearer and the ship-owner are the same government exactly where the charge would bite hardest, so adoption there is self-interested, not imposed. The rule also prices a melting Arctic route at almost nothing, the case the second instrument exists for, since a carbon price prices exactly what an exposure charge cannot see.

dock per transit-weight on the arcs the voyage traversed, the weights the released matrix carries, collected by the port state beside the fees it already collects, at the rate the covering coalition sets.

Verification is the margin on which the instrument can fail. The route is read off the ship's own transponder, a transponder can lie, and this water already carries a sanctioned fleet that turns its own off, so the answers must come from machinery the trade system already runs. Port-state inspection checks the voyage record against the track, and satellite radar covers the water where the track goes dark. Where neither reaches, a routing presumption on the origin-destination pair charges a voyage that cannot show its arc as if it took the dear one, an assignment that is incentive-compatible exactly where the radar layer makes a false track detectable.

The European scheme is the live instance of the thesis with the sign reversed, since it conditions a charge on the voyage but writes the charge on fuel. Priced on it at the no-transshipment bound and the 2024 allowance price, in Table S63, the representative crossing faces a surrender rate rising from 40 percent of covered emissions in 2024 through 70 in 2025 to 100 in 2026, raising both arcs without changing the ratio. A charge on an arc raises the cost of a voyage, and the Panama estimate measures what a cost of known size does to the trade that uses the arc, so the incidence runs partly into freight rates and partly into the quantity shipped.[21]

The reach of such a levy is bounded, since it decentralizes the arc charge of equation (6) only for voyages that terminate at a covered dock. Lanes that cross a gate without ever calling at a covering port, on this map the Asian services and the regional Red Sea feeders, are outside it. They are the residual an international instrument would have to cover, smaller than the whole because most of the exposure a crossing lane delivers is delivered by lanes that end in Europe or Asia. The feeders also floor what any toll can move, since a regional service inside the Red Sea pays every charge at the gate and has no detour to switch to, so its exposure is delivered at any price. That floor is one more reason to attach the instrument to the arc instead of the gate, the same reason the far end of the toll schedule already gave.

One arc price already exists, written on the very base the theory calls for by an agent with no interest in the exposure. War-risk underwriters quote an additional premium for entering listed water, and the southern Red Sea was listed throughout this window. That premium is an arc charge set by agents untouched by the exposure beneath the arc, and it moved the traffic. This water has carried a demonstration of the mechanism before, when the premium Somali piracy put on the Gulf of Aden raised the cost of every voyage that crossed it while leaving the traffic itself in place, at a welfare loss far past what the pirates collected (Besley, Fetzer and Mueller, 2015). The closure studied here is the case where the same kind of premium reached the quantity margin.

[21] Three alliances carry four fifths of the capacity between Asia and Europe over this window, and carriers price against the demand they face (Hummels, Lugovskyy and Skiba, 2009), with a leg's rate moving with the balance of the round trip (Wong, 2022). The split is a direction rather than a division, and the routing margin that carries every result here is unaffected.

10 Conclusion

The environmental cost of trade has been measured from what is produced and where. Yet the leg in between is a choice, and the agent who makes it faces no price on any of the damage the choice determines. When the Red Sea closed, the same goods moving between the same partners began arriving on a path 4,133 nautical miles longer, and their exhaust landed on a different coastline.

Everything the paper measures rests on corridors solved from the paths the diverted ships actually sail. The pollution moves with the ships and arrives on the wind, which is what makes the corridor weight a measure of plume incidence. On the estimation coasts, exposure came off 81 million people and landed on 37 million who earn two fifths as much, so the aggregate improves while the distribution worsens, and which of the two readings governs turns on a single distributional parameter. The efficient exposure charges on the two arcs stand at 2.59 to 1 against the Cape, a ratio set without ever pricing exposure itself. Yet every fuel-scaled instrument in force or under negotiation charges the opposite way. Where the distributional weight is light the correction at Suez runs the other way, since the canal's own carbon saving outweighs what a transit there delivers.

A population and a geometry are all the rule asks for, so it also prices passages where nothing has happened, read as the exposure a transit delivers at the gate itself. The gate price separates the heavily used passages by a factor of thirty-six and comes to effectively nothing at Bering, where no one lives within the plume's reach of the water.[22] A canal authority sees the base it collects tolls on. The Suez authority lost three fifths of that base, a number it already holds and one on which its own accounts and the satellite agree. The authority does not hold the exposure that moved or the trade that was displaced, and neither is small.

The general point does not depend on ships but on an emitter that can choose among arcs, arcs priced by agents who bear none of the damage downstream of them, and damage that depends on which arc is taken rather than only on how much is emitted. Long-haul trucking through mountain passes shares all three, and the Brenner corridor already tolls one arc above its parallels. So does aviation over structured routings, where the North Atlantic's organized tracks are redrawn daily with no price on who sits beneath them. Electricity is the instructive exception, since wholesale markets already price power at the node where it enters the grid, and the grid has taken the step this paper prices for water. The same design question arises wherever an emitter chooses among arcs and the damage depends on which one it takes.

[22] The construction there uses the distributed routing layer's own passage lines within reach of each gate. At the five gates where the solver's lanes also cross, the two geometries return the same ordering of loads. A transit also queues before it crosses, and a ship at anchor emits beside the gate city, so a charge built on transits alone understates what the two canals deliver.

References

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge. 2023. When Should You Adjust Standard Errors for Clustering? *Quarterly Journal of Economics* 138(1), 1-35.

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales. 2019. Shift-Share Designs: Theory and Inference. *Quarterly Journal of Economics* 134(4), 1949-2010.

Allen, Treb, and Costas Arkolakis. 2014. Trade and the Topography of the Spatial Economy. *Quarterly Journal of Economics* 129(3), 1085-1140.

Allen, Treb, and Costas Arkolakis. 2022. The Welfare Effects of Transportation Infrastructure Improvements. *Review of Economic Studies* 89(6), 2911-2957.

Alphaliner. 2024. Scrubber-Fitted Capacity at 10.3 Million TEU, 36.4 Percent of the Container Fleet. *Alphaliner Weekly Newsletter*, February 2024.

Anderson, Michael L. 2020. As the Wind Blows: The Effects of Long-Term Exposure to Air Pollution on Mortality. *Journal of the European Economic Association* 18(4), 1886-1927.

Anenberg, Susan C., Arash Mohegh, Daniel L. Goldberg, Gaige H. Kerr, Michael Brauer, Katrin Burkart, Perry Hystad, Andrew Larkin, Sarah Wozniak, and Lok Lamsal. 2022. Long-Term Trends in Urban NO_2 Concentrations and Associated Paediatric Asthma Incidence: Estimates from Global Datasets. *Lancet Planetary Health* 6(1), e49-e58.

Antweiler, Werner, Brian R. Copeland, and M. Scott Taylor. 2001. Is Free Trade Good for the Environment? *American Economic Review* 91(4), 877-908.

Arnott, Richard, André de Palma, and Robin Lindsey. 1993. A Structural Model of Peak-Period Congestion: A Traffic Bottleneck with Elastic Demand. *American Economic Review* 83(1), 161-179.

Atkinson, Anthony B. 1970. On the Measurement of Inequality. *Journal of Economic Theory* 2(3), 244-263.

Banzhaf, Spencer, Lala Ma, and Christopher Timmins. 2019. Environmental Justice: The Economics of Race, Place, and Pollution. *Journal of Economic Perspectives* 33(1), 185-208.

Baumol, William J., and Wallace E. Oates. 1988. *The Theory of Environmental Policy*. 2nd ed. Cambridge: Cambridge University Press.

Besley, Timothy, Thiemo Fetzer, and Hannes Mueller. 2015. The Welfare Cost of Lawlessness: Evidence from Somali Piracy. *Journal of the European Economic Association* 13(2), 203-239.

Böhringer, Christoph, Carolyn Fischer, Knut Einar Rosendahl, and Thomas Fox Rutherford. 2022. Potential Impacts and Challenges of Border Carbon Adjustments. *Nature Climate Change* 12(1), 22-29.

Bondarenko, Maksym, David Kerr, Alessandro Sorichetta, and Andrew J. Tatem. 2020. *Census/Projection-Disaggregated Gridded Population Datasets for 189 Countries in 2020 Using Built-Settlement Growth Model Outputs*. Southampton: WorldPop, University of Southampton.

Borusyak, Kirill, and Peter Hull. 2023. Nonrandom Exposure to Exogenous Shocks. *Econometrica* 91(6), 2155-2185.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. Revisiting Event-Study Designs: Robust and Efficient Estimation. *Review of Economic Studies* 91(6), 3253-3285.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2020. Geography, Transportation, and Endogenous Trade Costs. *Econometrica* 88(2), 657-691.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. 2024. Difference-in-Differences with a Continuous Treatment. NBER Working Paper 32117.

Cerdeiro, Diego A., Andras Komaromi, Yang Liu, and Mamoon Saeed. 2020. World Seaborne Trade in Real Time: A Proof of Concept for Building AIS-based Nowcasts from Scratch. IMF Working Paper 20/57.

Chay, Kenneth Y., and Michael Greenstone. 2005. Does Air Quality Matter? Evidence from the Housing Market. *Journal of Political Economy* 113(2), 376-424.

Chen, Jiafeng, and Jonathan Roth. 2024. Logs with Zeros? Some Problems and Solutions. *Quarterly Journal of Economics* 139(2), 891-936.

CMA CGM. 2025. Announcement that the INDAMEX service returns to a full Suez routing, 5 December 2025.

Colmer, Jonathan, Ian Hardman, Jay Shimshack, and John Voorheis. 2020. Disparities in $PM_{2.5}$ Air Pollution in the United States. *Science* 369(6503), 575-578.

Conley, Timothy G. 1999. GMM Estimation with Cross Sectional Dependence. *Journal of Econometrics* 92(1), 1-45.

Conley, Timothy G., and Christopher R. Taber. 2011. Inference with "Difference in Differences" with a Small Number of Policy Changes. *Review of Economics and Statistics* 93(1), 113-125.

Convention of Constantinople. 1888. Convention Respecting the Free Navigation of the Suez Maritime Canal, signed at Constantinople, 29 October 1888.

Copeland, Brian R., and M. Scott Taylor. 2004. Trade, Growth, and the Environment. *Journal of Economic Literature* 42(1), 7-71.

Copenhagen Convention. 1857. Treaty for the Redemption of the Sound Dues, signed at Copenhagen, 14 March 1857.

Corbett, James J., James J. Winebrake, Erin H. Green, Prasad Kasibhatla, Veronika Eyring, and Axel Lauer. 2007. Mortality from Ship Emissions: A Global Assessment. *Environmental Science and Technology* 41(24), 8512-8518.

Cristea, Anca, David Hummels, Laura Puzzello, and Misak Avetisyan. 2013. Trade and the Greenhouse Gas Emissions from International Freight Transport. *Journal of Environmental Economics and Management* 65(1), 153-173.

Currie, Janet, and Reed Walker. 2011. Traffic Congestion and Infant Health: Evidence from E-ZPass. *American Economic Journal: Applied Economics* 3(1), 65-90.

Currie, Janet, John Voorheis, and Reed Walker. 2023. What Caused Racial Disparities in Particulate Exposure to Fall? New Evidence from the Clean Air Act and Satellite-Based Measures of Air Quality. *American Economic Review* 113(1), 71-97.

de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2023. Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey. *Econometrics Journal* 26(3), C1-C30.

Deryugina, Tatyana, Garth Heutel, Nolan H. Miller, David Molitor, and Julian Reif. 2019. The Mortality and Medical Costs of Air Pollution: Evidence from Changes in Wind Direction. *American Economic Review* 109(12), 4178-4219.

Diamond, Michael S., and Lili F. Boss. 2025. Conflict-Induced Ship Traffic Disruptions Constrain Cloud Sensitivity to Stricter Marine Pollution Regulations. *Atmospheric Chemistry and Physics* 25(22), 16401-16409.

Donaldson, Dave. 2018. Railroads of the Raj: Estimating the Impact of Transportation Infrastructure. *American Economic Review* 108(4-5), 899-934.

Donaldson, Dave, and Richard Hornbeck. 2016. Railroads and American Economic Growth: A "Market Access" Approach. *Quarterly Journal of Economics* 131(2), 799-858.

Driscoll, John C., and Aart C. Kraay. 1998. Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data. *Review of Economics and Statistics* 80(4), 549-560.

Dunn, Jason, and Fernando Leibovici. 2024. Shipping Disruptions in the Red Sea: Local Shock, Global Impact. VoxEU column, Centre for Economic Policy Research, March 2024.

European Commission, Joint Research Centre. 2024. *GHS Urban Centre Database R2024A*, version 1.1. Global Human Settlement Layer. Ispra: European Commission.

European Parliament and Council. 2014. Regulation (EU) No 421/2014 Amending Directive 2003/87/EC. *Official Journal of the European Union* L 129: 1-4.

Fajgelbaum, Pablo D., and Edouard Schaal. 2020. Optimal Transport Networks in Spatial Equilibrium. *Econometrica* 88(4), 1411-1452.

Ferman, Bruno, and Cristine Pinto. 2019. Inference in Differences-in-Differences with Few Treated Groups and Heteroskedasticity. *Review of Economics and Statistics* 101(3), 452-467.

Feyrer, James. 2021. Distance, Trade, and Income: The 1967 to 1975 Closing of the Suez Canal as a Natural Experiment. *Journal of Development Economics* 153, 102708.

Fowlie, Meredith, and Nicholas Muller. 2019. Market-Based Emissions Regulation When Damages Vary across Sources: What Are the Gains from Differentiation? *Journal of the Association of Environmental and Resource Economists* 6(3), 593-632.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur. 2012. What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NO_x Trading Program. *American Economic Review* 102(2), 965-993.

Gambia Bureau of Statistics. 2024. *The Gambia 2024 Population and Housing Census: Preliminary Report*. Serrekunda: GBoS.

Ganapati, Sharat, Woan Foong Wong, and Oren Ziv. 2024. Entrepôt: Hubs, Scale, and Trade Costs. *American Economic Journal: Macroeconomics* 16(4), 239-278.

Gendron-Carrier, Nicolas, Marco Gonzalez-Navarro, Stefano Polloni, and Matthew A. Turner. 2022. Subways and Urban Air Pollution. *American Economic Journal: Applied Economics* 14(1), 164-196.

Georgoulias, Aristeidis K., K. Folkert Boersma, Jasper van Vliet, Xiumei Zhang, Ronald van der A, Prodromos Zanis, and Jos de Laat. 2020. Detection of NO_2 Pollution Plumes from Individual Ships with the TROPOMI/S5P Satellite Sensor. *Environmental Research Letters* 15(12), 124037.

Hammitt, James K., and Lisa A. Robinson. 2011. The Income Elasticity of the Value per Statistical Life: Transferring Estimates between High and Low Income Populations. *Journal of Benefit-Cost Analysis* 2(1), 1-29.

Hansen-Lewis, Jamie, and Michelle Marcus. 2025. Uncharted Waters: Effects of Maritime Emission Regulation. *American Economic Journal: Economic Policy* 17(1), 37-69.

Head, Keith, and Thierry Mayer. 2014. Gravity Equations: Workhorse, Toolkit, and Cookbook. In *Handbook of International Economics*, Vol. 4, edited by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, 131-195. Amsterdam: Elsevier.

Heiland, Inga, Andreas Moxnes, Karen Helene Ulltveit-Moe, and Yuan Zi. 2025. Trade From Space: Shipping Networks and the Global Implications of Local Shocks. *Review of Economics and Statistics*, forthcoming.

Hernández-Cortés, Danae, and Kyle C. Meng. 2023. Do Environmental Markets Cause Environmental Injustice? Evidence from California's Carbon Market. *Journal of Public Economics* 217, 104786.

Hersbach, Hans, Bill Bell, Paul Berrisford, Gionata Biavati, András Horányi, Joaquín Muñoz Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Iryna Rozum, Dinand Schepers, Adrian Simmons, Cornel Soci, Dick Dee, and Jean-Noël Thépaut. 2023. *ERA5 Hourly Data on Single Levels from 1940 to Present*. Copernicus Climate Change Service Climate Data Store.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. 2016. Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors. *American Economic Review* 106(12), 3700-3729.

Hsiang, Solomon, Paulina Oliva, and Reed Walker. 2019. The Distribution of Environmental Damages. *Review of Environmental Economics and Policy* 13(1), 83-103.

Huangfu, Peijue, and Richard Atkinson. 2020. Long-Term Exposure to NO_2 and O_3 and All-Cause and Respiratory Mortality: A Systematic Review and Meta-Analysis. *Environment International* 144, 105998.

Hummels, David. 2007. Transportation Costs and International Trade in the Second Era of Globalization. *Journal of Economic Perspectives* 21(3), 131-154.

Hummels, David, and Georg Schaur. 2013. Time as a Trade Barrier. *American Economic Review* 103(7), 2935-2959.

Hummels, David, Volodymyr Lugovskyy, and Alexandre Skiba. 2009. The Trade Reducing Effects of Market Power in International Shipping. *Journal of Development Economics* 89(1), 84-97.

International Maritime Organization. 2020. *Fourth IMO Greenhouse Gas Study 2020*. London: International Maritime Organization.

International Monetary Fund. 2024. *IMF PortWatch*. Washington: International Monetary Fund, University of Oxford and Delft University of Technology, accessed 2024.

Ito, Koichiro, and Shuang Zhang. 2020. Willingness to Pay for Clean Air: Evidence from Air Purifier Markets in China. *Journal of Political Economy* 128(5), 1627-1672.

Klotz, Richard, and Julia Berazneva. 2022. Local Standards, Behavioral Adjustments, and Welfare: Evaluating California's Ocean-Going Vessel Fuel Rule. *Journal of the Association of Environmental and Resource Economists* 9(3), 383-424.

Knittel, Christopher R., and Ryan Sandler. 2018. The Welfare Impact of Second-Best Uniform-Pigouvian Taxation: Evidence from Transportation. *American Economic Journal: Economic Policy* 10(4), 211-242.

Kummu, Matti, Maria Kosonen, and Sina Masoumzadeh Sayyar. 2025. Downscaled Gridded Global Dataset for Gross Domestic Product (GDP) per Capita PPP over 1990-2022. *Scientific Data* 12, 178.

Lee, Phillip. 2025. How Far Does a Climate Shock Travel? Effect of the Panama Canal Disruption on US West Coast Maritime Traffic and Air Pollution. SSRN Working Paper 5163721.

Lloyd's List. 2024. Red Sea War Risk Rates See Huge Jump in Wake of Yemen Airstrikes. London: Lloyd's List Intelligence, 17 January 2024.

Lorente, Alba, K. Folkert Boersma, Huan Yu, Steffen Dörner, Andreas Hilboll, Andreas Richter, Mengyao Liu, Lok N. Lamsal, Michael Barkley, Isabelle De Smedt, Michel Van Roozendael, Yang Wang, Thomas Wagner, Steffen Beirle, Jin-Tai Lin, Nickolay Krotkov, Piet Stammes, Ping Wang, Henk J. Eskes, and Maarten Krol. 2017. Structural Uncertainty in Air Mass Factor Calculation for NO_2 and HCHO Satellite Retrievals. *Atmospheric Measurement Techniques* 10(3), 759-782.

Maersk. 2025. Developments in the Red Sea: First Trans-Suez Sailing. Press release, 22 December 2025.

Maersk and Hapag-Lloyd. 2026. Structural Change to the Gemini Cooperation's ME11 Service. Joint announcement, 3 February 2026.

Montgomery, W. David. 1972. Markets in Licenses and Efficient Pollution Control Programs. *Journal of Economic Theory* 5(3), 395-418.

Montreux Convention. 1936. Convention Regarding the Régime of the Straits, signed at Montreux, 20 July 1936. *League of Nations Treaty Series* 173: 213.

Moretti, Enrico, and Matthew Neidell. 2011. Pollution, Health, and Avoidance Behavior: Evidence from the Ports of Los Angeles. *Journal of Human Resources* 46(1), 154-175.

Muller, Nicholas Z., and Robert Mendelsohn. 2009. Efficient Pollution Regulation: Getting the Prices Right. *American Economic Review* 99(5), 1714-1739.

Muller, Nicholas Z., Robert Mendelsohn, and William Nordhaus. 2011. Environmental Accounting for Pollution in the United States Economy. *American Economic Review* 101(5), 1649-1675.

Muñoz Sabater, Joaquín. 2019. *ERA5-Land Monthly Averaged Data from 1950 to Present*. Copernicus Climate Change Service Climate Data Store.

Natural Earth. 2022. *Ocean and Populated Places, 10m Physical and Cultural Vectors*, version 5.1.1. Public domain map dataset.

Notteboom, Theo, Hercules Haralambides, and Kevin Cullinane. 2024. The Red Sea Crisis: Ramifications for Vessel Operations, Shipping Networks, and Maritime Supply Chains. *Maritime Economics and Logistics* 26(1), 1-20.

Oster, Emily. 2019. Unobservable Selection and Coefficient Stability: Theory and Evidence. *Journal of Business and Economic Statistics* 37(2), 187-204.

Panama Canal Authority. 2024. *Advisories to Shipping*. Balboa: Autoridad del Canal de Panamá. Numbered advisories on the maximum authorized draft, the daily booking slots and the auction amounts, each carrying the limit and the date it took effect.

Peng, He, Meng Wang, and Chunjiang An. 2024. Implied Threats of the Red Sea Crisis to Global Maritime Transport: Amplified Carbon Emissions and Possible Carbon Pricing Dysfunction. *Environmental Research Letters* 19(7), 074053.

Pigou, Arthur C. 1920. *The Economics of Welfare*. London: Macmillan.

Pseftogkas, Andreas, Trissevgeni Stavrakou, Jean-François Müller, Maria-Elissavet Koukouli, Dimitris Balis, and Charikleia Meleti. 2024. Shifts in Maritime Trade Routes as a Result of Red Sea Shipping Crisis Detected in TROPOMI NO_2 Data. *Geophysical Research Letters* 51(20), e2024GL110491.

Roback, Jennifer. 1982. Wages, Rents, and the Quality of Life. *Journal of Political Economy* 90(6), 1257-1278.

Romano, Joseph P., and Michael Wolf. 2005. Exact and Approximate Stepdown Methods for Multiple Hypothesis Testing. *Journal of the American Statistical Association* 100(469), 94-108.

Saez, Emmanuel, and Stefanie Stantcheva. 2016. Generalized Social Marginal Welfare Weights for Optimal Tax Theory. *American Economic Review* 106(1), 24-45.

Santos Silva, J. M. C., and Silvana Tenreyro. 2006. The Log of Gravity. *Review of Economics and Statistics* 88(4), 641-658.

Schlenker, Wolfram, and W. Reed Walker. 2016. Airports, Air Pollution, and Contemporaneous Health. *Review of Economic Studies* 83(2), 768-809.

Security Council Report. 2026. The Red Sea. *July 2026 Monthly Forecast*, 1 July 2026.

Shapiro, Joseph S. 2016. Trade Costs, CO_2, and the Environment. *American Economic Journal: Economic Policy* 8(4), 220-254.

Shapiro, Joseph S. 2021. The Environmental Bias of Trade Policy. *Quarterly Journal of Economics* 136(2), 831-886.

Ship & Bunker. 2025. Algoa Bay Bunkering Set to Resume Following Shutdown in Sep 2023. News report, 30 January 2025.

Simonovska, Ina, and Michael E. Waugh. 2014. The Elasticity of Trade: Estimates and Evidence. *Journal of International Economics* 92(1), 34-50.

Sofiev, Mikhail, James J. Winebrake, Lasse Johansson, Edward W. Carr, Marje Prank, Joana Soares, Julius Vira, Rostislav Kouznetsov, Jukka-Pekka Jalkanen, and James J. Corbett. 2018. Cleaner Fuels for Ships Provide Public Health Benefits with Climate Tradeoffs. *Nature Communications* 9, 406.

Suez Canal Authority. 2023a. Circular No. 7/2023: Adjustment of Transit Dues. Ismailia, 16 October 2023.

Suez Canal Authority. 2023b. Circular No. 8/2023: Conditions of the Container Ship Rebate between North-West Europe and the Far East. Ismailia, October 2023.

Suez Canal Authority. 2025. Annual Navigation Statement for 2024. Statement of the Chairman, 16 April 2025.

Sun, Liyang, and Sarah Abraham. 2021. Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. *Journal of Econometrics* 225(2), 175-199.

Tietenberg, Thomas H. 1978. Spatially Differentiated Air Pollutant Emission Charges: An Economic and Legal Analysis. *Land Economics* 54(3), 265-277.

Tinbergen, Jan. 1952. *On the Theory of Economic Policy*. Amsterdam: North-Holland.

UNCTAD (United Nations Conference on Trade and Development). 2024. *Review of Maritime Transport 2024: Navigating Maritime Chokepoints*. Geneva: United Nations.

United Nations. 1982. *United Nations Convention on the Law of the Sea*. Montego Bay: United Nations. Article 26 on charges, Part III on straits used for international navigation.

United Nations Statistics Division. 2026. *UN Comtrade Database*, monthly trade data. New York: United Nations. Extracted July 2026.

U.S. Environmental Protection Agency. 2023a. *Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances*. Washington: U.S. Environmental Protection Agency.

U.S. Environmental Protection Agency. 2024. Mortality Risk Valuation Estimates. Appendix B of *Guidelines for Preparing Economic Analyses*, 3rd ed. Washington: U.S. EPA, National Center for Environmental Economics.

Van Damme, Martin, Lieven Clarisse, Simon Whitburn, Juliette Hadji-Lazaro, Daniel Hurtmans, Cathy Clerbaux, and Pierre-François Coheur. 2018. Industrial and Agricultural Ammonia Point Sources Exposed. *Nature* 564(7734), 99-103.

Veefkind, J. Pepijn, Ilse Aben, Kevin McMullan, Hendrik Förster, Jos de Vries, Gijsbertus Otter, Jacques Claas, Henk J. Eskes, Johan F. de Haan, Quintus Kleipool, Michiel van Weele, Onno Hasekamp, Ruud Hoogeveen, Jochen Landgraf, Ronald Snel, Paul Tol, Paul Ingmann, Robert Voors, Bob Kruizinga, Ruud Vink, Huib Visser, and Pieternel Levelt. 2012. TROPOMI on the ESA Sentinel-5 Precursor: A GMES Mission for Global Observations of the Atmospheric Composition for Climate, Air Quality and Ozone Layer Applications. *Remote Sensing of Environment* 120, 70-83.

Verschuur, Jasper, Johannes Lumma, and Jim W. Hall. 2025. Systemic Impacts of Disruptions at Maritime Chokepoints. *Nature Communications* 16, 10421.

Vickrey, William S. 1969. Congestion Theory and Transport Investment. *American Economic Review* 59(2), 251-260.

Viscusi, W. Kip, and Clayton J. Masterman. 2017. Income Elasticities and Global Values of a Statistical Life. *Journal of Benefit-Cost Analysis* 8(2), 226-250.

Wardrop, N. A., W. C. Jochem, T. J. Bird, H. R. Chamberlain, D. Clarke, D. Kerr, L. Bengtsson, S. Juran, V. Seaman, and A. J. Tatem. 2018. Spatially Disaggregated Population Estimates in the Absence of National Population and Housing Census Data. *Proceedings of the National Academy of Sciences* 115(14), 3529-3537.

Wong, Woan Foong. 2022. The Round Trip Effect: Endogenous Transport Costs and International Trade. *American Economic Journal: Applied Economics* 14(4), 127-166.

World Bank. 2024. *World Development Indicators*. Washington: World Bank.

WorldPop and Center for International Earth Science Information Network. 2018. *Global High Resolution Population Denominators Project: Age and Sex Structures*. Southampton: University of Southampton.

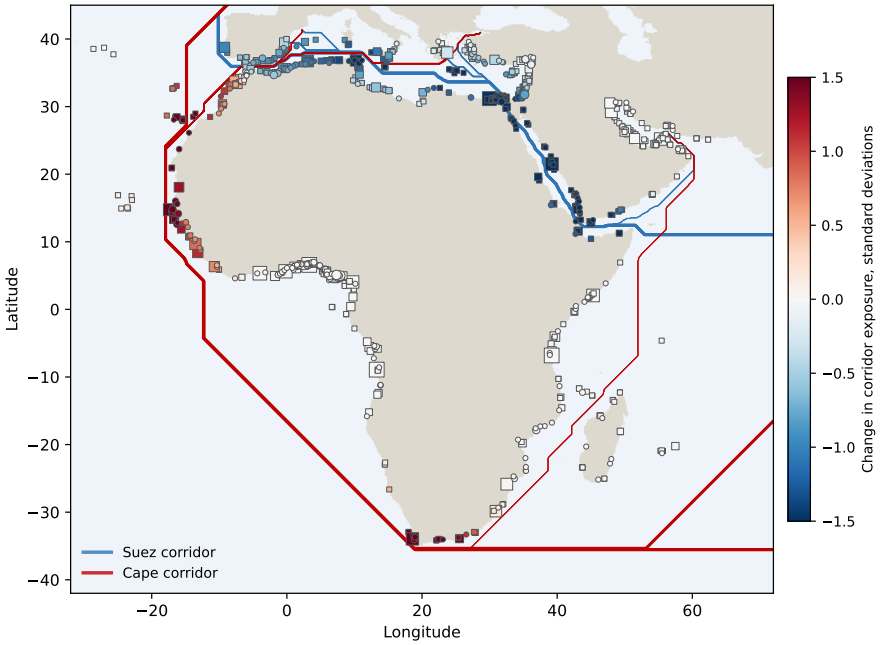
Xeneta. 2025. Largescale Return of Container Ships to Red Sea in 2026? Five Key Considerations for Shippers. Market analysis, citing Suez Canal container transit counts, 23 December 2025.

Zhao, Pengjun, Weiya Kong, Tingting Kang, Yang Chi, and Ming Wang. 2026. The Red Sea Crisis Exacerbates Air Pollution and Environmental Hazards in African Port Cities. *Ocean and Coastal Management* 280, 108300.

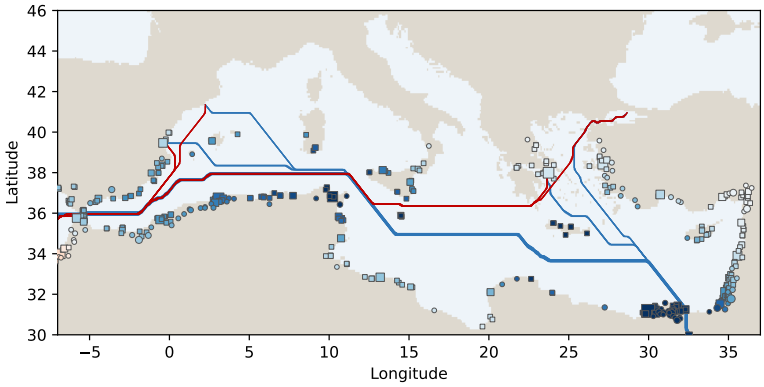
Zou, Eric Yongchen. 2021. Unwatched Pollution: The Effect of Intermittent Monitoring on Air Quality. *American Economic Review* 111(7), 2101-2126.

Figures and Tables

Figure 1: The two corridors and the coastal cities beside them



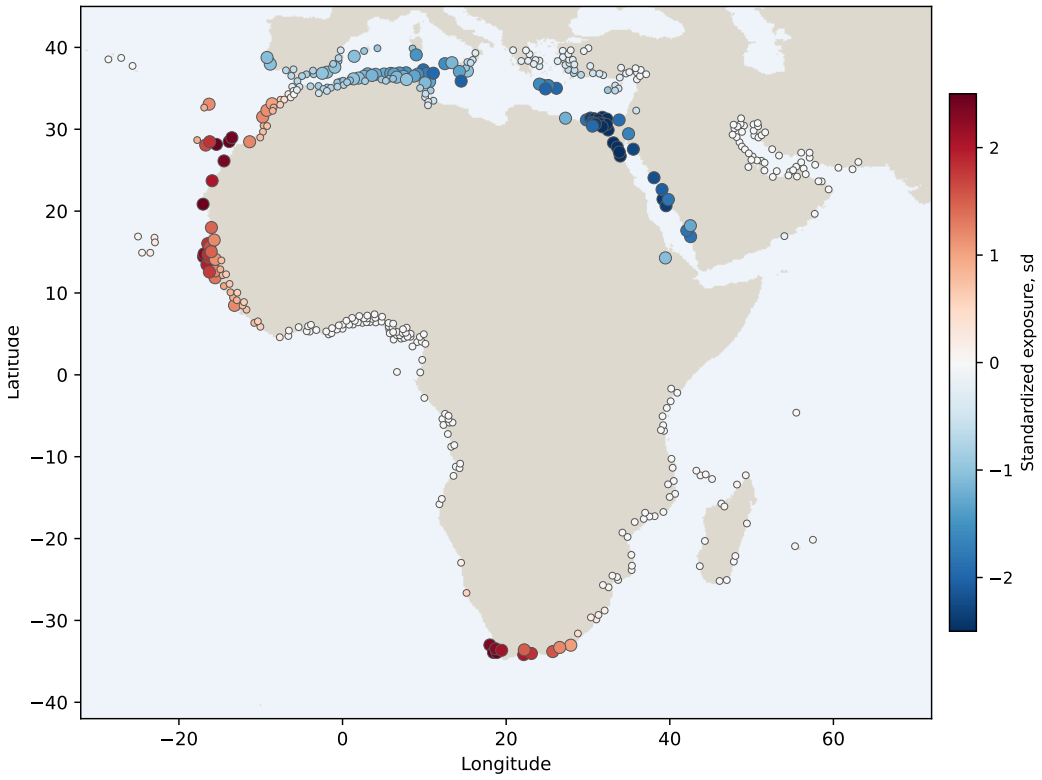
Panel A. The two corridors



Panel B. The Mediterranean, on its own scale

Notes. The colorbar is the change in corridor exposure in standard deviations of the coastal sample, the unit every coefficient is per. Markers are cities within 100 km of the coast, sized by population, squares where a port sits inside the call radius and circles where none does. The lower panel repeats the Mediterranean on its own scale. Coverage is the window drawn, 32 degrees west to 72 east and 42 south to 45 north. Unit of observation is the city. The map is descriptive and carries no standard errors.

Figure 2: The exposure gradient the design runs on



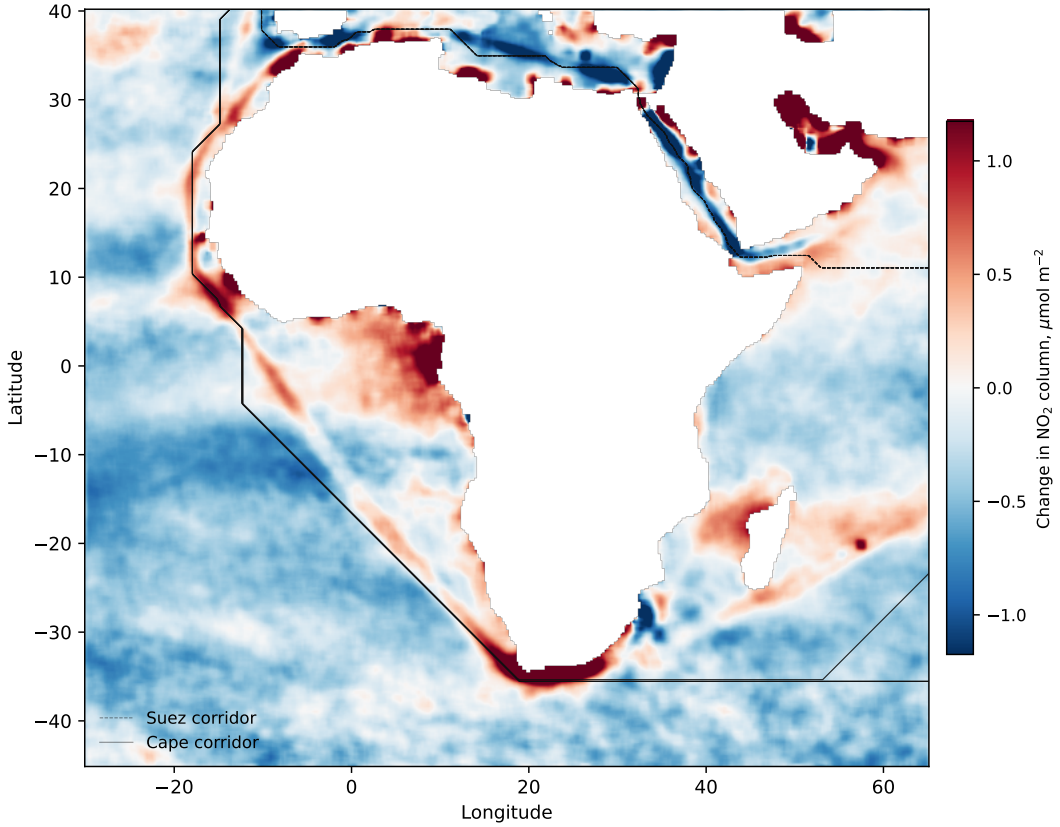
Notes. Every estimation-sample city at its standardized change in corridor exposure, positive where the Cape corridor came closer than the Suez corridor ever was and negative along the coasts the traffic left, capped at 2.5 standard deviations for the color scale, larger markers past one standard deviation. Coverage is the 433 cities of the estimation sample. Unit of observation is the city. The map is descriptive and carries no standard errors.

Table 1: Summary statistics for the three panels

	Mean	SD	10th	90th	Observations	
The coastal city panel, monthly from January 2019						
Nitrogen dioxide column	26.179	19.400	9.742	48.384	38,970	μmol m^{-2}
log nitrogen dioxide	-10.754	0.623	-11.539	-9.936	38,970	
Corridor exposure, standardized	-0.195	1.000	-1.361	0.894	38,970	
Distance to the coast	24.008	31.205	0.600	80.140	38,970	km
Port calls a day, monthly mean	1.880	2.561	0.065	4.467	38,430	
Population	541.413	1,054.526	58.244	1,239.698	38,970	thousands
The four chokepoints the design turns on, daily						
Transits a day	49.367	19.931	29.000	77.000	11,056	
Deadweight crossing a day	2.295	1.449	0.801	4.070	11,056	million tonnes
The trade panel, importer by exporter by chapter by month						
Trade value	8.117	107.772	0.000	5.849	10,072,517	$m
log trade value, dollars	10.968	3.729	5.766	15.582	10,072,517	
Canal dependence	0.656	1.933	0.000	2.621	10,072,517	000 nm
Draft shortfall	0.023	0.042	0.000	0.120	10,072,517	
The twenty-eight passages the world table prices						
People the passage reaches	6.598	12.026	0.000	16.628	28	millions
Transits a day	70.474	69.059	10.018	180.144	28	
Exposure a transit delivers	87.163	119.784	2.303	190.610	28	million person-weights

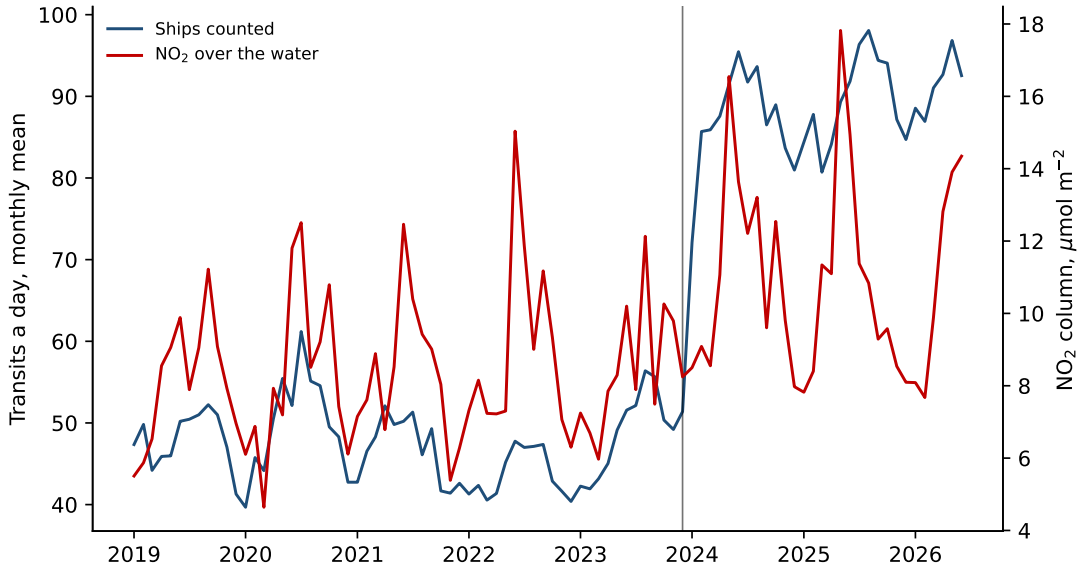
Notes. The city panel is cut to cities within 100 km of the coast. The column rows are in micromoles per square meter and the log rows are taken on the retrieval's own units of moles per square meter, which is why the two do not differ by a constant. Exposure is divided by its standard deviation among coastal cities. The trade rows count the cells carrying a positive trade value, 10,072,517 of the 17,544,694 the panel of Table S3 holds, before the fixed effects drop the cells they cannot identify. Coverage is each panel over its own window. Unit of observation is the city by month, the passage by day, the pair by product by month and the passage. The table is descriptive and carries no standard errors.

Figure 3: Nitrogen dioxide over the ocean, after the closure against before



Notes. Mean tropospheric NO_2 over January 2024 to June 2026 minus the mean over January 2019 to November 2023, at 25 km, with land masked by the ocean raster the solver used. Lines are the twelve heaviest solved lanes in each state. Coverage is the window the figure draws, 30 degrees west to 65 east and 45 south to 40 north, the satellite extraction's own footprint. Unit of observation is the 25 km ocean cell. The corridor water is enumerated rather than sampled, so the figure carries no standard errors.

Figure 4: Ships and air at the point where the ships are counted



Notes. Monthly means at the Cape of Good Hope, January 2019 to June 2026. The rule marks December 2023. Transits are the count PortWatch reports at that point, and the column is the mean over the water inside a 150 km box around it. Coverage is the ninety months the two series share. Unit of observation is the month. Standard errors are heteroskedasticity-robust, as in Table 2.

Table 2: Nitrogen dioxide over the water on ships counted there

	(1) log NO_2	(2) log NO_2	(3) log NO_2
log ships counted	0.413***	0.413***	0.381***
	(0.052)	(0.061)	(0.069)
Counting point	Cape of Good Hope	Cape of Good Hope	Cape of Good Hope
Pixels averaged	All water in the box	All water in the box	All water in the box
Standard errors	Robust	Newey-West, six lags	Newey-West, six lags
Window	2019-2026	2019-2026	Pandemic years out
Mean of dep. var.	-11.624	-11.624	-11.586
Observations	90	90	66
Adjusted R^2	0.632	0.632	0.625

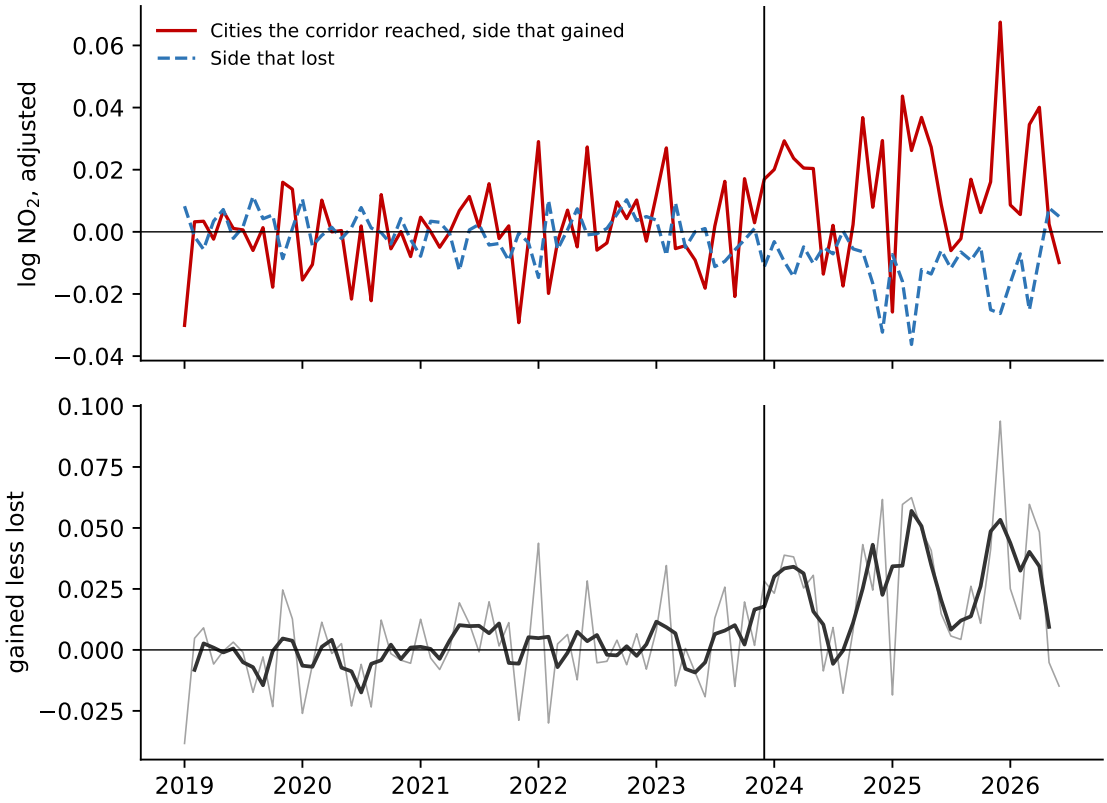
Notes. Month-of-year effects absorb the seasonal cycle. Column (2) replaces the heteroskedasticity correction with a Newey-West kernel at six lags. Column (3) drops 2020 and 2021. Coverage is the 90 months from January 2019 to June 2026 at the Cape of Good Hope. Unit of analysis is the month. Each column's standard error convention is named in its own row.

Table 3: Nitrogen dioxide on corridor exposure

	(1)	(2)	(3)
	Y: log NO_2 column		
Exposure × Post	0.035***		
	(0.005)		
Exposure × Own-gate traffic gone		0.047***	
		(0.007)	
Exposure × Post, side that gained			0.026***
			(0.010)
Exposure × Post, side that lost			0.043***
			(0.009)
City, country-month FE, city trends	Yes	Yes	Yes
Mean of dep. var.	-10.745	-10.745	-10.745
Observations	38,160	38,160	38,160
Adjusted R^2	0.944	0.944	0.944

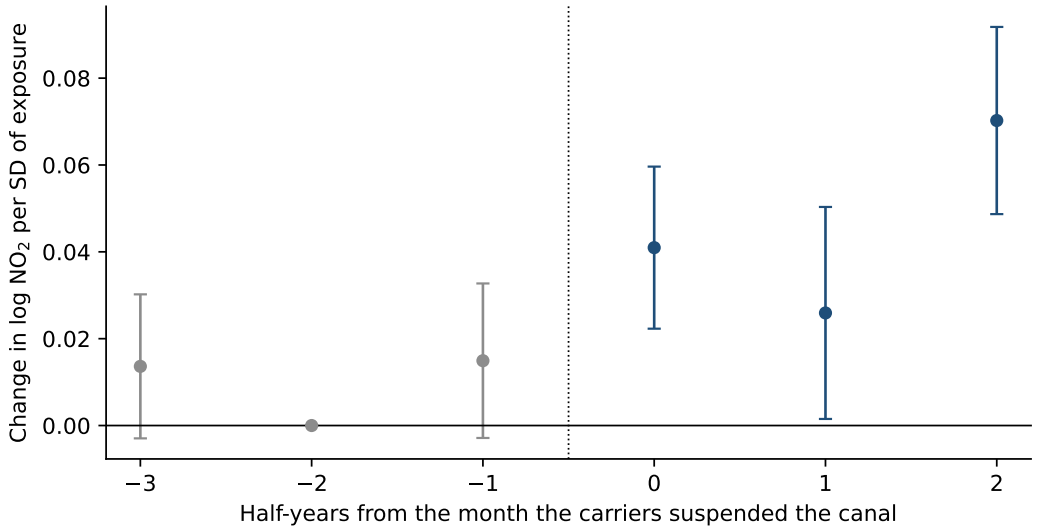
Notes. Equation (8). Column (1) enters the shock as a switch at December 2023, Column (2) the corridor share gone at Bab el-Mandeb by month, Column (3) the switch split by side, where a positive coefficient on the emptied side reads as a fall. Exposure is signed and standardized on the coastal sample's pooled distribution. Coverage is cities within 100 km of the coast with coasts at war excluded, January 2019 to June 2026. Unit of analysis is the city by month. Standard errors are clustered on the city.

Figure 5: Two coasts, month by month, before the regression



Notes. Each city's log column is demeaned within its country-month and against its own pre-closure line. The lines average cities with standardized exposure above one half on the gaining side and below minus one half on the losing side, each net of pre-period month-of-year means. The lower panel is that difference with its three-month centered mean, and the vertical line is December 2023. Coverage is the estimation sample. Unit of observation is the city-month, averaged to the group-month, and the series are drawn rather than estimated, so no standard errors are carried.

Figure 6: The response by half-year around the diversion



Notes. Each point is corridor exposure interacted with a half-year of event time, from a regression with city and country by month effects and the calendar month interacted with exposure, weighted by city population. The base is the half-year ending in May 2023. Bars are 95 percent intervals. Coverage is cities within 100 km of the coast with coasts at war excluded, over January 2019 to June 2026 with 2020 and 2021 removed. Unit of analysis is the city by month. Standard errors are clustered on the city.

Table 4: The response against the wind on each satellite pass

	(1)	(2)	(3)	(4)	(5)
Exposure × Post	0.022***	0.014	0.020***	0.012	
	(0.006)	(0.008)	(0.006)	(0.009)	
Exposure × Post × Downwind		0.017**		0.017*	
		(0.008)		(0.009)	
Exposure × Post × Wind from the lane			0.016***		
			(0.002)		
Exposure × Downwind, pre-closure passes					0.005
					(0.018)
Effect on a downwind pass	n.a.	0.030 (0.006)	n.a.	0.029 (0.006)	n.a.
Effect on an upwind pass	n.a.	0.014 (0.008)	n.a.	0.012 (0.009)	n.a.
Mean of dep. var.	-10.310	-10.310	-10.310	-10.316	-10.314
Observations	349,960	349,960	349,960	324,454	229,368
Adjusted R^2	0.657	0.658	0.659	0.658	0.661

Notes. Downwind is an indicator for the wind projected onto the bearing from the nearest solved lane to the city being positive, and the wind from the lane is that projection divided by its standard deviation. Column (4) drops the calmest passes, 7.3 percent of them, and Column (5) reads the pre-closure passes alone. Coverage is the 120 coastal cities the corridor reaches most strongly, January 2019 to June 2026. Standard errors are clustered on the city. Unit of analysis is the city by satellite overpass with any valid retrieval in the box.

Table 5: The column at cities the ships pass against cities they stop at

	log NO_2
Exposure × Post × Pass-by	0.036***
	(0.008)
Exposure × Post × Call-at	0.034***
	(0.006)
Difference, pass-by less call-at	0.002
standard error	0.008
p	0.822
Mean of dep. var.	-10.745
Observations	38,160
Adjusted R^2	0.944

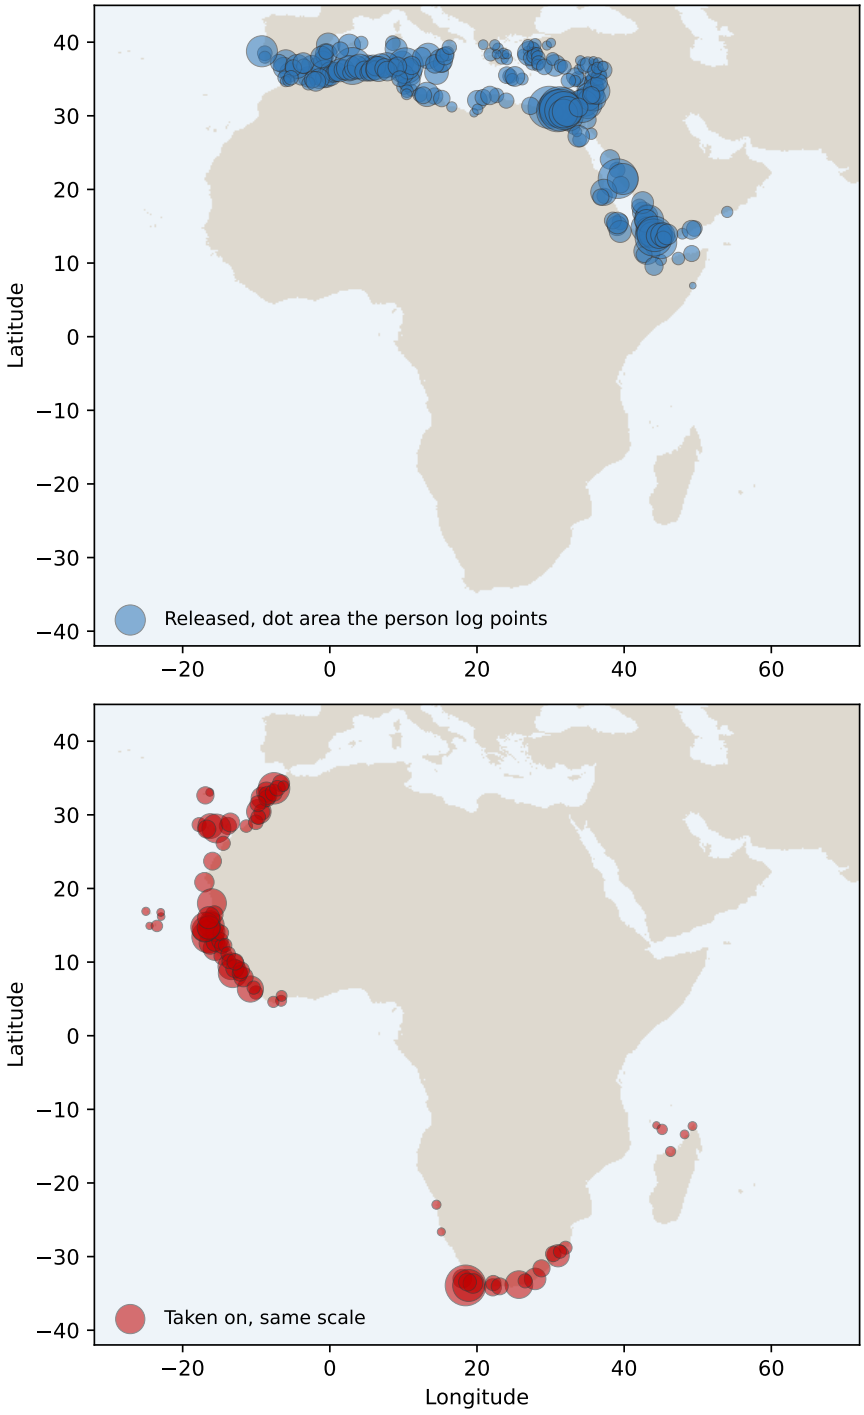
Notes. A city is call-at when a port sits within 40 km of it and pass-by otherwise, a proximity rule rather than an ownership one, so a city that shares its harbor with a neighbor is call-at on both sides of the pair. Coverage is cities within 100 km of the coast with coasts at war excluded. Unit of analysis is the city by month. Standard errors are clustered on the city.

Table 6: The population on each side of the reallocation

	Cities	People, m	Mean Δ log NO_2	Person-weighted
Gained the corridor	88	36.7	0.0352	1,388,613
of which no port of their own	26	5.4	0.0421	248,173
Lost it	233	119.2	-0.0374	-4,990,798
Neither	186	119.0	0.0001	10,878

Notes. The coefficient from Table 3 applied to every city within 100 km of the coast, weighted by the gridded population count inside its box. A city counts as gaining or losing the corridor when standardized exposure lies outside 0.05. The mean is an unweighted average across the cities of a side and the person-weighted total sums each city's own response times its own people, so the two middle columns do not multiply into the last one. Coverage is the 507 coastal cities carrying a population count. Unit of analysis is the city. The quantities are enumerated rather than sampled, so no standard errors are carried.

Figure 7: The reallocation drawn city by city along the coastline



Notes. Every city of the incidence account outside the 0.05 cut, the released side in the upper panel and the taken side in the lower one, dot area proportional to the person log points the account assigns the city and the two panels on one scale. Coverage is the cities of Table 6, coasts at war included since the account counts them. Unit of analysis is the city. Drawn rather than estimated, so no standard errors are carried.

Table 7: The income of the people the plume moved between

	Cities	People, m	Income per head	Below the median	Below $2,000
Gained the corridor	85	36.4	$4,725	65 percent	41 percent
Lost it	209	109.0	$11,748	16 percent	13 percent
Neither	185	118.7	$7,689	82 percent	26 percent

Notes. The median column counts people at or below the coastline's population-weighted median income. The below-median column is the share of that population below that income. The final column repeats the count against an absolute line of 2,000 dollars. Income per head is the country's, from the World Development Indicators (World Bank, 2024) for 2022, and figures are population-weighted. Coverage is 479 coastal cities with both a population count and an income figure. Unit of analysis is the city. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table 8: The sign of the reallocation, at each weight on the poor

Inequality aversion	Taken on	Released	Net
0.0	1.39	-4.22	-2.83
0.5	2.62	-5.46	-2.84
1.0	5.99	-8.23	-2.24
1.5	15.63	-14.20	1.43
2.0	44.33	-28.47	15.85
2.5	132.75	-67.17	65.58
3.0	412.50	-183.15	229.35
Aversion at which the net turns positive		1.38	
The same point with exposure in level units		2.34	

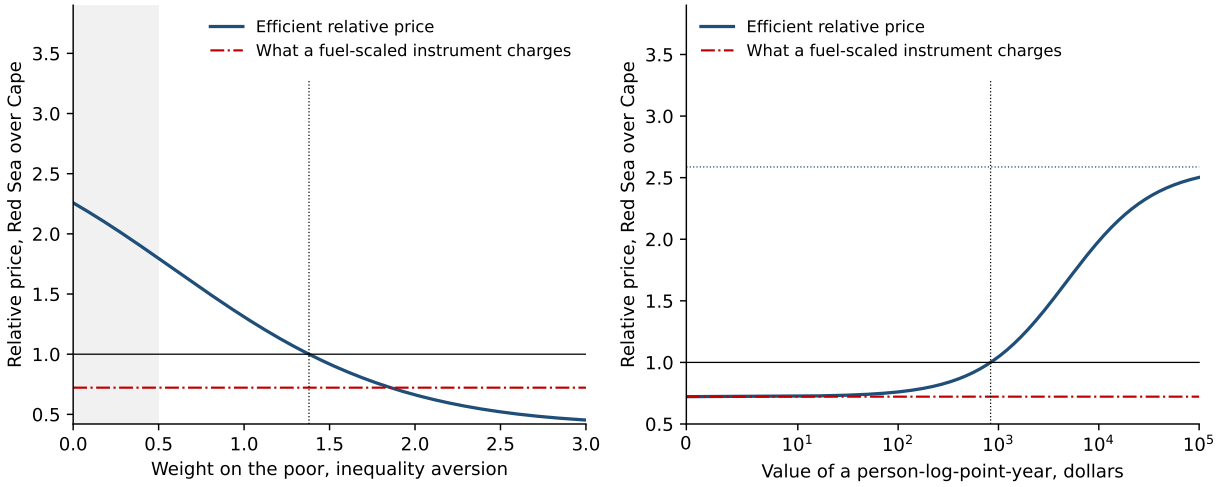
Notes. Person log points of exposure, in millions, from the estimate in Table 3 applied to every coastal city, weighted by population. Inequality aversion enters as a city's income relative to the population-weighted mean of this coastline raised to minus that power, so zero is the unweighted sum and higher values count the poor for more. Coverage is the 479 coastal cities with both a population count and an income figure. Unit of analysis is the city. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table 9: The fuel and carbon the extra distance burns

Extra distance a crossing voyage sails, nm	4,133
Extra deadweight crossing the Cape a day, tonnes	2,546,238
Extra ships crossing a day	40.3
Days since the closure	969
Deadweight-nautical-miles added, 10^{12}	10.2
Carbon dioxide added, million tonnes	64.4
Carbon dioxide added a year, million tonnes	24.2
attributed to this closure	15.8
worth a year at 190 dollars a tonne, billions	3.0
Exposure the diversion removes, million person-log-points	3.6
attributed to this closure	2.4
Value of a person-log-point-year at which the two balance, dollars	1,279

Notes. Extra distance is the difference between the two solved states averaged over the crossing lanes with the container weights, against the mean over January 2019 to November 2023. Carbon dioxide per deadweight tonne-nautical-mile is 6.31 grams, the Annual Efficiency Ratio the Fourth IMO Greenhouse Gas Study reports for the international fleet in 2018. Coverage is the 969 days the transit series counts, from the closure to late July 2026. Unit of analysis is the day. The quantities are enumerated rather than sampled, so no standard errors are carried.

Figure 8: The relative price of the two arcs, by income weighting and by the value of exposure



Notes. Both panels plot the price of the Red Sea arc relative to the Cape arc. The left panel is the income-weighted exposure charge as the weight a planner puts on the poor rises, with a shaded band to an aversion of one half. The right panel is the total efficient charge, carbon included, as the value of exposure rises against carbon at equal weights. Coverage is the 479 coastal cities carrying an income figure in the left panel and all 507 in the right. The dash-dotted line is what every fuel-scaled instrument now in force charges. Dotted verticals mark where that ratio passes one. The dotted horizontal on the right is the exposure ratio the price approaches once exposure dominates the ledger. Unit of analysis is the arc. The coastline is enumerated rather than sampled, so both panels carry no standard errors.

Table 10: The efficient exposure charge on the two arcs

	Red Sea arc	Cape arc
Efficient exposure-charge ratio per tonne of fuel, same hull either arc	3.73	1
at the fitted 225 kilometer length	3.54	1
Efficient exposure-charge ratio per transit, Cape = 1	2.59	1
at the fitted 225 kilometer length	2.44	1
the part from reaching more people	1.39	1
the part from running closer to them	1.86	1
People the arc reaches, millions	133.6	96.0
of whom both arcs reach	63.7	63.7
Exposure delivered to people, millions of person-weights	112.8	43.6
Mean weight on a person the arc reaches	0.845	0.454
Efficient ratio per transit, the fleet each arc carried	1.96	1
with emissions scaling as deadweight to the 0.65	2.16	1
Price a fuel levy implies, Cape = 1	0.69	1

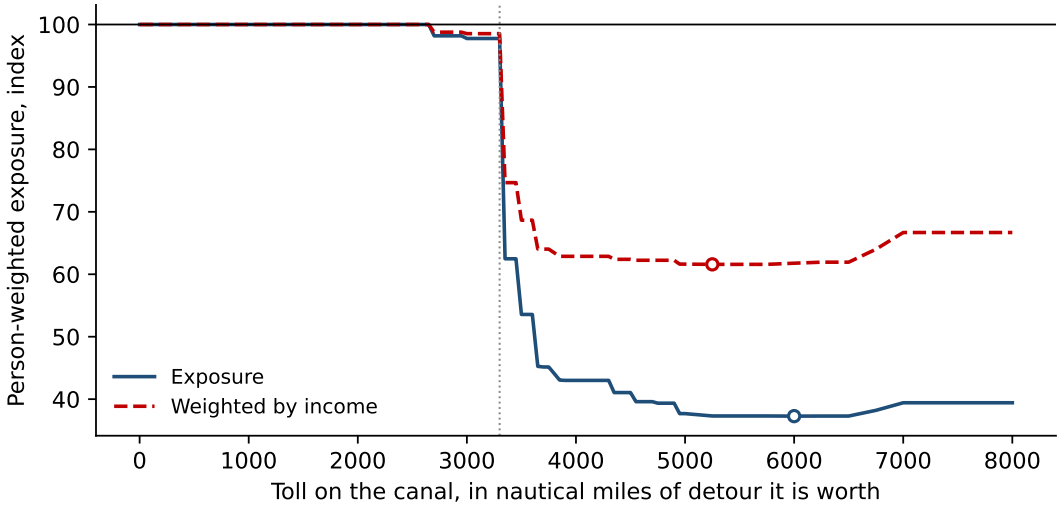
Notes. Exposure delivered is the population of each coastal city times its corridor weight, summed over the sample, and a person-weight is one person times the kernel weight at their city. Reached means a corridor weight above a twentieth of its own standard deviation across the two arcs. Coverage is the 507 coastal cities the incidence table uses, with the gridded population count inside each city box. Unit of analysis is the city. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table 11: The efficient exposure charge on a transit at every monitored passage

Passage	People at the gate, m	Along the passage, m	Transits a day	Exposure a transit delivers	Efficient charge	Across the decay's fit
Sunda Strait	44.4	74.1	27.1	497.8	100	100-100
Taiwan Strait	13.6	74.6	245.9	399.6	80	70-100
Suez Canal	45.6	57.8	60.4	199.4	40	37-43
Malacca Strait	9.7	26.4	182.4	186.8	38	36-46
Mindoro Strait	0.2	39.2	47.9	181.0	36	36-40
Luzon Strait	0.0	48.3	69.6	168.0	34	25-51
Korea Strait	7.6	12.0	227.2	128.7	26	18-41
Gibraltar Strait	4.1	19.1	129.8	102.9	21	18-24
Bohai Strait	6.3	28.9	179.2	85.1	17	13-24
Lombok Strait	4.7	9.3	35.0	83.3	17	12-23
Dover Strait	15.3	26.7	164.1	71.8	14	13-16
Panama Canal	2.2	2.4	31.8	50.8	10	8-16
Tsugaru Strait	0.8	3.7	45.9	43.6	8.8	5-15
Bab el-Mandeb Strait	3.6	8.2	60.9	38.6	7.8	7.1-9.1
Bosporus Strait	19.8	8.4	97.0	28.6	5.7	5.0-7.3
Windward Passage	1.3	4.5	15.8	28.4	5.7	4.2-8.1
Oresund Strait	1.9	3.5	51.0	26.0	5.2	3.6-8.5
Mona Passage	0.5	7.3	10.5	25.8	5.2	4.0-6.8
Cape of Good Hope	0.1	7.1	47.5	22.3	4.5	3.9-5.7
Strait of Hormuz	1.0	2.8	74.8	21.0	4.2	3.5-6.0
Makassar Strait	0.0	1.2	53.0	14.4	2.9	2.1-4.5
Yucatan Channel	0.6	2.1	43.9	13.7	2.7	1.9-4.1
Ombai Strait	0.7	2.5	11.4	9.8	2.0	1.7-2.9
Balabac Strait	0.0	0.7	10.9	7.8	1.6	0.8-3.4
Kerch Strait†	0.6	–	36.9	2.3	0.5	0.2-1.1
Torres Strait	0.0	0.4	9.0	2.2	0.4	0.3-0.7
Magellan Strait	0.3	0.1	3.7	0.6	0.1	0.1-0.2
Bering Strait	0.0	–	0.8	0.0	0.0	0.0-0.0

Notes. The efficient exposure charge is each urban center's population against the corridor kernel's line integral, indexed with Sunda's at 100. Headcount columns count people within 150 kilometers of the gate point and the passage line. The dagger marks the war-collapsed Kerch Strait, and a dash an along-passage cell with no passage line near the gate. Coverage is the 11,422 urban centers of the global settlement database against the twenty-eight monitored passages. Unit of analysis is the passage. The quantities are enumerated rather than sampled, so no standard errors are carried.

Figure 9: Exposure under a canal toll, weighted and unweighted



Notes. At each toll the lanes whose detour costs less than it divert, the weights are rebuilt, and exposure is each coastal city's population times its weight. Both series are indexed to the toll of zero. Hollow circles mark each series' floor and the dotted vertical the revenue peak. The weighted series divides exposure by income relative to the population-weighted mean. Coverage is the 507 coastal cities the incidence table uses and the 140 crossing lanes. Unit of observation is the toll. The quantities are enumerated rather than sampled, so no standard errors are carried.

Online Appendix

S1 Threats to Identification

The estimate of Table 3 could appear without a plume in three ways, through the choices behind the specification, through the inference, and through a real correlation with something else on the same coastline. One contrast carries the paper's estimate wherever it is quoted, the switch on the full sample. Every entry in the checks Table S64 collects on one page is a named departure from it, and the four pollutants, the one family the paper tests jointly, carry their own stepdown correction.

S1.1 Specification

Table S65 prices a lane three ways. The first runs along every mile of its path, the second reads its nearest approach only, and the third scales the path by how often the air moves from the lane to the city. The three read 0.035, 0.029 and 0.031, so the deepest choice in the design is not a choice the result rests on. Taken jointly, the rest of the specification carries no more weight. The decay length, the inland cutoff, the canal country, the port ranking, the form of the shock and the clustering are each defensible more than one way. With the city trend, the post month and the population weights varied beside them there are 2,880 distinct specifications. Across every one of them the median in Figure S12 is 0.039, the tenth and ninetieth are 0.016 and 0.053, every one is positive, and nine in ten are significant on their own, with the reported 0.035 below the median, in Table S66. With the city trend, the season interacted with exposure and the population weights varied one at a time in Table S67, neither the season nor the weights moves the estimate by more than 0.008.

The trend is the one choice with content, and city-specific seasonal cycles read 0.026. Yet exposure is fixed geography, so a seasonal cycle correlated with it is partly the plume itself, arriving more in the months the wind serves the lane, and absorbing every city's own calendar over-absorbs treatment as well as threat. A city-specific quadratic beside the linear trend leaves 0.028, and opening the window after the sensor's along-track pixel shrank or after 2019 leaves 0.030 and 0.035, in Table S68.

The estimate without a city trend is 0.018 against the event study's 0.046 over its post bins in Table S69, two no-trend readings that differ because the bins face a base one half-year deep while the switch absorbs five pre-years of the shortfall. Fitting the trend on the pre-period alone gives 0.036, the whole window 0.035. The headline's trend is a correction whose form the fifteen-year OMI gradient motivates, estimated at -0.0031. The first post bin tolerates a departure from parallel trends 1.6 times the largest pre-period deviation before losing separability, between 1.5 and 1.75 without the open-ended terminal bin. At half that benchmark, the post-period average carries a robust confidence set of 0.021-0.090. The leads without trends test parallel trends and are jointly flat in the figure's own machinery.

Designs that carry a within-window city trend return 0.035 on the column and 0.029 on tanker calls, while designs that read the post-period against the whole pre-period without one return half

that, because the cities the corridor reaches most were rising more slowly than their coastlines before the ships moved. Therefore the pre-fitted reading, which no post-treatment month informs, is the largest of the family's five and the strongest answer the paper can give. The trend is also older than the panel, since over 2005-2018 an earlier instrument reads the same within-country gradient at -0.003 log points a year, in Table S70, two thirds of the Sentinel-period size. The correction therefore removes a gradient older than the instrument, the pandemic and the closure by more than a decade. One coefficient scales both sides of every ratio, so no crossing and no ordering turns on which reading a reader prefers.

In Table S71 the fixed effects enter in steps instead of varying, and the estimate rises from 0.008 with city and month effects alone to 0.035 once country-specific month effects and city trends are in. The coarser designs understate the response rather than manufacture it, the direction attenuation predicts, since a common month effect cannot absorb a national shock that lands on one coast. Under selection on unobservables proportional to the selection the observables reveal, the Oster (2019) adjustment at $\delta = 1$ and $R_{\max} = 1$, the bias-adjusted estimate is 0.049. The moderate benchmark of 1.3 times the observed fit caps at the same bound, since the fit already stands at 0.944. Thus if the unobserved selection runs the way the observed selection does, the headline is the understatement.

The rule that decides which cities enter carries no weight either. I keep two cities whose retrieval boxes overlap as one, because the incidence account cannot let the same air stand behind two populations, and that rule removes 797 of the 1,304 coastal cities the sensor covers. On the weights the account holds, the box rule puts 346.6 million people on this coastline and the exclusive assignment puts 275.0 million, so a fifth of the coastline was counted twice, in Table S72, with one stretch of the western Mediterranean drawn under the three accountings in Figure S13.[23] Every population in this paper is the exclusive count, and every level the incidence account reports is computed on it.

A regression needs no such rule, since overlapping boxes make neighbors' errors correlated, which the clustering already allows for, and Table S73 puts every removed city back. The estimate is 0.035 on the full set against 0.035 on the thinned one, unchanged when the clustering is widened to a one-degree cell, and 0.036 on the removed cities by themselves. A coarser dose, the corridor-wide share, puts one number on every coast and reads 0.052 against the gate-matched 0.047, in Table S74. The Mediterranean alone reads 0.069 and the coasts the diverted ships took 0.059, so matching the gate moves the dose by half a percentage point and changes no conclusion.

S1.2 Inference

Cities on one coastline are not independent draws, since the corridor runs past them in sequence and the plume that reaches one reaches the next. The corrections written for shift-share designs with many independent shocks (Adão, Kolesár and Morales, 2019) do not apply to a design with two arcs, so I ask instead how often a coastline arranged at random produces an estimate this large,

[23] The box narrows with latitude, 64 kilometers of east-west width at 35 degrees north against 75 at 15, so the count is each city's own catchment.

and I read the parametric corrections for dependence in space and in time beside it, the clustering cells widened to five and ten degrees in Table S75 and the panel collapsed to one pre and one post mean per city in Table S76.

The exposure values are reassigned among the cities of the same country a thousand times in Table S77, destroying only the correspondence between a city and its distance to the corridor. The randomized coastlines spread 0.006 around zero and none of the thousand reaches the estimate, and the conservative parametric reading points the same way, the headline clustered on the countries instead of the cities, and it stays significant there. Because neighboring cities share exposure, the sharper permutation rotates each country's whole exposure vector along its coastline by a random offset, so contiguous runs stay contiguous, and no rotation of a thousand produces an estimate as large. The dependence the design implies follows neither the city nor the country but distance, and both permutations preserve that structure. That bound extends to the falsification family jointly, since the probability that a rearranged coastline reproduces the whole observed pattern, the headline with the wind increment, the equalities and the nulls beside it, is no larger than the probability it reproduces the headline alone. None of the thousand does.

The dose specification raises its own inference problem, since with a continuous treatment the two-way fixed effects estimand need not be a convex combination of dose responses (de Chaisemartin and D'Haultfœuille, 2023; Callaway, Goodman-Bacon and Sant'Anna, 2024). Read as a binary contrast, the top tercile of exposure against the bottom, the dose carries 0.051 in Table S78. I rotate the dose path through all ninety offsets of the window, which keeps its serial structure and breaks only its alignment with the closure. Exactly one rotation of ninety produces an estimate as large, and that rotation is the true alignment itself. The hub removal below is run on the dose sample as well as the switch sample.

S1.3 Named Confounders

The candidates a country by month effect cannot absorb are cross-city and correlated with exposure, and I enter each of them directly. The first is the port itself, in two forms. Congestion comes first, and controlling for the ocean state and for port scale interacted with the date moves nothing, in Table S79. I price the South African port crisis, whose queues parked ships beside the strongest gained-side cities, directly in Table S80, and the estimate holds between 0.029 and 0.035 with South Africa, Durban or Cape Town removed. The harbor economy comes second, and on the corridor that filled, the cities with no port of their own carry 0.053 against 0.034 at the port cities, in Table S81, the ordering the plume reading predicts. The most informative test is the lost-side split between passing and calling cities, which has the power to detect what it rules out.

The second is Egypt, where the lost corridor concentrates and a second shock lands on the same date, the Red Sea tourism collapse. Its resorts sit on the Red Sea coast rather than the Mediterranean one. Dropping them leaves 0.032 in Table S82. Dropping Egypt altogether leaves 0.034 against 0.035 with the country in, so the most exposed coastline in the sample is not carrying the coefficient. The sample's three most exposed boxes are the canal-bank cities themselves, Port Said, Al Adabiyah and Ismailia, whose columns carry queueing and convoying as well as passage.

Dropping them raises the estimate to 0.036 with the lost side at 0.046, and Table S49 consolidates the coefficient, the ratio and the crossing under the Egypt and canal-bank cuts in one place. The estimate does not rest on the anchorage.

The third is capacity arriving where the ships did. Off Senegal the Sangomar field delivered first oil in June 2024 and the Orange Basin off Walvis Bay was being drilled, and Table S83 drops the cities that carry them, at 0.034, the gained side alone at 0.037.[24]

The fourth is the calendar, where a month effect cannot hold two movers, the fast that advances eleven days a year across a largely Muslim lost coastline, and the weather. I enter them one at a time and jointly, in Table S84, and the estimate runs 0.033, 0.034 and 0.033. The fast's own term carries 0.036, the size the corridor carries, and the corridor term holds steady under its entry, so the two are separate things. On the pre-closure months alone the seasonal absorption leaves it at 0.016, which the data cannot separate from zero, in Table S85, a control the design was right to carry and not a channel the closure could run through. Two calendar readings remain, and Table S86 closes both. Exposure predicts no differential response to Eid al-Adha, and absorbing a break specific to exposed Egyptian cities at the pound's collapse leaves the corridor term at 0.030.

The fifth is the European trading scheme, which reached shipping one month after the closure on exactly the Union coast. Table S87 splits the response there, at 0.046 off the scheme's coasts against 0.021 on them. That gap runs the wrong way for a scheme story. The African and Anatolian Mediterranean, where no scheme charges a call, carries 0.051. The bloc is not short of variation that could have carried a scheme effect, since exposure varies as much within it as outside it. As a result, the cleaner test is exclusion, and dropping every scheme coast entirely leaves the switch at 0.046, and it remains significant on the cities that stay in. The scheme's other margin, transshipment calls moved to boundary hubs, is priced in Table S88, where removing the five hubs leaves 0.036 on the gate-dose sample, 0.036 on the baseline sample, and 0.054 on the arrivals either way.

The sixth is the air itself. South of twenty degrees south, where the filled corridor runs through air with no lightning belt and no biomass season, the estimate is 0.073, twice the pooled reading, in Table S89. The southern reading shows where the signal is cleanest, not where the exposure is largest. Whether the column stands for what people breathe can be checked only in Europe, where nine cities of the sample carry a station sited long before the closure. Within a city the street tracks the column at a slope of 0.90, falling to 0.29 once every common month is swept out. Read one coastline at a time the same slope is 0.51 where the corridor emptied and 0.71 where it filled, in Table S90. A box centered on a coastal city is also part water, so Table S91 rebuilds the outcome from land pixels alone, at 0.033, and holds the count of valid retrievals fixed. Nothing moves. Because the panel is balanced with no missing and no non-positive city-month, which monthly compositing over a 0.7 degree box essentially guarantees, the log creates no selection. The count-form estimator on the level column returns 0.044 with the trends in, in Table S92.

The variation within a single country carries the estimate. The two countries with four or more cities on each side, Spain and Morocco, the only designs nothing national can enter, give 0.018 and 0.022 on the most demanding cut the data allow. Country by month effects leave 0.272 of the

[24]The Dangote refinery at Lekki started over the same window and is not on the list.

exposure variance standing, and the country removals read what that residual carries. Figure S14 removes each of the 53 countries in turn. The estimate runs from 0.029 without South Africa to 0.043 without Spain, and it is statistically distinguishable from zero every time. Table S93 groups the removals.

The design can also be carried east whole, since the closure substituted the Sunda Strait for the Malacca approach inside Indonesian water. Of the two margins there, the counts validate one and reject the other, since Sunda's rise matches the solved substitution, that strait carrying intra-Asian traffic the closure never touched and the Singapore hub the diversion itself fed. On the validated margin Table S94 finds a null against a realized Sunda dose of 0.009 at the headline response.

S1.4 Placebos and the Measurement of Exposure

Every null in the paper carries the smallest effect the design could have found, in Table S95, and those bounds put the gas or the outcome below the treatment the same design finds elsewhere. In Table S96 the Red Sea weight shows nothing on the pre-period level while the Cape weight leans negative, so the coast the traffic moved onto was if anything the cleaner one before it moved. City effects absorb any level, so the margin that must be balanced is the pre-period trend, and there the exposure carries -0.005 log points a year against a coastline that trends up 1.9 percent a year. Both margins run against the estimate, and on the trend margin the composite nets two signed components, the Cape weight leaning negative and the Red Sea weight positive. A variant that hands every side-level trend to fixed effects outright, side by month effects beside the country by month ones, leaves the estimate at 0.035, so the fixed effects absorb the component imbalance.

Across the 31 countries holding four or more coastal cities the exposure of a city and the one next to it correlate at 0.39, so most of the difference between two neighbors comes from their own geometry. Dropping Spain and Morocco, the two coasts that sit beside both corridors at once, raises the estimate to 0.046, in Table S68, and the same table enters, beside a city's own exposure, the mean exposure of its neighbors within 300 kilometers, where the neighbors' term is zero while the city's own holds. Instrumenting the line integral with the nearest-approach kernel inside each country moves Spain's 0.018 to 0.017 and Morocco's 0.022 to 0.023, with reliability ratios at one, so the solver's error does not hold those readings down, and the pooled estimate rises when the two countries leave because the sample changes, not because the measure improves.

Moved to the two clean Decembers before the real one, the shock finds 0.006 and 0.002 in Table S97, with confidence intervals that include zero. With the shock held fixed, the geometry moves instead in Table S98, with the same kernel built from the thirty lanes whose solved distance is identical in both states. Entered together, the real measure holds 0.034 and the unchanged lanes carry -0.001, so the estimate is not a generic feature of where ships sail. Entered alone, those lanes carry the capacity withdrawal itself, in Table S99.

The diversion absorbed fleet capacity and carriers canceled sailings on strings whose distance had not moved, so the cities beside them lost port calls at -0.019 and lost column at -0.013, a second treatment of the opposite sign. I therefore remove it from the comparison group directly. Dropping every city whose unchanged-lane kernel weight exceeds a twentieth of that weight's spread, 132

cities, raises the estimate to 0.047, and dropping the comparison cities within 150 kilometers of a top-tercile treated city leaves it at 0.035. As a result, the withdrawal beside the unchanged lanes was holding the headline down, not propping it up, and the incidence account of the paper prices that same withdrawal as an uncounted improvement running in the same direction. The joint entry of the real and the control kernels, 0.034 beside -0.001, is the conservative statement the paper rests on.

The window is bracketed by two events, and the reading survives truncation at either end. Suez raised its tolls in January 2024 with the crossing lanes exempt through June, and the switch estimated on the window ending in May 2024 alone, while the exemption still held, is 0.038, in Table S13, so the response does not wait for canal pricing to reach the treated lanes. At the other end the strikes at Hormuz in 2026 sent Gulf traffic around the Cape for reasons of their own, so in Table S100 I re-estimate the switch, the dose and the trade term on the months before March 2026. The switch moves from 0.035 to 0.034, the dose holds at 0.101, the trade term at -0.007. Nothing read off the closure rests on the months the second conflict occupies.

Moving the dates and thinning the sample leave the estimate where it was. Under the rule that excludes the pandemic years, the continuous dose reads 0.051 against the full window's 0.052, in Table S101. Moving the month the post period begins across the dates a reader could defend, December through February, runs the estimate from 0.035 to 0.041. Opening the window in October or November counts attack-free months as treated and returns 0.024 and 0.027, dilution rather than a smaller response, in Table S102. Widening the coastal band from 25-300 kilometers moves it from 0.036 to 0.029, drift from cities the corridor barely reaches, in Table S103. Dropping the most exposed cities, ranking the exposure instead of measuring it, or trimming the outcome's own tails leaves the estimate between 0.033 and 0.041 in every cut, in Table S104.

The remaining checks vary how exposure is measured, not when it is read. Rebuilt on every hull class and allowed to empty at its two ends in the proportion the transits there fell, the corridor reads 0.030 and 0.023 in Table S105. The endpoint set is rebuilt five ways in Table S106, merged terminals, twenty ports a side, equal lane weights, length-normalized weights, and the Alexandria agglomeration forced into the western set. Across the five the estimates run 0.035-0.038 with arc ratios 2.14-2.59, while rebuilding the weights on each pre-year alone moves the ratio only between 2.57 and 2.61, in Table S107.[25] I sweep the decay scale across 50-500 kilometers in Table S108, and the estimate runs from 0.018 to 0.053, monotone in the scale as a plume weight should be.[26]

The shock can also be entered as the plume itself rather than as a switch, in Table S110, where it raises nitrogen dioxide by 0.101 per unit. Split by corridor, the emptied corridor's own traffic carries the movement, at 0.196 with Egypt in and 0.186 with it out, while the filled corridor's

[25] A liner string that calls at several ports enters as several lanes over the same water, and the gate-count validation is invariant to that structure, since a transit is counted once at a gate however many port pairs its string serves. No Egyptian hub ranks among the western ten because the ranking counts berthings and Egyptian container traffic splits across four ports, none of which alone reaches the tenth port's calls. Much of Port Said's throughput is transshipment on strings that transit without berthing there. Forced in, the Alexandria agglomeration returns 0.035 and a ratio of 2.41.

[26] The length is itself fitted, and integrating over its own sampling distribution returns 0.040 against the reported 0.035, in Table S109. The solved paths, the standoff and the lane weights enter as data, and Table S52 sweeps them one at a time.

transits carry nothing on their own. The two paths imply that asymmetry. The Cape filled once and stayed full, so its identifying variation is the switch, not the month-to-month movement, and the switch's own split column carries that side at 0.026. In contrast, the Suez path rose and partly returned, and that movement is the variation a traffic term can price.

Inside the kernel, the decay length is measured on the water itself, on an instrument sensitive enough that the plume of a single large ship has been picked out of the column (Georgoulias et al., 2020). Fitting an exponential to the change in the column against distance from the solved path, over the ocean cells of the filled corridor's own open water and no city at all, gives 225 kilometers, in Figure S15 and Table S111. A refit without the bins where the South Atlantic distributary lanes enter returns a length inside the fit's own error. The same profile validates the solver, since the solved paths enter nowhere in the column, and its enhancement peaks at zero distance from the solved corridor and dies away from it.

The two corridors do not share one length. Run instead on each corridor's exclusive water, the same fit returns 225 kilometers on the Cape's open ocean and 31 in the Red Sea's confined basin, and the short fit is the physics rather than the basin's truncated fitting domain, since a synthetic enhancement decaying at exactly 225 kilometers laid on the Red Sea's own cells comes back from the identical procedure at 226. A confined basin concentrates what open ocean disperses, so the single length understates the Red Sea arc's exposure. Exposure rebuilt with each state's paths at its own fitted length returns a response of 0.022 beside the baseline's 0.035, in Table S112, and significant under the same design. The kernel reads a column while the damage stands at the surface, and entering the mixing depth between them directly leaves the estimate at 0.035 in Table S113, with the split halves at 0.032 and 0.037, one reading at either depth.

S2 The Trade Cost of a Chokepoint

The estimating equation is

$$\ln v_{iekt} \;=\; \phi\left(\Delta_{ie} \times \text{Short}_t\right) \;+\; \mu_{iek} \;+\; \nu_{et} \;+\; \pi_{it} \;+\; \rho_{kt} \;+\; \varepsilon_{iekt} \tag{12}$$

for importer i, exporter e, chapter k and month t, where Δ_{ie} is the sailing distance the pair loses without the canal, in thousands of nautical miles, solved on the same ocean grid with the isthmus as an arc that opens and closes, Short_t is one minus the authorized draft over its unconstrained limit, and v_{iekt} is the monthly import value in current dollars. The triple effect μ_{iek} holds the pair and the product together, so nothing is identified off which countries trade what. Three time effects hold each exporter, each importer and each chapter in each month, so a shock to a country's exports, to a country's demand or to a product's price in any month is absorbed. The variation left is the same pair and product against another pair with less to lose from the isthmus. Inference is clustered on the country pair. The panel is imports, since a reporter's import record covers goods arriving from every partner while its export record covers only its own shipments, and the flows the drought hit arrive rather than leave.

Pricing an arc needs a shock to its cost whose size is known. Because the Red Sea closed at one date for every lane at once, an elasticity there is identified by one pair against another over a common time path, a cross-sectional comparison in panel form. I estimate it at -0.006 to -0.010 per thousand nautical miles of canal dependence across the fixed-effect ladder of Table S114. The -0.010 comes from the column whose window ends before the American tariff round of 2025. That round lands on pairs whose canal dependence correlates with the tariff list, and the truncated window keeps every month the closure ran before that shock arrives. A placebo break in the middle of the pre-period carries 0.001. Rotating the calendar supplies the inference for both readings, with two of 88 shifted calendars producing a coefficient as large on the full window and one of 73 on the pre-tariff window. Figure S16 draws the path by half-year, on the conventions the paper's event studies share. The leads sit within 0.002 of zero. The trade on the added-distance pairs falls on the closure and fades as the freight market normalizes. The Red Sea coefficient prices a posted detour. Panama's -0.070 prices a unit of quantity rationing, the queue whose shadow cost the slot auctions revealed, and the factor of seven between the two measures how much harder rationing binds than a posted price.

The drought over Gatun Lake cut the transits at Panama from 31.9 a day over the five years to October 2023 to 21.7 across November 2023 to January 2024, with the trough at 19.5 in January, and let them back up over the following year. The two canals, the two detours, and the three passages that sit near one of the canals and on neither detour are drawn together in Figure S17. The cut was to draft and never to access, the mirror of the expansion of 2016, whose distance shock runs the other way and carries its own trade literature (Heiland et al., 2025). As the lake fell the authority reduced the maximum authorized draft in the Neopanamax locks step by step from 15.24 meters to 13.41, announcing each cut as a numbered advisory that gave the new limit and its effective date (Panama Canal Authority, 2024). A vessel arriving deeper is required to trim or

off-load cargo before it is allowed through, so the constraint is on cargo per transit and binds on every ship that crosses, and the lake sets the limit. The first cuts came in March 2023, stepped down to the 13.41 meter floor by June, held there through the following spring and were lifted stepwise to the full 15.24 meters by advisory A-28-2024 in August 2024, the date the treatment window closes, so the treatment path spans eighteen months, not one winter.

Suez closed and stayed closed, so a pair is treated from December 2023 onward or it is not. Panama did not close. Instead, exposure rises and falls within the treated group, and the same pairs supply their own control period on the way out, so a trend in Pacific trade would have to bend twice, each time in the right direction, to reproduce the estimate. That two-sided path makes the second episode worth as much as the first, and it runs over the same months as the closure on a geography that shares nothing with the Red Sea, the shock the pricing question needs.

At the deepest cut a pair like the United States and China, which loses 4,362 nautical miles without the canal, trades 3.6 percent less, and the estimate holds where I stress the design. It runs from -0.073 to -0.070 across the fixed-effect ladder in Table S115. A placebo that moves the draft path two years early carries 0.016 in Table S116, and the same shift serves as the anticipation test. Scaling dependence as a route share or in logs leaves both the sign and the magnitude unchanged, in Table S117. Estimated to June 2026 instead of to the end of the restoration year, the coefficient is -0.047 in Table S118, the value that appending a year and a half of treatment-free months should return. Entered together in the same table, the transit shortfall goes to -0.003 and the draft keeps -0.066, so the margin that moves trade is the depth a ship may load to and not the number of ships the canal admits. Value weighting and clustering on the pair and the month at once change nothing, and pair-specific trends read -0.035, the four specifications that matter most sitting together in Table S119.

The El Niño of 2023 and 2024 put the drought there and moved Pacific and Latin American agriculture, fisheries and hydropower over the same months. Yet it moved them at the level of whole countries, and the exporter and importer time effects in equation (12) absorb shocks at exactly that level. The draft coefficient comes from the difference between two pairs in the same month that differ in how much the isthmus saves them. Entering the published ocean index against the same dependence leaves the point where it is, in Table S120, as two measures of one cause should when both enter one regression and compete for the same variation. The index and the shortfall record the same drought, and the canal's version carries the dates and magnitudes a carrier saw. On the warm phase alone, where the lake actually fell, the estimate is -0.084.

Exposure here is geography and the shortfall path is the shock, the case Borusyak and Hull (2023) treat, so I re-randomize the shock path and hold the exposure fixed. Rotating the shortfall through every circular offset of the window, I find that four of seventy-one rotations produce a coefficient as large, and the inference rests on that count. Because the coefficient is identified by one national time path interacted with a cross-sectional distance, an error built on the cell count treats one national path as many independent draws (Conley and Taber, 2011; Ferman and Pinto, 2019). A trend fitted through the treated months of a non-monotone shortfall absorbs part of the treatment itself, so the reading I carry fits the trend on the months before the shortfall alone.

Refitting each pair's line on the months before the shortfall alone and extrapolating it forward, the form that cannot load on treatment by construction, leaves -0.057, and it is significant. The gap to -0.035 measures how much of the treatment a within-window trend absorbs. Every dollar object downstream is therefore priced on the band of readings, not on the point. Because the two shocks share their months, the baseline of Table S121 enters the Red Sea's own exposure beside the draft. The draft term holds at -0.069 with the closure in and the two pair exposures correlate at -0.15, so the overlapping months separate the two shocks instead of confounding them.[27]

The outcome carries the freight inside it, so the bias comes with a sign. Import values are recorded with cost, insurance and freight included for most reporters, so a drought that raises freight on exactly the dependent pairs pads the recorded value of what still arrives while the quantity falls. The United States and the other free-on-board reporters record no such padding, so the bound binds on the CIF subset and is merely absent, not reversed, elsewhere. Ocean freight runs a twentieth of import value on long containerized hauls (Hummels, 2007), so even a shortfall severe enough to raise delivered freight by half would hide two and a half percentage points of the fall, leaving the 3.6 at the deepest cut a floor. Every dollar figure priced with the elasticity inherits that floor. Most of the response to Panama was not rerouting. Two extra transits a day appeared in the Strait of Magellan against eleven a day missing at the canal, so a ship that could not get a slot mostly waited or bid for one. A treatment built on the transit count is attenuated by that whole margin, in Figure S18. The draft is not.[28]

The restriction brings its own falsification, because it did not apply to the whole canal. The lake constrained the draft of the Neopanamax chambers, which carry the container ships and the larger liquid-bulk carriers, while the dry bulk crosses in the older and shallower Panamax chambers whose limit the drought never moved. Sorting goods by the chamber they pass through therefore yields a prediction that turns on lock geometry and not on what the goods are, and the contrast that decides it is containerized against dry bulk. In Table S122 the containerized chapters carry -0.078 and the dry bulk 0.038, with a confidence interval that includes zero. The contrast itself is a difference of 0.115, in Table S123. A stricter chamber mapping there reassigns the two chapters able to travel containerized and moves the containerized cell not at all. The liquid-bulk chapters carry -0.080, in line with the containerized cell. A shock to demand for the pairs that use the isthmus would not sort itself by lock chamber, so this split is also the paper's answer to the El Niño demand story.

The shortfall entered as its five-month mean returns -0.075 against the single month's -0.070, in Table S124, so shipments that merely waited do not drive the estimate. Split at the trough, the falling segment carries -0.073 and the recovering one -0.066, in Table S125, and the two agree, the symmetry that carrying the estimate to a permanent change would require. The other canal's own shock leaves the estimate where it is and takes -0.010 of its own, in Table S126, the same number

[27] The shortfall is continuous and non-monotone, and the rotation reads it whole, comparing the observed path against every rotation of itself with no weighting argument.

[28] The authority's advisories publish the queue itself. That series would price the waiting margin directly, but the advisories circulate as notices rather than as a machine-readable archive, so the draft, the constraint the queue forms on, serves as the treatment at daily resolution.

the Red Sea ladder's upper reading gives when I estimate that term in its own design, so two designs close on one coefficient. Filling the cells with no trade in them leaves the magnitude where it is, in Table S127. The count form the gravity literature prefers when the zeros carry information (Santos Silva and Tenreyro, 2006) holds the magnitude near -0.080 whether the empty cells are kept or dropped. Reading the product dimension twelve times finer returns -0.038 at the heading level in Table S128, and the loss still concentrates in the goods that cross the deep chambers.

The two elasticities go into one money through the hull's immersion relation, which turns a meter of draft into deadweight, and the published cut, 1.83 meters off 15.24, forfeits a sixth of a Neopanamax's cargo, raising the cost per delivered tonne by a fifth while the shortfall binds. Ocean freight is a twentieth of delivered value on these hauls, so the deepest cut raised delivered cost by one percent. The 3.6 percent trade fall it produced is an elasticity near minus four to delivered cost, squarely in the range the gravity literature recovers (Head and Mayer, 2014; Simonovska and Waugh, 2014). On the Red Sea side, priced as fuel alone, a thousand nautical miles is half a percent of delivered value and the -0.010 it produces reads near minus two. With the carrier's own time priced in at published charter rates, the cost doubles and the reading halves. The two shocks thus price trade in the same direction and within an order of each other, and each elasticity prices its own domain, the draft pricing the displaced-trade band and the closure pricing its own row of the ledger.

The authority rations slots and auctions the rest, and the bidding in the driest weeks cleared at a record of 3.975 million dollars for one slot in November 2023, an episode the Hormuz surge of 2026 repeated. From January 2024, by advisory A-01-2024, it set the opening amount at 100,000 dollars for a Neopanamax slot and 110,000 on the busiest day of the week. That record is the one arc price in the paper quoted rather than inferred, a published price for one transit through one arc, posted by the single agent on this map who prices an arc while bearing none of its plume, so the price is set by receipts instead of by damage. It cleared at seven times the posted toll, so at the trough the binding price of the arc was the rationing wedge and not the toll, which is exactly the margin the draft-based treatment reads.

The elasticity and the published draft path give the proportional loss on every pair in every month it bound, and I multiply by what that pair actually traded in the year before the lake fell to turn it into money, with nothing calibrated from outside. A chokepoint that loses a tenth of the draft it can offer for a year and a half displaces 162 billion dollars of trade from the routes that depend on it, in Table S129, the total read at the baseline elasticity. Across the specification readings, from the pair-trend -0.035 to the baseline -0.070, the same arithmetic runs from 83 to 162 billion, and the sampling interval around the baseline runs 43-277 in the same table.[29] That figure is the quantity an authority setting the price of an arc does not have, and it is a gross reallocation, not a welfare loss. The importer by month effects mean the coefficient measures substitution across

[29] Scaling dependence by the distance the canal saves, the deepest cut costs the most dependent pairs a few percent of their trade, while scaling by the share of the route the canal saves implies losses near a tenth for a fully reliant pair. The two scalings agree on which pairs lose and disagree on how steeply loss rises with reliance, and the band is quoted on the flatter, smaller-dollar reading.

suppliers inside an importer, so the figure counts trade moved to a different origin, and the welfare cost of moving it is the wedge and not the flow.

The customs ledger of one importer shows the same displacement directly for the closure. American ports sit on two coasts, and from an origin east of Suez the Atlantic coast is reached through the canal while the Pacific coast never is, so the closure lengthened exactly one of an origin's two legs. In Table S130 the share of each Asian origin's containerized imports that lands on the Atlantic side falls by 2.6 percentage points on a base of 67 where that origin's Atlantic leg runs through Suez, with the window ended before the American tariff round of 2025, the reading the design rests on, and by 3.6 points over the full window. The fall is measured against origins that reach the Atlantic through Panama as easily as through Suez, and the share falls by as much when each origin is weighted by what it ships. Since the comparison stays within Asia, whatever moved all Asian cargo between the American coasts nets out.

The share is flat for the two years before the closure, crosses into negative territory in the closure quarter, and stays negative through the two years after it, deepest in the quarters the American tariff round also moves. That is why the pre-tariff window carries the reading. The reading is unmoved on the window ending before Hormuz and holds at -0.034 with China excluded from the origins entirely, so the tariff round could not produce the fall even by composition. Every ingredient here is mechanical, the origin fixed by geography east of Suez, the two American coasts fixed by which canal reaches them, and the customs record independent of every satellite and transit series in the paper. On the trade side this is the clearest reading of the closure. The lengthened arc moved trade between coasts and did not stop it.

S3 Robustness of the Arc Price Ratio

A sweep of the decay length is not a robustness check here. It is the disagreement between the two instruments written as one parameter. As the length grows a city's weight stops depending on where the lane ran and comes to depend only on how long it was, so the ratio walks to the ratio of the two route lengths, 0.69, the price a charge proportional to fuel sets. Figure S19 draws the whole walk and Table S131 lists its solved points. The ratio is 2.99 at 50 kilometers, 2.59 at the 150 I read everything at, and 2.44 at the 225 the water itself fits. It reaches one just past two thousand kilometers.

The single length also averages over two waters that fit different lengths. On each corridor's exclusive water the Cape profile returns 225 kilometers while the Red Sea's own decline dies within 31. Per transit moved, the fall at the Red Sea lane is eleven times the rise at the Cape lane, which is the confined-basin arithmetic the framework's two-length fit reads in kilometers and this section reads in concentration. Repricing each arc with its own fitted pair returns 2.70 against the common kernel's 2.59, in Table S132, so the single length the paper uses is the smaller of the two readings its own water supports.

Netting out the water both routings cross moves the ratio from 2.59 to 2.41, in Table S133, the routing-margin object Assumption 3 defines. Pricing each lane by its closest point moves the pair to 2.53 and 2.31. The shared water thus accounts for 0.18 of the ratio, not for the ratio itself.[30] One chosen parameter of the solver, the standoff, moves the ratio and moves it in only one direction. Rebuilt at every standoff setting the calibration admits, in Table S134, the ratio runs from 2.11 at no standoff to 4.09 at the heaviest, because pushing carriers offshore empties the open coast's kernel faster than the confined sea's.

The baseline is the exposure a transit delivers, summed over every coastal city with a population count, at the 150 kilometer decay length every coefficient in this paper is estimated at, per transit of a given hull. Everything else in this section departs from it by one named choice. It is 2.59 times the Cape arc. The coefficient that establishes the weight as plume incidence is estimated at 150, so the pricing reads at 150 with the 225 pair printed beside it everywhere, 2.44 per hull and 3.54 per tonne against the headline's 2.59 and 3.73. Collected in Table S53, the variants change one thing each, with Figure S20 drawing the family against the 0.69 a fuel charge sets, every reading on the far side of parity from the instruments in force.

Every base a charge could be written on leaves the same arc dear. Charged per transit at the ship sizes actually observed, with the given-hull clause of Assumption 1 relaxed, the ratio is 1.96, because the hulls that round the Cape are the larger ones. It is 2.16 when emissions scale with deadweight to the 0.65, and 3.73 charged per tonne of fuel, the unit the instruments in force operate in. A charge written per transit would also sort the fleet, pushing small hulls onto the canal and drifting the ratio toward the per-hull 2.59, while the per-tonne base does not move when

[30] The netting is also why the western Mediterranean carries almost no exposure while Casablanca, on the Atlantic, carries a gain, visible on the corridor map. A lane from Asia to Rotterdam passes Valencia whether the canal is open or shut, while the Atlantic approach passes Casablanca only in the closed state. The near-zero shows the measure working as designed, and the netted figure is the one an instrument would act on, reported beside the baseline because the theory is defined on the whole routing.

hulls sort. The unit the instruments already use is therefore also the one robust to the sorting a charge would set off. Moreover, fuel per mile rises with the square of speed, so a per-tonne charge on the confined arc induces slow steaming through exactly the water where exposure is dearest, an abatement channel a per-transit charge does not touch.

Reweighting who is counted leaves the ordering untouched. Letting the response differ by city size leaves the ratio where it stands, letting it differ by density leaves 2.61, and weighting a person by whether they are over sixty-five leaves 2.49, since the two coastlines carry almost the same share of older people. Weighting jointly by age and inverse income leaves 1.62, all in Table S53. One sampled input runs through every variant, the decay length, whose sampling band both headline objects inherit. Scaling each city's weight by the share of hours the wind actually blows from the lane toward it leaves 2.16, so the kernel's isotropy is worth four tenths of the ratio and none of its sign.

Curvature in the damage function moves the ratio in a direction the baselines pin down. The Red Sea coastline is the dirtier baseline, so damage convex in concentration moves the ratio up. A marginal damage inverse to the baseline column relaxes the linearity of Assumption 2 toward a bound rather than a chemistry, since it puts the most weight on the cleanest air the reallocation touches. Under it the ratio is 1.76, the worst case the saturating class can produce. The toll-setter's own nationals bear 9.2 percent of the Red Sea arc's load, and netting them out leaves 2.35.

Nothing in the sample or the solver moves the ordering either. Dropping any single lane of the 140 moves the ratio between 2.55 and 2.63, and dropping any single country's cities moves it between 2.35, without Egypt, and 3.02, without South Africa. Every reading leaves the same side dear, and the pair is an extremal one by construction, the narrowest heavily used sea on the map against the most open water. The one construction input that moves the ratio is the standoff. Over the settings the gate rule's algorithm admits, a rule that never sees the ratio, the ratio runs from 2.11 to 3.03, and the 2.59 I carry sits near the conservative end. Re-solving every lane on a sixteen-neighbor grid that admits the knight's moves leaves the city weights correlated at 0.996. In Table S52 the constructions of the measure span 2.11-3.45, and one departure at a time across its constructions, bases and curvatures runs 1.76-3.73. No reading in the table crosses one.

The water itself shows the confined-basin asymmetry in its own units. The column over the corridor the ships left fell by 0.142 micromoles against the far field, while the corridor they moved to carries -0.016, well inside its own error. Emission spread over open ocean at forty degrees south produces a signal that small. The counting box at the Cape rises with the traffic, the corridor stands 0.405 above the far field the decay is fitted against, and the same corridor read raw moves by less than its own error, three readings of the same water that appear to disagree and in fact measure the same change. That is because the raw level carries a field-wide fall of four tenths of a micromole over the window. Put on one baseline the three agree, and the reason the paper differences within country and month everywhere is that a raw level over this window mixes the corridor's change with that field-wide drift.

The Red Sea's confined air clears exhaust far more slowly than the windswept ocean off the Cape, since concentration is emission over the volume of air that carries it away. So an instrument

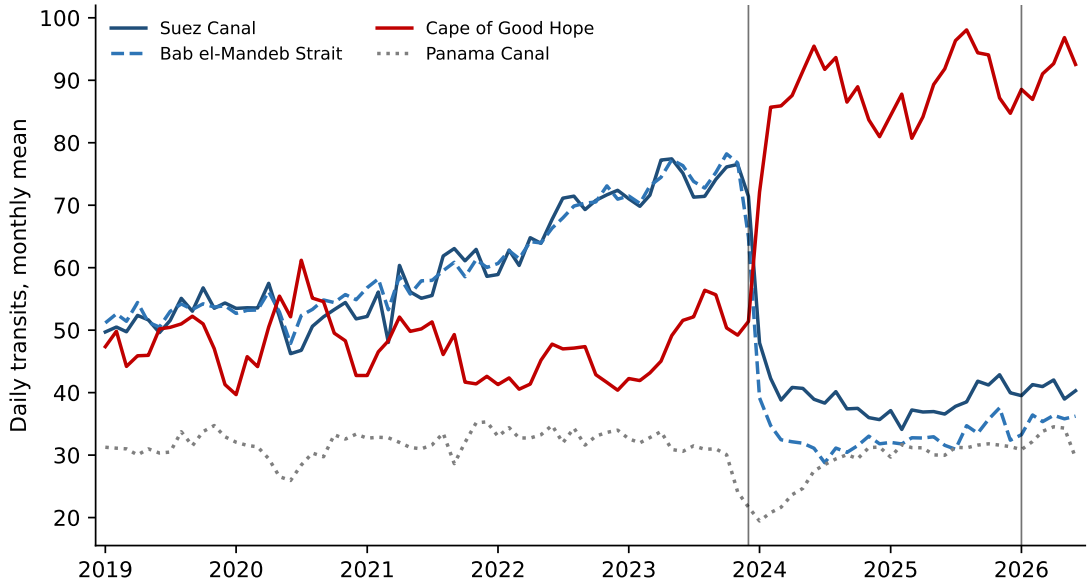
that priced the air's failure to clear would move the ratio further in the direction this one already points. To price that direction, I replace the kernel with a plume model in the ERA5 fields of the closure year. Every point of both corridors spreads through its own month-by-month wind rose, diluted by the boundary layer and the wind, thinning geometrically, decaying over hours of travel instead of kilometers of map. The air along the Red Sea routing moved at 5.7 meters a second under a 620 meter layer against 7.1 under 770 along the Cape, and the model returns 2.74 against the kernel's 2.59 on the same domain and cities, in Table S52. Letting the damage ride that wind for 24 hours before it lands, falling on the people under the day-downwind footprint rather than beside the lane, the model returns 2.30, and the ordering survives the displacement as well.

Against the model that lands the plume beside the lane, the kernel is the conservative reading of its own physics, and Table S135 stacks the margins in both directions. Stacked one way, own fitted kernels at 2.70, plume transport at 2.74 and the ammonia geography heavier still carry a reader toward three. Stacked the other way, the wind weighting, the netted water and the saturating damage bound jointly leave 1.17 per transit of a given hull and 1.70 per tonne of fuel. On the base the instruments in force are written in, every stack of named margins, favorable or unfavorable, leaves the same arc dear. It is one more reason the per-tonne unit leads Table 10.

S4 Additional Figures

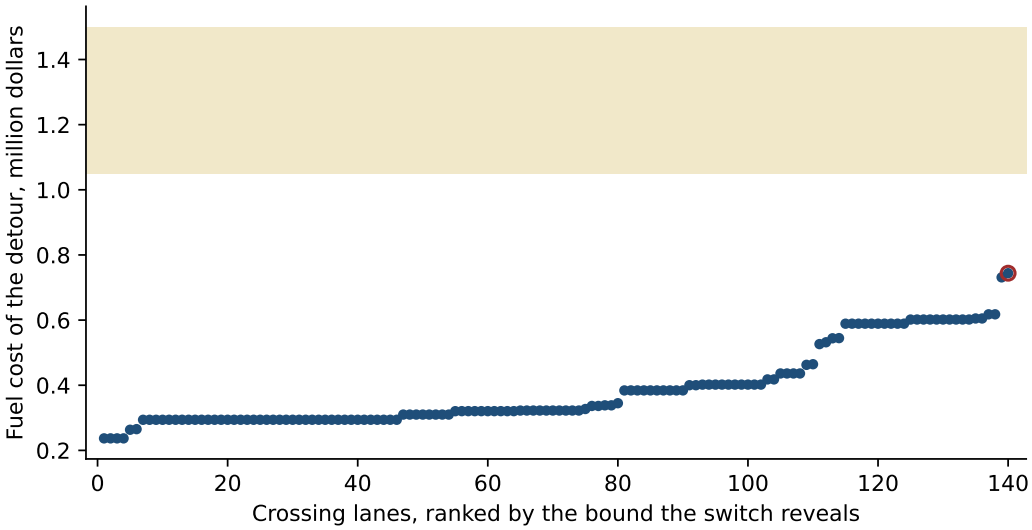
S4.1 Data and the Solved Corridors

Figure S1: Daily transits at the chokepoints the design turns on



Notes. Monthly means of daily transit counts, January 2019 to July 2026, the daily series running twenty-six days past the monthly panel's end. The rules mark December 2023, when the corridor closed, and January 2026, when carriers began returning. Coverage is the 28 chokepoints monitored, of which four are drawn. Unit of analysis is the chokepoint by month. The quantities are enumerated rather than sampled, so no standard errors are carried.

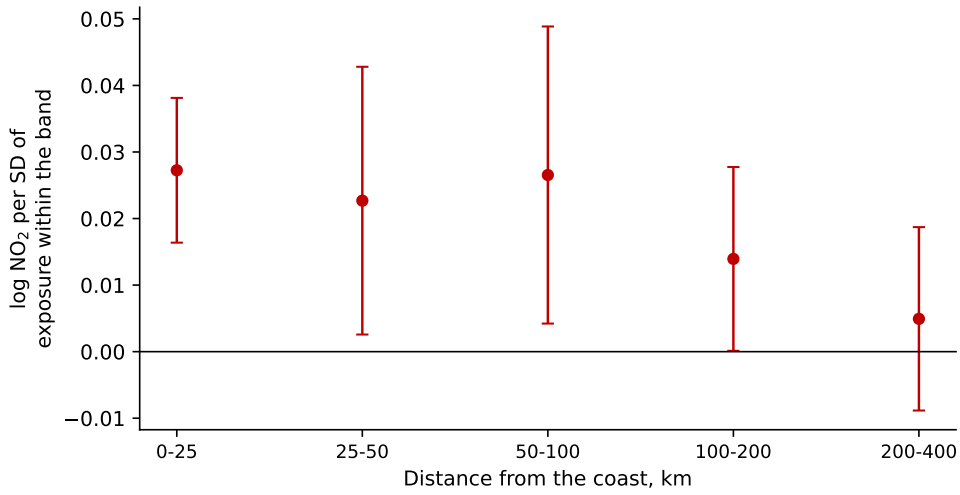
Figure S2: The bound every switching lane reveals, against the premium



Notes. Each point is one of the 140 crossing lanes, ranked by the fuel cost of its own detour around the Cape, 611 tonnes over a 4,133 mile average detour at 600 dollars a tonne. The shaded band is the war-risk premium from December 2023, seven tenths to one percent of a 150 million dollar hull, and the circled point is the Jebel Ali to Ambarli lane. Coverage is every crossing lane solved in both states. Unit of observation is the lane, and the bounds are computed rather than estimated, so no standard errors are carried.

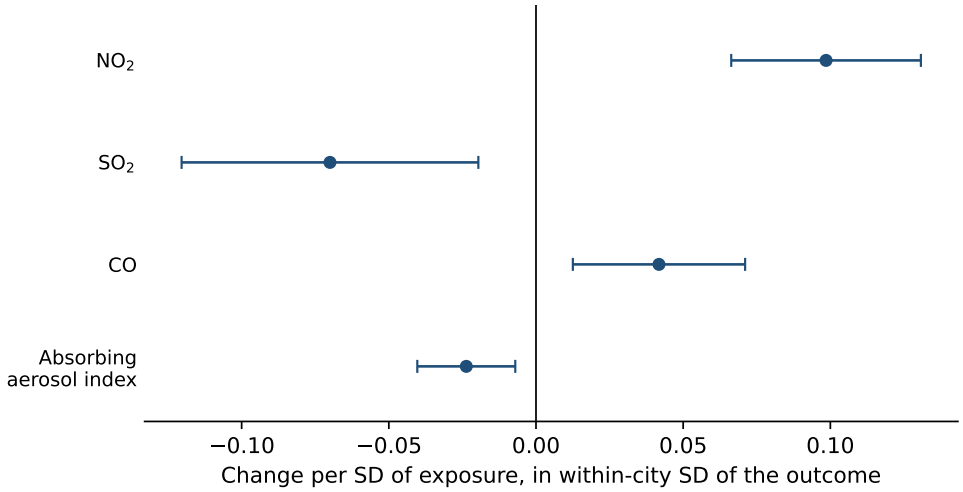
S4.2 The Diversion and Coastal Air

Figure S3: The response by distance from the coast



Notes. Each coefficient from the distance-band regression is rescaled by the exposure spread the band itself has, because a coastal standard deviation lies far outside the support of an inland band. Egypt is excluded. Coverage is cities within 400 km of the coast, which is where the bands stop. Unit of analysis is the city by month. Standard errors are clustered on the city.

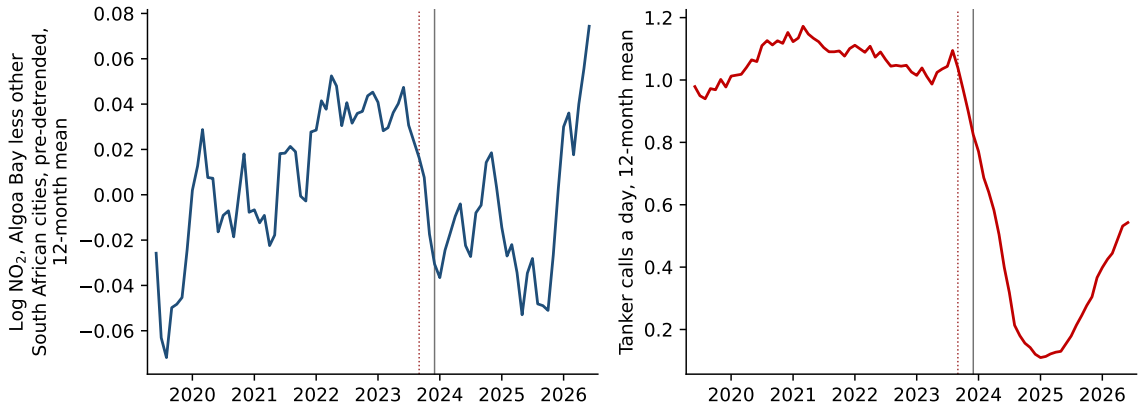
Figure S4: The pollutant signature, on one scale



Notes. Each pollutant is divided by its own within-city standard deviation so the four are on one scale. Bars are 95 percent intervals. Coverage, unit of analysis and clustering are as in Table 3. Ship exhaust after the 2020 sulfur cap is high in nitrogen oxides and low in everything else, while biomass smoke is high in carbon monoxide and in the absorbing aerosol index, and heavy fuel burned on land carries sulfur dioxide with it, which is the chemistry the text reads the four species against. The sulfur coefficient is read as a sign, and the product flags volcanic months.

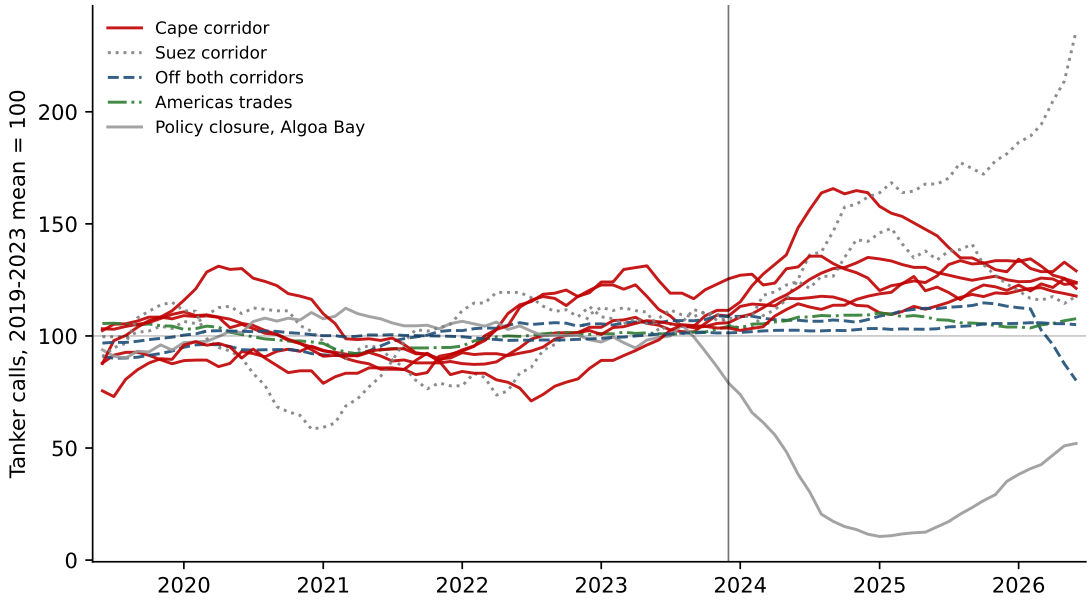
S4.3 The Plume against the Port

Figure S5: Algoa Bay, where bunkering stopped and the plume arrived



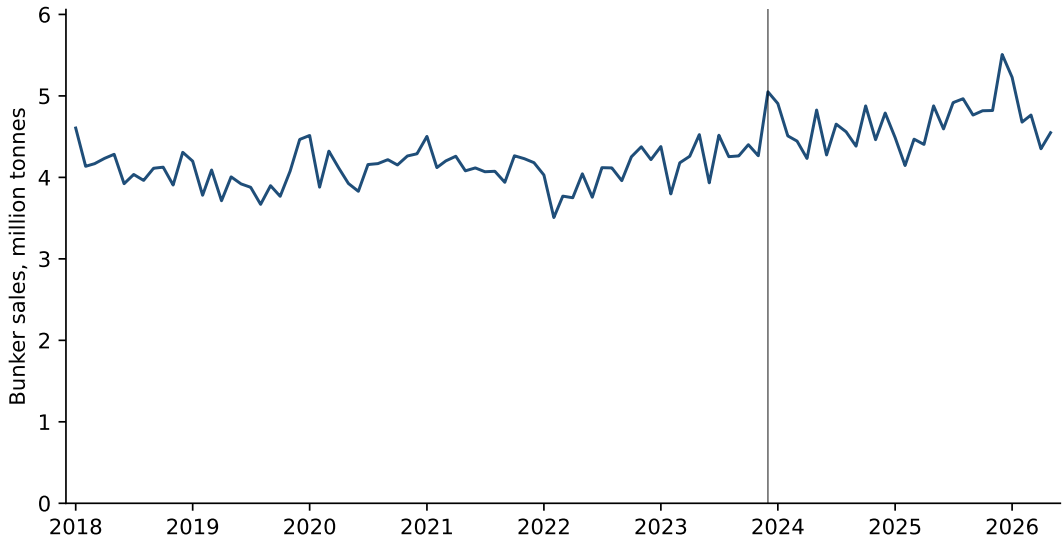
Notes. The left panel is the twelve-month mean log nitrogen dioxide column over the Algoa Bay box less the mean over other South African coastal cities, net of that difference's pre-period linear trend. The right panel is the twelve-month mean of daily tanker calls at Port Elizabeth. The dotted line marks September 2023, when bunkering stopped, and the solid line December 2023. Coverage is one port and its city box over January 2019 to June 2026. Unit of observation is the month. The series are drawn rather than estimated, so no standard errors are carried.

Figure S6: Tanker calls at the ports the diversion was expected to move



Notes. Twelve-month moving means of daily tanker calls, each port indexed to its own 2019-2023 average. Coverage is eleven ports, the Cape corridor bunkering ports, the two Suez-side ports, Singapore and Fujairah off both corridors, and Houston on the Americas trades the closure never touched. Unit of analysis is the port by month. The Algoa Bay series is a policy rather than a corridor and is drawn apart from the rest. The quantities are enumerated rather than sampled, so no standard errors are carried.

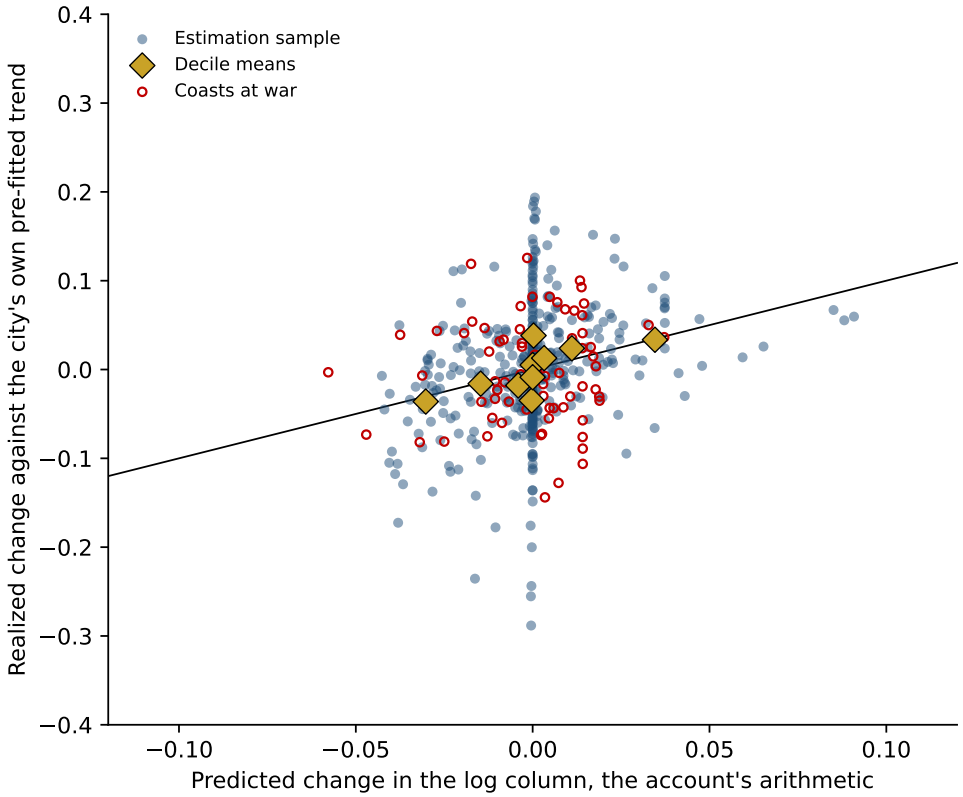
Figure S7: Bunker fuel sold at Singapore



Notes. Monthly bunker sales in million tonnes, from the Maritime and Port Authority of Singapore through the national statistics service. The line marks December 2023. Coverage is January 2018 to the most recent month the authority publishes, and the unit of observation is the month. Nothing in this series is derived from satellite data or from the transit counts, and Singapore sits on neither corridor. The series is enumerated rather than sampled, so the figure carries no standard errors.

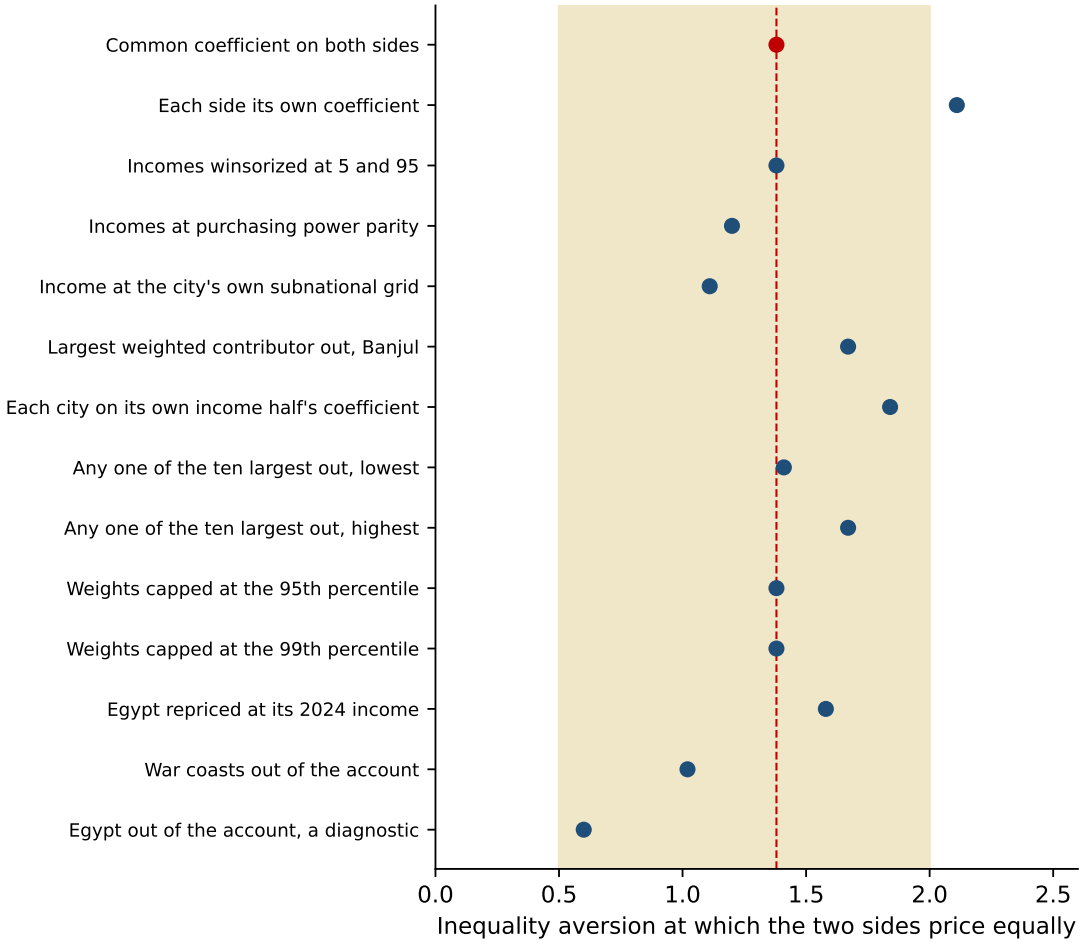
S4.4 Incidence

Figure S8: The predicted change against the observed change, city by city



Notes. Each point is one coastal city. The vertical axis is the post-period mean of the log column against the city's own trend fitted on the pre-period, the horizontal the coefficient times the city's standardized exposure change, both demeaned within country. The line is the diagonal, the diamonds are decile means of the estimation sample, and the open markers are the coasts at war. Coverage is every coastal city with two pre-years and six post months of retrievals. Unit of analysis is the city.

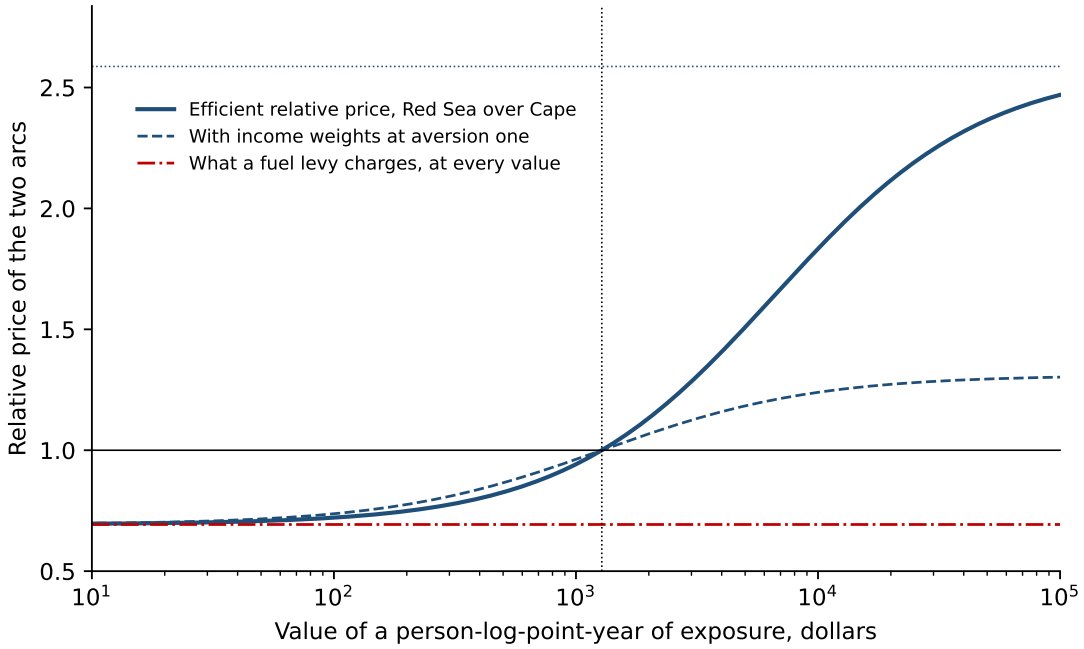
Figure S9: The crossing family on one axis



Notes. Every crossing variant of Table S47 with the war-coasts-out reading of Table S48 and the Egypt-out diagnostic of Table S49 appended, read straight from those tables. The shaded band is inequality aversion from one half to two, the range the distributional literature treats as standard, and the dashed line is the baseline crossing. The two lowest rows delete populations from the account and are diagnostics rather than candidate readings. The quantities are enumerated rather than sampled, so no standard errors are carried.

S4.5 The Arc Price

Figure S10: The relative price of the two arcs, as a function of what exposure is worth



Notes. The total efficient charge is θ times person-weighted exposure delivered plus 190 dollars a tonne times carbon burned. The solid curve is the ratio of the two charges as θ varies and the dashed weights each person by inverse relative income at an aversion of one. The dash-dotted horizontal is the ratio of fuel burned, the dotted horizontal above it the exposure ratio the curve approaches, and the dotted vertical the indifference value of Table 9. Coverage is the 507 coastal cities of the incidence account. Unit of observation is the value of exposure. The coastline is enumerated rather than sampled, so the curve carries no standard errors.

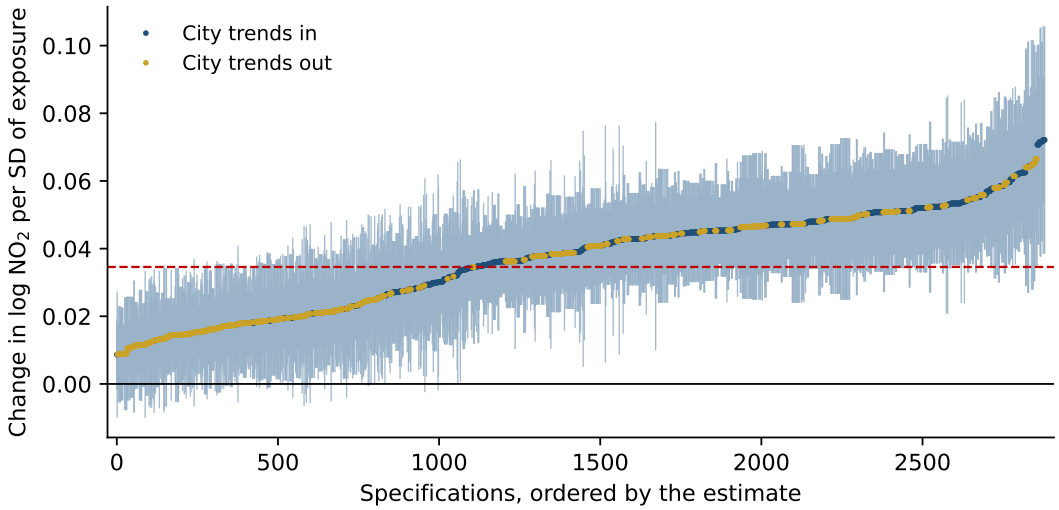
Figure S11: The exposure charge at every passage on the map



Notes. One disc per passage, its area the exposure a transit there delivers under the kernel of Table 11. Discs drawn in outline are the two passages at which a toll beyond cost-recovery service dues is actually collected. Straits regimes such as the Turkish Straits collect charges for specific services rendered, which the law permits and which price no exposure. Coverage is the 28 passages the transit product reports. Unit of analysis is the passage. The centers and the passages are enumerated rather than sampled, so the map carries no standard errors.

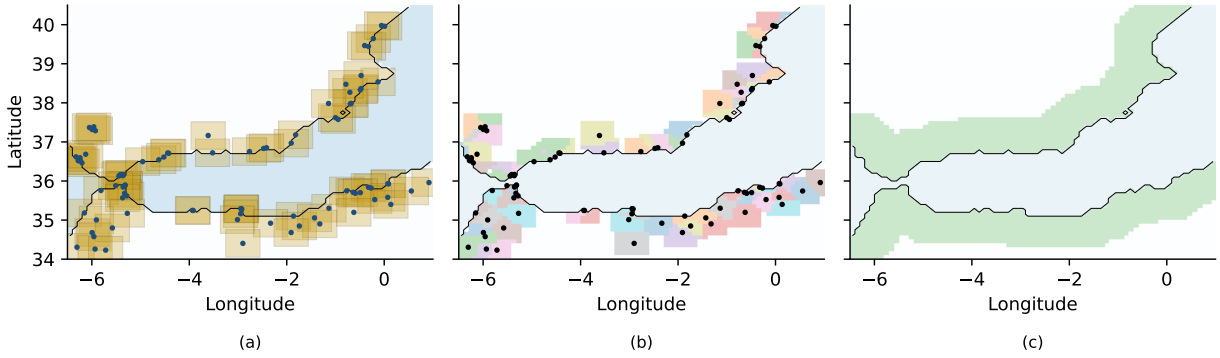
S4.6 Specification of the Estimating Equation

Figure S12: The estimate under every defensible specification



Notes. Each point is one combination of a decay length among 75, 100, 150, 200 and 300 km, a coastal band among 50, 100 and 200 km, Egypt in or out, endpoint ports ranked on container calls or on calls of every hull class, the shock as a switch or as the observed share of the corridor gone, errors clustered on the city or on the country, the city trend in or out, the post period opened in November, December or January, and population weights on or off. Bars are 95 percent intervals. The dashed line is the estimate the paper reports. Coverage and unit of analysis are as in Table 3, except where the specification changes them.

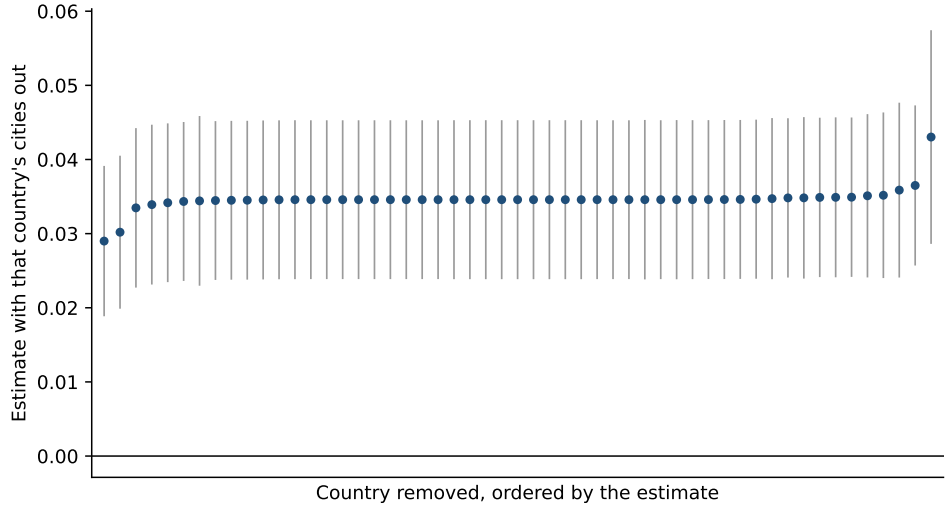
Figure S13: One coast under the three population accountings



Notes. The left panel draws the retrieval box around each city center, 0.7 degrees on a side, with shading deepening where boxes overlap. The middle panel is the exclusive assignment, each land cell given to its nearest city center, the account every population in the paper carries. The right panel is the gridded strip within 100 kilometers of the water on the solver's grid, the account of Table S41. Coverage is the cities of the register inside the drawn window. The panels are drawn rather than estimated, so no standard errors are carried.

S4.7 Named Confounders

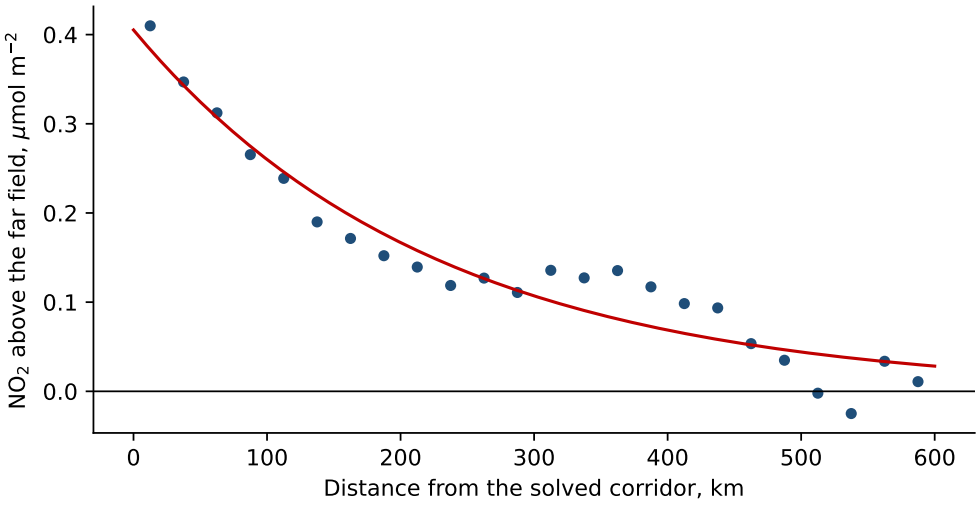
Figure S14: Every coastline removed in turn



Notes. Each point is equation (8) estimated with one country's coastal cities removed, ordered by the estimate. Bars are 95 percent intervals. Coverage is the 53 countries of the estimation sample. Unit of analysis is the city by month. Standard errors are clustered on the city.

S4.8 Placebos and the Measurement of Exposure

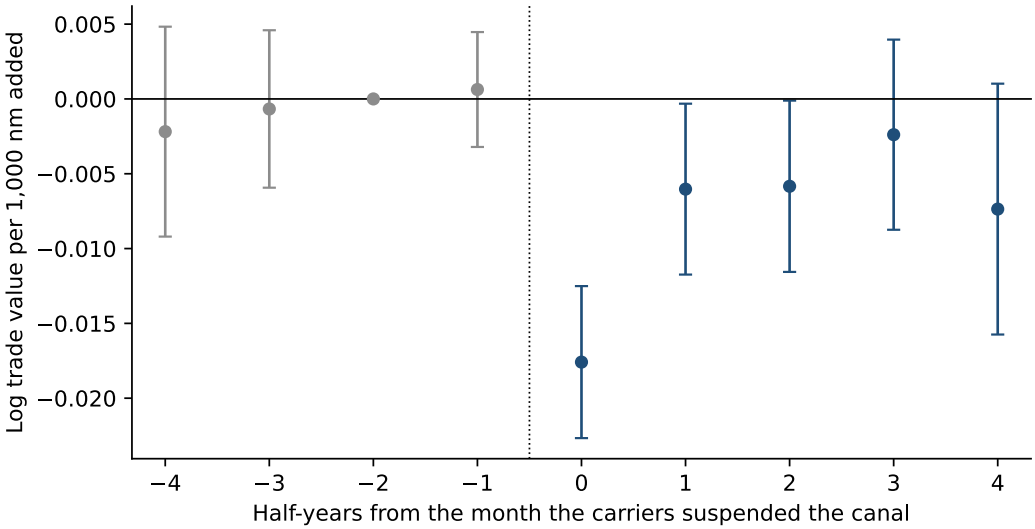
Figure S15: The plume over the water, against distance from the corridor



Notes. Points are means of the change in the column within each 25 km band of distance from the solved path, drawn above the fitted far-field level. The line is the fitted exponential. Coverage and unit of observation are as in Table S111, which also carries the standard error on the fitted decay length.

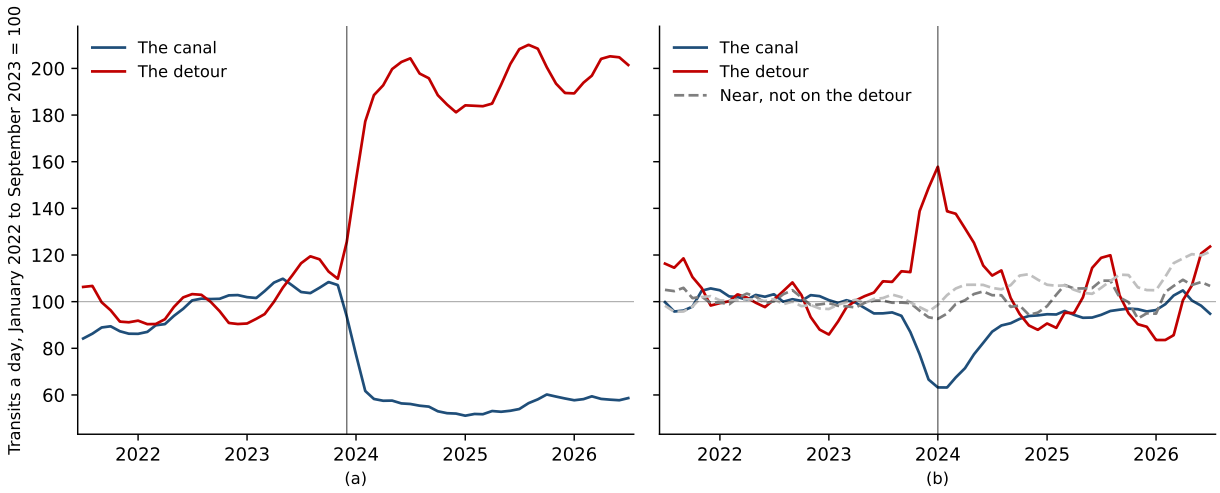
S4.9 The Trade Cost of a Chokepoint

Figure S16: The trade response by half-year around the diversion



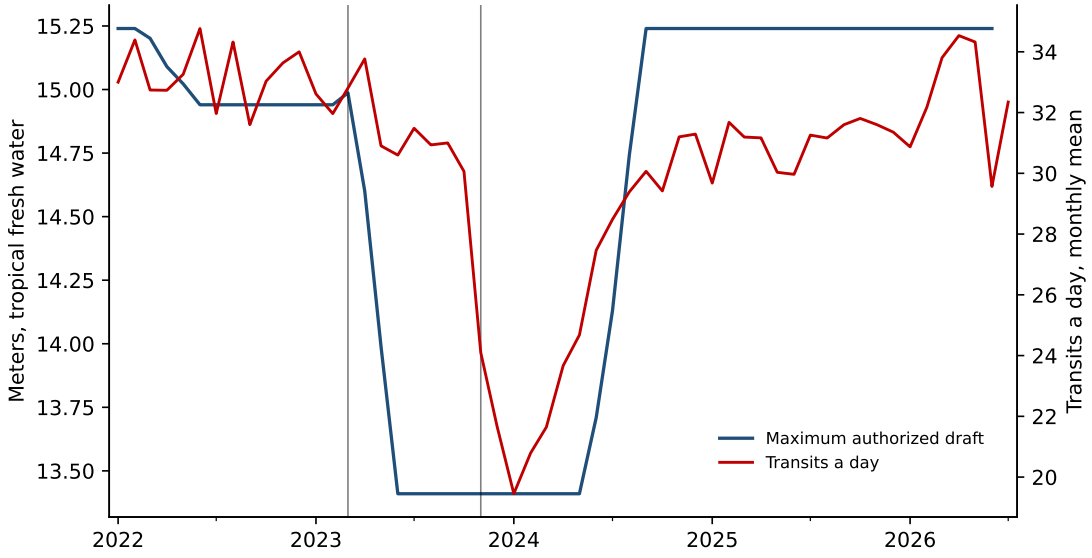
Notes. Each point is a pair's added solved distance in thousands of nautical miles interacted with a half-year of event time, on the global bilateral chapter-month panel with pair-chapter, exporter-month, importer-month and chapter-month effects. The base is the half-year ending in May 2023, and the half-year holding the November attacks is estimated. The pandemic years are removed, and the first bin pools everything before it. Bars are 95 percent intervals. Unit of analysis is the importer by exporter by chapter by month. Standard errors are clustered at the pair level.

Figure S17: Two canals, two detours, and the passages in between



Notes. Monthly means of daily transit counts, indexed to each chokepoint's own mean over January 2022 to September 2023 and smoothed with a centered three-month mean, since the Strait of Magellan carries four transits a day against seventy at Suez. The rules mark December 2023 and January 2024. Panel (a) is Suez against the Cape of Good Hope and Panel (b) is Panama against the Strait of Magellan, with the Windward Passage and the Yucatan Channel drawn beside it. Coverage is the 28 chokepoints monitored, of which six are drawn. The quantities are enumerated rather than sampled, so no standard errors are carried. Unit of analysis is the chokepoint by month.

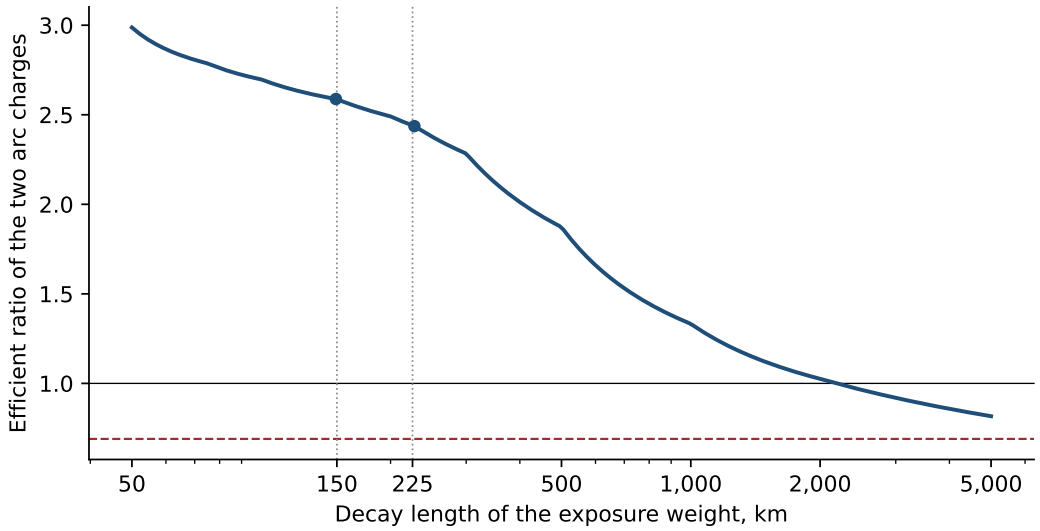
Figure S18: The authorized draft at Panama against monthly transits



Notes. Maximum authorized draft in the Neopanamax locks, read from the Canal Authority's numbered advisories and carried forward from each effective date, against monthly mean daily transits. The rules mark March 2023, when the drought cuts began, and November 2023, when the transit count first moved. Coverage is January 2022 to June 2026. Unit of observation is the month. The advisories are enumerated rather than sampled, so the series carries no standard errors.

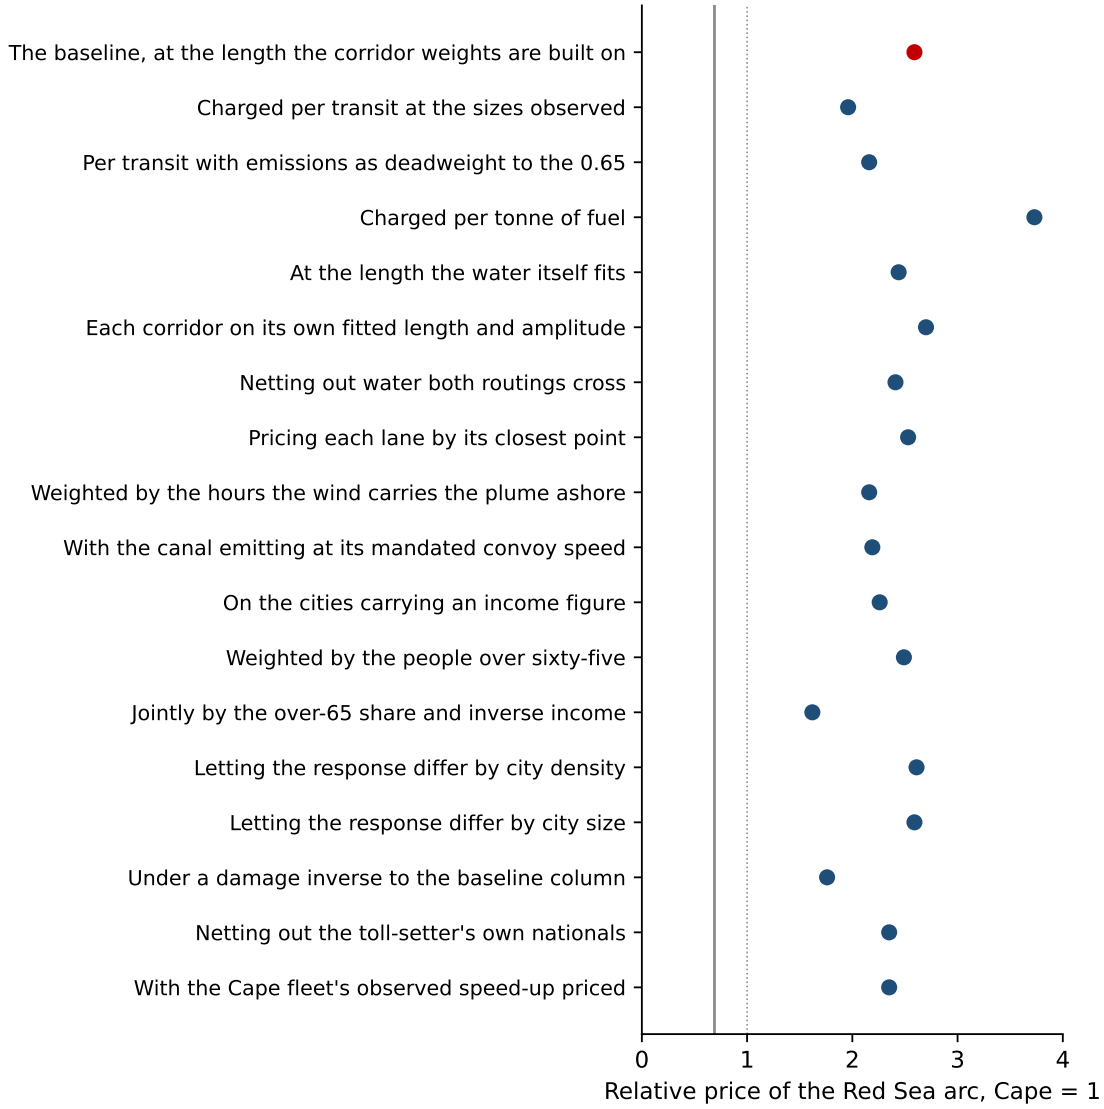
S4.10 Robustness of the Ratio

Figure S19: The efficient exposure ratio at every decay length



Notes. The curve prices the two arcs at every decay length of the exposure weight, interpolated between the eleven solved scales, with the population account held fixed. The dotted verticals mark 150 kilometers, the length every coefficient in the paper is estimated at, and 225, the length the water itself fits. The dashed horizontal is the ratio of the two route lengths, which is the price a fuel-scaled charge sets and the curve's limit as the decay length grows without bound. The thin solid line marks one, where the two arcs price equally. Coverage is the 507 coastal cities of the incidence account. The curve is computed rather than estimated, so no standard errors are carried.

Figure S20: The ratio family on one axis



Notes. Every named variant of Table S53, read straight from that table. The solid vertical line is the 0.69 a fuel-scaled charge sets, the relative price every instrument now in force implies, and the dotted line is parity. The gap between the cloud and the solid line is the disagreement the paper prices. The quantities are enumerated rather than sampled, so no standard errors are carried.

S5 Additional Tables

S5.1 Data and the Solved Corridors

Table S1: The ports the lanes run between

Side	Port	Container calls a day
East of the canal	Singapore	37.8
	Hong Kong	36.7
	Busan	33.1
	Shanghai	31.7
	Ningbo	26.3
	Kaohsiung	21.1
	Qingdao	18.9
	Port Klang	18.6
	Xiamen	14.3
	Ho Chi Minh City	14.1
	Guangzhou (Nansha)	14.0
	Yangshan	11.9
	Jebel Ali	11.8
	Tanjung Priok	11.4
West of the canal	Rotterdam	19.8
	Antwerp	12.6
	Hamburg	9.8
	Piraeus	8.1
	Algeciras	8.0
	Valencia	7.9
	Tanger Med	7.8
	Bremerhaven	7.2
	Ambarli	6.6
	Barcelona	6.1

Notes. Container calls a day are the mean over January 2019 to November 2023, from the transit data. The fourteen busiest east of the canal and the ten busiest west of it give 140 crossing lanes, and the six busiest on each side, paired within their own side, give the thirty lanes of the control corridor. The western set carries no Levantine or Egyptian hub, and the endpoint variants of Table S106 widen the set. Lane weights are products of container calls, which stand in for the liner strings that call at several ports. Coverage is the endpoint ports of the crossing lanes. Unit of observation is the port. The list is descriptive and carries no standard errors.

Table S2: The route solver against the points at which transits are counted

Stand-off		Distance to the counting point			Asia to North Europe			All lanes	
Penalty	km	Cape	Bab el-Mandeb	Suez	Open	Closed	Extra	Extra	Used
0.0	30	11	13	10	8,593	12,162	3,570	4,170	
0.5	30	26	7	10	8,653	12,235	3,581	4,166	
1.0	30	47	4	10	8,748	12,295	3,547	4,133	Yes
2.0	30	70	4	10	8,817	12,356	3,540	4,126	
3.0	30	80	13	10	8,849	12,370	3,522	4,119	
0.5	60	70	7	10	8,713	12,308	3,595	4,193	
1.0	60	114	12	10	8,827	12,409	3,582	4,170	
2.0	60	158	4	10	8,991	12,475	3,484	4,106	
3.0	60	180	13	10	9,078	12,501	3,423	4,180	

Notes. Each row solves both states at one setting of the stand-off penalty. Distances are from the solved path to the point at which a transit is counted. The traffic-weighted extra distance is the quantity equation (2) uses. Coverage is every crossing lane, read at the Singapore to Rotterdam lane in the middle columns and averaged over all of them, weighted by container traffic, in the last. Unit of analysis is the solved path. The solver is deterministic, so the table carries no standard errors.

Table S3: Every panel the paper estimates on

Panel	One row is	Rows	Units	Window	Selection rule
The city panel	City by month	38,970	433	2019-01 to 2026-06	Urban centres and ports within 100 km of the coast, outside the war coasts
The event window	City by month	28,578	433	2019-01 to 2026-06	The city panel with the two pandemic years out
The satellite-pass panel	City by overpass	349,960	120	2019-01 to 2026-06	One row per city per cloud-free overpass, with the wind at that hour
The Panama panel	Pair by chapter by month	17,544,694	5,869	2019-01 to 2024-12	Country pairs whose least-cost path uses the canal, by trade chapter
The corridor water panel	Corridor by month	180	2	2019-01 to 2026-06	The solved lanes' own water, on the retrieval grid
The chokepoint panel, monthly	Passage by month	705	8	2019-01 to 2026-06	The passages the transit product covers
The Suez trade panel	Pair by chapter by month	6,727,481	2,702	2019-01 to 2026-05	Customs-reported trade for the reporters publishing monthly
The customs-district panel	District by origin group by month	263,694	23	2019-01 to 2026-05	United States customs districts, by their import origin mix
The ground-station panel	City by month	1,168	14	2019-01 to 2026-06	European monitoring stations, averaged to the city they sit in

Notes. Counts are read off the processed file each panel is estimated on, before the fixed effects drop singletons, so a table's own reported sample is at or below the row here. The ground-station panel holds every monitored coastal city the European download returns, and nine of the fourteen match a city of the satellite sample, which is the nine the validation of Table S90 runs on. Coverage is every panel the paper estimates on. Unit of analysis differs by row and is named in the second column. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S4: The countries the corridors run past

Country	Cities	People, m	Low	High	In	Country	Cities	People, m	Low	High	In
Nigeria	30	35.5	0.00	0.00	Yes	Qatar	3	2.1	0.00	0.00	Yes
Egypt	20	21.4	-2.08	-0.77	Yes	Guinea	5	2.1	0.38	0.64	Yes
Angola	12	15.9	0.00	0.01	Yes	Bahrain	1	2.0	0.00	0.00	Yes
Algeria	28	15.8	-0.96	-0.47	Yes	Iraq	3	1.9	0.00	0.00	Yes
Syria	9	12.1	-0.27	-0.05	No	Liberia	4	1.7	0.12	0.52	Yes
Morocco	25	11.1	-0.48	0.78	Yes	Greece	14	1.6	-1.28	-0.07	Yes
Spain	25	10.8	-0.68	1.65	Yes	Jordan	2	1.5	-1.07	-0.34	Yes
Turkey	31	10.7	-0.47	-0.04	Yes	The Gambia	1	1.3	1.25	1.25	Yes
South Africa	17	10.6	0.02	1.76	Yes	Mauritania	2	1.3	0.93	1.57	Yes
Yemen	19	9.2	-2.07	-0.73	No	Oman	6	1.2	-0.13	0.00	Yes
Ghana	10	7.4	0.00	0.00	Yes	Kuwait	2	1.0	0.00	0.00	Yes
Saudi Arabia	16	7.3	-1.53	0.00	Yes	Eritrea	5	1.0	-1.93	-0.66	No
Tanzania	5	7.1	0.00	0.00	Yes	Mauritius	1	0.9	0.00	0.00	Yes
Mozambique	20	6.8	0.00	0.03	Yes	Gabon	2	0.9	0.00	0.00	Yes
Côte d'Ivoire	8	6.3	0.00	0.06	Yes	Guinea-Bissau	3	0.8	0.41	0.85	Yes
Tunisia	15	5.9	-1.25	-0.21	Yes	Sudan	3	0.8	-1.18	-0.81	No
United Arab Emirates	9	4.3	0.00	0.00	Yes	Democratic Republic of the Congo	4	0.7	0.00	0.00	Yes
Senegal	13	4.3	0.48	1.50	Yes	Cyprus	4	0.7	-0.61	-0.33	Yes
Cameroon	5	4.2	0.00	0.00	Yes	Comoros	3	0.7	0.02	0.04	Yes
Palestine	3	4.1	-0.95	-0.48	No	Djibouti	2	0.7	-1.26	-0.87	No
Benin	4	3.8	0.00	0.00	Yes	Ethiopia	1	0.5	-0.68	-0.68	Yes
Italy	11	3.6	-0.99	-0.19	Yes	Réunion	1	0.4	0.00	0.00	Yes
Lebanon	2	3.5	-0.29	-0.17	No	Cabo Verde	5	0.4	0.04	0.15	Yes
Iran	19	3.4	0.00	0.00	Yes	Malta	1	0.4	-1.30	-1.30	Yes
Libya	13	3.2	-1.02	-0.17	No	Mayotte	1	0.2	0.06	0.06	Yes
Israel	4	3.1	-0.70	-0.35	No	Equatorial Guinea	2	0.2	0.00	0.00	Yes
Togo	3	3.1	0.00	0.00	Yes	Sao Tome and Principe	1	0.2	0.00	0.00	Yes
Somalia	14	2.7	-1.39	0.00	No	Pakistan	2	0.2	0.00	0.00	Yes
Portugal	8	2.6	-0.68	0.75	Yes	Western Sahara	2	0.1	1.30	1.52	Yes
Sierra Leone	5	2.5	0.38	0.77	Yes	Namibia	2	0.1	0.09	0.36	Yes
Madagascar	11	2.4	0.00	0.04	Yes	Seychelles	1	0.1	0.01	0.01	Yes
Kenya	4	2.3	0.00	0.00	Yes						
Total	507	275.0									

Notes. Every country holding a city within 100 km of the coast in the window the corridors cross, with the number of such cities, the people the gridded count puts inside their boxes, the range of the exposure measure across them, and whether the country enters the regressions. The ten coasts on which a war was fought over the same window are counted in the incidence account and excluded from every regression, which removes 74 cities. Coverage is every country holding a city within 100 kilometers of the coast in the window the corridors cross. Unit of analysis is the country. The list is descriptive and carries no standard errors.

Table S5: The retrieval's own processing against the closure

Version	In force from	Against the closure
1.0.0	2019-01	ends before it
1.3.0	2019-04	ends before it
1.4.0	2020-12	ends before it
1.5.0	2021-08	ends before it
1.6.0	2022-08	spans it
1.7.0	2024-12	begins after it
1.8.0	2025-12	begins after it

Notes. The algorithm version carried by one nitrogen dioxide scene a month, read from the mission archive, and the month each version first appears. A reprocessing step changes the level of a retrieval and would enter a before-and-after design as a break, so the table dates every step against December 2023. Coverage is January 2019 to June 2026 of the offline product. Unit of observation is the month. The versions are read rather than estimated, so the table carries no standard errors.

S5.2 The Diversion and Coastal Air

Table S6: Seasonality, a trend, and the window that excludes the pandemic years

	(1)	(2)	(3)	(4)
log ships counted	0.495***	0.413***	0.310***	0.389***
	(0.073)	(0.052)	(0.090)	(0.126)
Month of year	No	Yes	Yes	Yes
Linear trend	No	No	Yes	Yes
Window	2019-2026	2019-2026	2019-2026	2021-2026
Mean of dep. var.	-11.624	-11.624	-11.624	-11.582
Observations	90	90	90	66
Adjusted R^2	0.328	0.632	0.637	0.616

Notes. Coverage, unit of analysis and standard errors are as in Table 2, except in Column (4), which starts the window in January 2021 and so drops the pandemic months the other columns keep.

Table S7: The regression at successive distances from the coast

	(1) All water	(2) Beyond 25 km	(3) Beyond 50 km	(4) Beyond 75 km
log ships counted	0.413***	0.412***	0.391***	0.372***
	(0.052)	(0.055)	(0.060)	(0.065)
Share of box averaged	69.8%	58.3%	41.4%	24.5%
Mean of dep. var.	-11.624	-11.644	-11.696	-11.748
Observations	90	90	90	90
Adjusted R^2	0.632	0.637	0.617	0.577

Notes. Each column averages a smaller and more distant part of the same box, first all of its water and then only water beyond 25, 50 and 75 kilometers from the nearest land. Coverage, unit of analysis and standard errors are as in Table 2.

Table S8: The hull classes against one column

Panel A

	(1) Container	(2) Tanker	(3) Dry bulk
log container ships	0.207***		
	(0.027)		
log tankers		0.426***	
		(0.058)	
log dry bulk			0.531***
			(0.070)
Change in transits	254.7%	63.4%	53.5%
Mean of dep. var.	-11.624	-11.624	-11.624
Observations	90	90	90
Adjusted R^2	0.635	0.589	0.614

Panel B

	(1)	(2)	(3)
log ships counted	0.413***		
	(0.052)		
log capacity crossing		0.396***	
		(0.052)	
log container capacity			0.138***
			(0.018)
Traffic measured as	Hulls counted	Capacity crossing	Container capacity
Mean of dep. var.	-11.624	-11.624	-11.624
Observations	90	90	90
Adjusted R^2	0.632	0.633	0.638

Notes. Coverage, unit of analysis and standard errors are as in Table 2. The three elasticities of the upper panel each fit the same rise in the column, so that panel is one relationship read across the hull mix rather than three identifications. Exhaust follows the fuel a ship burns rather than the number of hulls, which is why the lower panel enters capacity in place of the count.

Table S9: The passages whose traffic did not move

	(1) Cape of Good Hope	(2) Windward Passage	(3) Mona Passage	(4) Yucatan Channel
log ships counted	0.413***	-0.131	0.099	-0.180
	(0.052)	(0.169)	(0.186)	(0.232)
Change in transits	84.6%	-2.6%	4.3%	10.3%
Mean of dep. var.	-11.624	-11.417	-11.583	-11.489
Observations	90	90	90	90
Adjusted R^2	0.632	0.227	0.057	0.399

Notes. The three Caribbean passages are open water and are monitored by the same counting system as the Cape. Coverage, unit of analysis and standard errors are as in Table 2. The smallest slopes each column could have detected at eighty percent power are in Table S95. The transit change is read off the monthly series weighted by the days in each month, against the same window before December 2023, which differs from the daily series of Table S11 only in the day weighting.

Table S10: The change over the water on each corridor

	Change, μmol m^{-2}	Standard error	Blocks
The corridor the ships left	-0.142	0.079	242
The corridor they moved to	-0.016	0.022	698

Notes. Coverage is the ocean cells within 150 km of each solved corridor. The change is the one drawn in Figure 3, averaged over ocean cells within 150 km of one corridor's solved paths and more than 150 km from the other's, so the water both routings cross belongs to neither. The standard error of each mean is taken over 125 km spatial blocks. Unit of observation is the 25 km ocean cell.

Table S11: The solved detour at the gates it was never fit to

	Change, percent
Bab el-Mandeb, on the closed arc	-43.6
Suez, on the closed arc	-33.1
Cape of Good Hope, on the detour	84.9
Sunda, the entry the solved detour substitutes toward	19.2
Median absolute change at the 20 untouched passages	5.0
the same median with its sign	-2.6

Notes. Daily transits from December 2023 against the five years before, per passage. The solver was calibrated to put the paths through the three counting points and nothing else, so the Sunda row and the untouched rows are out of sample. The untouched set is the passages the solved lanes cross in neither state, which leaves out Malacca and Gibraltar, the doors the open-state lanes themselves use, and it also excludes Kerch, whose collapse is the Black Sea war, and Panama, whose drought is the second case. Coverage is the 28 monitored chokepoints. Unit of analysis is the chokepoint day. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S12: Every reading of the same response

Reading	Estimate	Where the variation comes from
The switch, one date for every treated city	0.035 (0.005)	Cities near the corridor against cities far from it, before and after
The corridor's counted traffic	0.052 (0.006)	Month-to-month movement in the transits at Bab el-Mandeb
The switch, the corridor that gained	0.026 (0.010)	The coast the corridor filled, on its own
The switch, the corridor that lost	0.043 (0.009)	The coast the corridor emptied, on its own
One satellite pass at a time	0.022 (0.006)	The same switch on 120 cities read pass by pass
Pass by pass, downwind only	0.030 (0.006)	Passes when the air moved from the lane to the city
Cities the ships pass	0.036 (0.008)	Coastal cities with no port inside the call radius
Cities the ships call at	0.034 (0.006)	Coastal cities with a port inside it

Notes. Each row is the coastline's response to the reallocation under a different reading of the same design, with the variation that identifies it named beside it. The first four rows are the monthly city panel, the next two break the same switch to the individual satellite pass on the 120 cities the corridor reaches most strongly, and the last two split the monthly panel by whether a port sits inside the call radius. Coverage, unit of analysis and clustering are as in Table 3 except where the row's own reading changes them. Standard errors are in parentheses beside each estimate, clustered as in the table the row is drawn from.

Table S13: The shock on the window the event figure uses

	(1) log NO_2	(2) log NO_2	(3) log NO_2
Exposure × Post	0.023^{***}	0.033^{***}	0.038^{***}
	(0.004)	(0.005)	(0.011)
Sample	Pandemic years out	Pandemic years out	Post through May 2024
Mean of dep. var.	-10.733	-10.733	-10.754
Observations	27,984	27,984	27,560
Adjusted R^2	0.958	0.944	0.945

Notes. Column (1) carries the fixed effects of Figure 6 and Column (2) those of Table 3. Coverage, unit of analysis and clustering are as in Figure 6. Column (3) keeps the pandemic years and ends the post period in May 2024, the months over which the container exemption held the canal's own pricing fixed for the treated lanes.

Table S14: The traffic that left, and the part of it that returned

	(1)	(2)	(3)	(4)
Exposure × Post	0.035***		-0.022	
	(0.005)		(0.022)	
Exposure × Share of the corridor gone		0.052***	0.076***	0.090*
		(0.006)	(0.026)	(0.046)
Window	Whole window	Whole window	Whole window	After the closure
Mean of dep. var.	-10.745	-10.745	-10.745	-10.719
Observations	38,160	38,160	38,160	13,144
Adjusted R^2	0.944	0.944	0.944	0.945

Notes. The share of the corridor gone is one minus the transits counted at Bab el-Mandeb in that month over their mean from January 2019 to November 2023. It is zero before the attacks, reaches 0.53 in July 2024 and falls back to 0.45 by the end of the window. Column (3) enters it against the switch, so the switch is left with whatever a step can explain that the path cannot. Column (4) drops every month before the closure. Coverage, unit of analysis and clustering are as in Table 3.

Table S15: The distance-band regression

	(1) log NO_2	(2) log NO_2
0-25 km × Exposure × Post	0.026***	0.028***
	(0.005)	(0.005)
25-50 km × Exposure × Post	0.023**	0.021**
	(0.011)	(0.008)
50-100 km × Exposure × Post	0.029**	0.023**
	(0.012)	(0.009)
100-200 km × Exposure × Post	0.027**	0.035***
	(0.013)	(0.010)
200-400 km × Exposure × Post	0.025	0.067***
	(0.035)	(0.022)
Sample	Egypt out	Egypt in
Mean of dep. var.	-10.835	-10.806
Observations	75,780	79,290
Adjusted R^2	0.924	0.927

Notes. Exposure interacted with an indicator for the city's distance from the coast. A coefficient is the effect of a full coastal standard deviation, and Figure S3 rescales each band by the spread the band actually has. The bands stop at 400 km because past it the measure holds a fiftieth of its coastal spread. Column (1) excludes Egypt and is the reading the text carries, and Column (2) keeps it. Detectable effects at eighty percent power are in Table S95. Coverage is cities within 400 km of the coast. Unit of analysis is the city by month. Standard errors are clustered on the city.

Table S16: City size and built-up density

	(1) Larger half	(2) Smaller half	(3) Denser half	(4) Sparser half
Exposure × Post	0.031***	0.031***	0.035***	0.030***
	(0.006)	(0.008)	(0.006)	(0.010)
Mean of dep. var.	-10.546	-10.929	-10.496	-10.965
Observations	18,180	18,630	17,820	18,630
Adjusted R^2	0.956	0.929	0.954	0.933

Notes. Splits at the median of the 2020 gridded population count inside the tile and of the 2020 built-up share. Coverage, unit of analysis and clustering are as in Table 3.

Table S17: The improved bearing and the lagged wind

	(1)	(2)	(3)	(4)
Exposure × Post	0.014	0.009	0.027***	0.026***
	(0.008)	(0.007)	(0.006)	(0.005)
× carrying wind	0.017**	0.027***	-0.008	-0.005
	(0.008)	(0.006)	(0.005)	(0.005)
Bearing and timing	Nearest point the hour's wind	Kernel-weighted the hour's wind	Nearest point the day before	Nearest point two days before
Observations	349,960	349,960	349,836	349,719

Notes. The outcome is the log column on one satellite pass, with the fixed effects of Table 4. The carrying-wind row is the premium on a pass whose wind projects positively onto the stated bearing, entered with its own lower-order terms. Column (1) is the paper's bearing, the direction from the nearest solved lane. Column (2) replaces it with the kernel-weighted mean bearing over the whole corridor, which is the direction the exposure measure itself implies. Columns (3) and (4) keep the nearest-point bearing and match each pass to the wind of one and two days earlier. Coverage is the overpass panel of Table 4, unit of analysis the city by satellite pass, inference clustered on the city.

Table S18: The six days of 2021, on the same gradient

	(1)	(2)	(3)	(4)
Exposure × Blocked days, 2021	0.080***			
	(0.027)			
Exposure × Blocked days, 2021 downwind		0.138***		
		(0.033)		
Exposure × Blocked days, 2019			0.012	
			(0.018)	
Exposure × Blocked days, 2022				0.013
				(0.015)
Spring of	2021	2021	2019	2022
Sample	All passes	Downwind passes	All passes	All passes
Mean of dep. var.	-10.338	-10.424	-10.381	-10.323
Observations	13,101	8,109	12,905	12,813
Adjusted R^2	0.719	0.679	0.682	0.701

Notes. The unit is one city on one satellite pass, over a February to mid-May window each year, with city and day effects. Blocked days are the six the canal was shut by the grounding. The last two columns run the identical six calendar days of 2019 and 2022, the clean years either side, with the pandemic spring of 2020 set aside. Coverage is the 120 coastal cities of the overpass panel. Standard errors are clustered on the city. Unit of analysis is the city by satellite pass.

Table S19: The two episodes as one dose response

Response to the 2021 blockage, per month fully shut	0.080
Share of the corridor gone after December 2023	0.437
What the 2021 response implies for that share	0.035
The 2023 switch, estimated	0.035
Ratio of the two	1.00
95 percent interval	0.28-1.72
Days the 2021 canal stood shut	6
calendar days in the regression window	7

Notes. The 2021 regressor is the share of the month the canal stood shut, so its coefficient is the effect of a month fully closed. Scaling it by the share of the corridor's traffic gone after December 2023 gives what the smaller episode implies for the larger one. Coverage is the overpass panel of Table S18 for the 2021 reading and the estimation sample of Table 3 for the 2023 one. Unit of analysis is the derived quantity, and the inputs carry the standard errors of Tables 3 and S18.

Table S20: The water and the coast, tied on one chain

Ship share of the column over the counting box	0.413
Column over that box, μmol m^{-2}	8.9
Median coastal city column, μmol m^{-2}	20.9
Corridor weight, median city within 50 km of a lane	1.86
Corridor weight, median coastal city	0.30
Distance from the median coastal city to the nearest lane, km	248
One standard deviation of the reallocation, weight units	0.64
A corridor's plume at the median city, percent of its column	
anchored on the box elasticity	2.8
anchored on the far-field enhancement	0.6
the estimate, rescaled from one standard deviation	1.6

Notes. A corridor weight is the kernel summed over both corridors and read at one city. The box anchor multiplies the ship share by the ratio of the median city's weight to a lane-side city's weight and by the ratio of the box column to the city column. The far-field anchor decays the fitted enhancement of Table S111 to the median city's distance. Coverage is the estimation sample. Unit of analysis is the city. The chain is arithmetic on estimated inputs and carries no standard errors of its own.

Table S21: The four pollutants

	(1) NO_2	(2) SO_2	(3) CO	(4) Aerosol index
Exposure × Post	0.099***	-0.070***	0.042***	-0.024***
	(0.016)	(0.026)	(0.015)	(0.008)
Mean of dep. var.	0.038	0.007	0.045	0.041
Observations	32,138	32,138	32,138	32,138
Adjusted R^2	0.691	0.665	0.936	0.953

Notes. Each pollutant is divided by its own within-city standard deviation before estimation, so the four coefficients are on one scale. The four are read on the 32,138 city-months where every outcome is defined, so that a contrast across the pollutants is not partly a contrast across coastlines. The binding margin is the sulfur column, whose retrieval falls below the detection floor in a seventh of the balanced nitrogen dioxide panel's city-months, and that selection is itself priced in Table S24. Table S22 tests the four together rather than one at a time. Coverage, unit of analysis and clustering are as in Table 3.

Table S22: The four pollutants tested as one family

Outcome	Family-corrected p
NO_2	0.001
SO_2	0.024
CO	0.024
Absorbing aerosol index	0.024

Notes. The four outcomes are resampled together in the stepdown manner of Romano and Wolf (2005), which keeps the correlation across four retrievals that share a sensor and a sky and retests the outcomes that survive as a smaller family at each step. The resampling unit is the city, matching the clustering in each equation, over a thousand draws. Coverage, unit of analysis and clustering are as in Table 3.

Table S23: The sulfur fall, by side and by rule

	(1)	(2)	(3)
Exposure × Post	-0.072***		
	(0.026)		
Exposure × Post, corridor that gained		-0.178***	-0.178***
		(0.043)	(0.043)
Exposure × Post, corridor that lost		0.029	
		(0.030)	
lost corridor, before the control area			0.021
			(0.030)
lost corridor, under the control area			0.069
			(0.056)
Mean of dep. var.	-9.341	-9.341	-9.341
Observations	32,138	32,138	32,138
Adjusted R^2	0.857	0.857	0.857

Notes. The outcome is the log total sulfur dioxide column, read on the city-months where all four retrievals are usable. Column (3) splits the lost corridor's term at May 2025, when the Mediterranean sulfur control area began to bind on that coast. Detectable effects at eighty percent power are in Table S95. Coverage, unit of analysis and clustering are as in Table 3.

Table S24: Missing retrievals on the treatment

Exposure × Post	0.013***
	(0.004)
Missing retrieval of	Sulfur dioxide
Mean of dep. var.	0.152
Observations	38,160
Adjusted R^2	0.653

Notes. The outcome is an indicator for a city-month with no usable retrieval of the named species, on the full balanced panel, so the estimate is a linear probability of dropping out. Nitrogen dioxide, carbon monoxide and the aerosol index are never missing on this panel, so sulfur dioxide is the only species with an extensive margin. City and country by month effects are absorbed. Coverage and unit of analysis are as in Table 3, inference clustered on the city.

Table S25: The four species through the same distance bands

Species	0-50 km	50-100 km	100-200 km	200-400 km
NO_2	0.041	0.040	0.021	0.008
	(0.008)	(0.017)	(0.011)	(0.011)
SO_2	-0.070	-0.110	-0.064	0.021
	(0.016)	(0.033)	(0.026)	(0.028)
CO	0.025	0.029	0.009	0.015
	(0.008)	(0.017)	(0.007)	(0.007)
Absorbing aerosol index	-0.020	-0.007	-0.016	-0.017
	(0.007)	(0.015)	(0.011)	(0.011)
Share of the coastal effect surviving to the far band		NO_2 0.20, CO 0.58		
Observations in the band	28,260	8,190	15,480	23,850

Notes. Each species standardized on the coastal sample, so the columns are comparable across rows, and each band's coefficient rescaled by the exposure spread inside that band, so they are comparable across columns. The surviving-share row is the 200-400 kilometer coefficient over the 0-50 kilometer one, reported for the two species whose coastal reading is positive. Coverage, unit of analysis and clustering are as in Table 3, with Egypt excluded and every distance from the coast included.

Table S26: The negative control and the signed prediction

	(1)	(2)	(3)	(4)
Exposure × Post	0.026***	0.003	-0.003***	
	(0.004)	(0.005)	(0.001)	
Exposure × Post, corridor that gained				-0.006***
				(0.001)
Exposure × Post, corridor that lost				0.000
				(0.001)
Outcome	Nitrogen dioxide	Formaldehyde	Ozone	Ozone
Mean of dep. var.	-10.745	-9.170	-2.041	-2.041
Observations	38,160	38,155	38,160	38,160
Adjusted R^2	0.978	0.972	0.992	0.992

Notes. Formaldehyde is the tropospheric column and ozone the total column, each read over the same city boxes as the main outcome and entered in logs. The stratosphere carries most of a total ozone column, so the ozone columns are a signed prediction and not a magnitude. Formaldehyde and ozone carry a strong city-level seasonal cycle of their own, biogenic for the first and photochemical for the second, so every column absorbs city-by-calendar-month effects, and Column (1) shows the headline estimate under the same absorption. Column (4) splits exposure by the side of the reallocation. Coverage, unit of analysis and clustering are as in Table 3.

Table S27: The scene itself, as an outcome

	(1) Cloud fraction	(2) Surface albedo
Exposure × Post	0.0010	-0.0019**
	(0.0032)	(0.0008)
Mean of dep. var.	0.313	0.203
Observations	38,160	38,160
Adjusted R^2	0.890	0.959

Notes. The outcomes are the monthly mean cloud fraction and surface albedo of the mission's own cloud product, read over the same city boxes as the column, under the specification of Table 3. A change in the retrieval scene that moved every product at once would appear here. Coverage, unit of analysis and clustering are as in Table 3.

Table S28: The aerosol loading conditioned on directly

	(1) log NO_2	(2) log NO_2
Exposure × Post	0.035***	0.037***
	(0.005)	(0.005)
Absorbing aerosol index		0.150***
		(0.023)
Sample	Index retrieved	Index retrieved
Mean of dep. var.	-10.745	-10.745
Observations	38,160	38,160
Adjusted R^2	0.944	0.945

Notes. The outcome is the log nitrogen dioxide column under the specification of Table 3, on the city-months where the absorbing aerosol index is retrieved. Column (2) adds the index as a control. The index is itself moved by the traffic, so the column conditions on an outcome and bounds the interference channel rather than estimating a preferred specification. Coverage, unit of analysis and clustering are as in Table 3.

S5.3 The Plume against the Port

Table S29: Port calls and tanker calls on the corridor

	(1) log port calls	(2) log tanker calls
Exposure × Post	0.029**	0.029**
	(0.013)	(0.012)
Clustered on one-degree cells	(0.013)	(0.012)
Sample	Cities with a port	Cities with a port
Mean of dep. var.	0.762	0.308
Observations	13,398	13,398
Adjusted R^2	0.968	0.944

Notes. Equation (8). The outcome is the log of one plus the monthly mean of daily calls at the city's own port, so a city-month with no call carries a zero rather than falling out. Coverage is the 218 cities within 100 km of the coast with a port inside the 40 kilometer call radius, with coasts at war and the pandemic years excluded. Unit of analysis is the city by month. Standard errors are clustered on the city, and the second parenthesis row clusters on one-degree cells instead, which agrees with the city clustering to three decimals in both columns.

Table S30: The city trend, on both outcomes

	(1) log calls	(2) log calls	(3) log tanker calls	(4) log tanker calls
Exposure × Post	0.030***	0.029**	0.012	0.029**
	(0.011)	(0.013)	(0.008)	(0.012)
City trend	No	Yes	No	Yes
Clustered on	City	City	City	City
Mean of dep. var.	0.762	0.762	0.308	0.308
Observations	13,398	13,398	13,398	13,398
Adjusted R^2	0.960	0.968	0.929	0.944

Notes. Columns (1) and (2) are port calls and Columns (3) and (4) tanker calls, each without and then with the city trend. Coverage, unit of analysis and clustering are as in Table S29.

Table S31: The placebo dates on port calls and tanker calls

	(1) log calls	(2) log calls	(3) log tanker calls	(4) log tanker calls
Exposure × Placebo post	-0.017	0.008	-0.011	0.005
	(0.016)	(0.010)	(0.014)	(0.009)
Placebo date	12/2019	12/2022	12/2019	12/2022
Mean of dep. var.	0.748	0.748	0.303	0.303
Observations	7,105	7,105	7,105	7,105
Adjusted R^2	0.981	0.981	0.967	0.967

Notes. Each column moves the shock to a December before the real one and keeps only the months before the real one, so the exposure measure holds its variation and the date is the only thing that moves. Columns (1) and (2) are port calls and Columns (3) and (4) tanker calls. Coverage, unit of analysis and clustering are as in Table S29.

Table S32: The arrivals counted rather than logged

	(1) Calls	(2) Tanker calls	(3) Trade	(4) Calls	(5) Tanker calls	(6) Trade
Exposure × Post	0.102^{***}	0.075	0.018^{***}	0.054^{***}	0.032	-0.016
	(0.017)	(0.050)	(0.006)	(0.019)	(0.020)	(0.130)
Estimator	Poisson	Poisson	Poisson	Log	Log	Log
Mean of dep. var.	1.499	0.394	19,891	0.684	0.270	7.621
Observations	5,092	5,006	5,090	5,214	5,214	5,214

Notes. Port calls, tanker calls and trade volume are counts with zeros, so the first three columns estimate them by Poisson with the same effects, which drops the city-months whose outcome is zero in every month of a fixed-effect cell, 848, 934 and 850 of the 5,940 the gained side supplies. The last three columns repeat the outcomes in logs. Coverage is the coastal cities of the corridor that gained the traffic, over the window with the two pandemic years out. Unit of analysis is the city by month. Standard errors are clustered on the city.

Table S33: Port calls at cities the ships pass against cities they stop at

	(1) log port calls	(2) log tanker calls
Exposure × Post × Pass-by	0.000	-0.002
	(0.005)	(0.006)
Exposure × Post × Call-at	0.057***	0.022*
	(0.018)	(0.013)
Equality of the two	Rejected	Not rejected
p of that test	0.001	0.122
Mean of dep. var.	0.287	0.109
Observations	13,596	13,596
Adjusted R^2	0.991	0.970

Notes. The call outcomes split on whether a city owns a port inside the 40 kilometer call radius, over every city on the corridor that gained the traffic rather than over the port-owning cities alone, since a split needs both groups present. Fuel is sold and cargo is handled where a ship stops, so an arrival effect that is real sits on the call-at term and is absent from the pass-by term. Coverage is the gained-side coastal cities outside the war coasts with the pandemic years removed. Unit of analysis is the city by month. Standard errors are clustered on the city. Detectable effects at eighty percent power are in Table S95.

Table S34: Trade volume at the port against port calls

	log trade volume
Exposure × Post	-0.033
	(0.079)
Mean of dep. var.	7.900
Observations	13,398
Adjusted R^2	0.915

Notes. Trade volume is the reported import plus export estimate at the port. Coverage, unit of analysis and clustering are as in Table S29.

S5.4 Incidence

Table S35: Every population and every net the paper carries

Reading	Value
Released-to-taken ratio, person-weighted	
Canonical, the estimation coasts at the pooled coefficient	2.63
war coasts carried in	3.59
cities carrying an income figure	3.04
the 100-200 km band at its own coefficient	5.1
the band at the Egypt-out coefficient	4.8
Egyptian band cities removed	3.8
the gridded strip, no city definition	4.85
People, millions	
Taken on, the coasts the estimate covers	36.7
Released, the coasts the estimate covers	80.9
Released, war coasts carried in	119.2
The whole coastline, exclusive assignment	275.0
the same coastline under the box rule	346.6
The gridded strip, taken on and released	89.2 and 422.3
Net exposure removed, million person log points a year	
The account, cities carrying an income figure	2.8
The account, every city with a population count	3.6
The gridded strip, no city definition at all	11.7
Attributable to this closure at the uniform scalar	2.4
discounting only the Cape's addition	4.1
At the downwind premium alone, the plume's own floor	1.7
Cities	
The incidence account enumerates	507
of which on the ten war coasts	74
Enter the estimation sample	433
of which own a port inside the call radius	218

Notes. Each row is read live from the file or table that owns it, the incidence account for the estimation-coast populations and nets, the exclusive-assignment table for the coastline totals, the gridded-strip table for the account with no city definition, and the attribution share of the carbon account for the attributable rows. The city block counts the register's rows under the rules named beside them. The figures are enumerated rather than sampled, so the table carries no standard errors.

Table S36: The account extended one band inland

	People, m	Person-weighted
Within 100 km		
Taken on	36.7	1,388,613
Released	119.2	-4,990,798
ratio, released to taken		3.6
With the 100-200 km band		
Taken on	43.5	1,475,196
Released	174.9	-7,574,091
ratio, released to taken		5.1
The band at the Egypt-out coefficient		
Taken on	43.5	1,455,405
Released	174.9	-6,983,624
ratio, released to taken		4.8
Egyptian band cities removed		
Taken on	43.5	1,475,196
Released	145.0	-5,623,237
ratio, released to taken		3.8

Notes. The base sample is the war-inclusive coastline. The second band enters at the coefficient the distance profile estimates for 100-200 km, and the bands beyond it are excluded. The lower block reprices the band with Egypt out of the band sample. Coverage is every city within 200 km of the coast with a population count, coasts at war excluded from the estimation and counted in the account. Unit of analysis is the city. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S37: The twenty most populous cities the reallocation moved

City	Country	People, m	Exposure, sd	Income
The ten largest cities the corridor arrived on				
Casablanca	MAR	3.9	0.60	3,463
Durban	ZAF	2.8	0.20	6,534
Cape Town	ZAF	2.5	2.65	6,534
Freetown	SLE	1.4	1.15	860
Monrovia	LBR	1.4	0.78	745
Banjul	GMB	1.3	1.88	836
Conakry	GIN	1.3	0.90	1,417
Nouakchott	MRT	1.2	1.40	1,961
Kenitra	MAR	1.0	0.18	3,463
Stellenbosch	ZAF	1.0	2.45	6,534
The ten largest cities the corridor left				
Alexandria	EGY	4.2	-1.80	4,233
Damascus	SYR	3.7	-0.31	1,057
Algiers	DZA	3.6	-1.23	4,960
Aleppo	SYR	3.2	-0.07	1,057
Tanta	EGY	3.2	-2.25	4,233
Mit Ghamr	EGY	2.8	-2.59	4,233
Izmir	TUR	2.7	-0.29	10,898
Sidon	LBN	2.6	-0.43	3,654
Jeddah	SAU	2.5	-2.29	38,510
Mansoura	EGY	2.5	-2.68	4,233

Notes. The ten most populous cities on each side of the reallocation at the 0.05 cut on standardized exposure. Income is the country figure the distribution account uses. Coverage and unit of analysis are as in Table 6. The cities are enumerated rather than sampled, so the table carries no standard errors.

Table S38: The account on the population over sixty-five

	Taken on	Released
Share of the population over sixty-five	0.053	0.082
People over sixty-five, millions	1.9	9.7
Person-weighted, over sixty-five only, thousands	78	-343

Notes. The base sample is the war-inclusive coastline. Age structure is the gridded age and sex product (WorldPop, 2018) summed over each city box, the eight bands from sixty-five up. What is taken from the raster is each city's over-65 share and the level is the account's own. The last row weights each city's person log points by its over-65 share. Coverage is the incidence cities with an age raster. Unit of analysis is the city. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S39: The account carried east of its window

	People, m	Person log points
Taken on east of the window	94.6	3,237,142
Released east of the window	179.2	-4,775,272
Inside the 0.05 cut, on neither side	834.5	-29,525
Net east of the window		-1,567,654

Notes. The base sample is the eastern extension of the war-inclusive coastline, while the canonical account of Section 8 runs on the estimation sample. The same kernel, the same population raster and the same estimated response are carried to every urban center between 65 and 130 east that either state's corridor reaches. The coefficient is the one estimated on the western sample. Coverage is the 2,850 eastern urban centers the corridor reaches in either state. Unit of analysis is the city. The coastline is enumerated rather than sampled, so the table carries no standard errors.

Table S40: The incidence account with no thinning at all

	Cities	People, m	Mean $\Delta \log NO_2$	Person-weighted
Gained the corridor	257	136.6	0.0080	1,549,313
Lost it	889	301.9	-0.0116	-3,957,684
Neither	513	267.7	0.0000	0

Notes. Every urban center of the global settlement database within 250 km of either corridor. Population is the database's own count for the center, a narrower object than the box populations of Table 6, so the comparison is of shares rather than levels. A city counts as gaining or losing on the same cut as Table 6. Unit of analysis is the urban center. Coverage is every urban center on the coast with no thinning. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S41: The account with no city definition at all

	People, m	Person-weighted
Taken on	89.2	3,041,741
Released	422.3	-14,755,728

Notes. WorldPop's global 2020 mosaic summed onto the solver's tenth-of-a-degree grid, the kernel computed at every populated node within 200 km of the sea, and the account the sum of population times the response times the standardized shock, with the same 0.05 cut the incidence table uses. The kernel varies over a hundred and fifty kilometers, so the grid loses nothing the kernel can see, and the population inside a cell is summed rather than sampled. Coverage is every populated coastal cell in the study window, 1,333 million people. Unit of analysis is the grid cell. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S42: The account at two calendars of one population product

	2020 calendar	2025 calendar
People on the side taken on, millions	211.5	235.9
People on the side released, millions	255.5	277.7
Person-weighted released over taken	4.44	4.40
Net exposure removed, millions a year	7.7	8.3

Notes. Both columns read the constrained one kilometer global mosaic of the same WorldPop release over the account's own city boxes, with every cell given to the nearest city alone, so moving between them moves the calendar and nothing else. The rows split every city of the account by the sign of its predicted change with no 0.05 cut applied, war coasts in. Coverage is the incidence account's cities. Unit of analysis is the city. The counts are enumerated rather than sampled, so the table carries no standard errors.

Table S43: The same comparison with the geography removed

	Cities	Gained	Lost	Ratio
All countries	294	$4,725	$11,748	40 percent
Income deciles trimmed	261	$4,798	$8,100	59 percent
Africa only	172	$3,405	$4,375	78 percent
Excluding the Gulf and Israel	280	$4,725	$9,134	52 percent
The 3 largest on each side out	288	$4,557	$12,725	36 percent
The 10 largest on each side out	274	$5,702	$13,649	42 percent

Notes. Income per head is the 2022 country figure weighted by the population bearing the exposure. The second row drops the top and bottom income deciles of the sample, the third keeps only cities on African coasts, and the fourth drops the Gulf states and Israel. The last two take the largest cities out of each side, since a ratio of two population-weighted means can rest on a few places. Coverage is the coastal cities with both a population count and an income figure, split at the same cut as the incidence table. Unit of analysis is the city. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S44: The response split at the median income

	(1)	(2)
Exposure × Post	0.035***	
	(0.005)	
Exposure × Post, above median income		0.039***
		(0.007)
Exposure × Post, below median income		0.021*
		(0.011)
Coefficient	Pooled	Split at the median
Difference, rich less poor		0.018
standard error		0.012
p		0.159
Mean of dep. var.	-10.741	-10.741
Observations	37,980	37,980
Adjusted R^2	0.944	0.944

Notes. The incidence account applies one coefficient to every city and then weights by income, so whether the coefficient itself varies with income decides whether that weighting compounds two things or one. The split is at the median income of the coastal sample, on the cities carrying an income figure. Coverage, unit of analysis and clustering are as in Table 3.

Table S45: The schedule at named weights

Ratio at $\eta = 0.0$	2.26
Ratio at $\eta = 0.5$	1.80
Ratio at $\eta = 1.0$	1.31
Ratio at $\eta = 1.5$	0.92
Ratio at $\eta = 2.0$	0.66
Ratio at $\eta = 2.5$	0.52
Ratio at $\eta = 3.0$	0.45
The η at which the two arcs price equally	1.38

Notes. The schedule of Figure 8, left panel, read at named values of the inequality aversion. Coverage and unit of analysis are as in Table 8. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S46: A mortality translation of the account

	Deaths a year	At the risk interval
Added where the corridor arrived	107	54-212
Avoided where it left	394	198-780
Net avoided	287	144-568

Notes. The per-city log change is chained onto the pre-closure surface concentration of the surface nitrogen dioxide product (Anenberg et al., 2022), converted at 1.88 micrograms per cubic meter a part per billion, and carried through the meta-analytic all-cause relative risk of 1.02 for 10 micrograms per cubic meter (Huangfu and Atkinson, 2020) with each country's crude death rate from the World Development Indicators. Coverage is the 485 account cities carrying a surface level and a death rate. Unit of analysis is the city. The interval column reads the chain at the risk function's interval, 1.01 to 1.04.

Table S47: The crossing, under each reading of the account

	The crossing η^{*}
Common coefficient on both sides	1.38
Each side its own coefficient	2.11
Incomes winsorized at 5 and 95	1.38
Incomes at purchasing power parity	1.20
Income at the city's own subnational grid	1.11
Largest weighted contributor out, Banjul	1.67
Each city on its own income half's coefficient	1.84
Any one of the ten largest out, lowest	1.41
highest	1.67
Weights capped at the 95th percentile	1.38
Weights capped at the 99th percentile	1.38
Egypt repriced at its 2024 income	1.58

Notes. Each row recomputes the inequality aversion at which the weighted released equals the weighted taken. The side-specific row applies the split estimates of Table 3 to each side. The subnational row replaces each city's national income with the mean of the downscaled admin-2 grid of Kummu, Kosonen and Masoumzadeh Sayyar (2025) over the city's own box, at purchasing power parity for 2022. The Egyptian repricing row swaps the 2022 national income for the 2024 one, after the March 2024 float, with every other country held at the 2022 vintage. Coverage and unit of analysis are as in Table 8. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S48: The account under the income data's own gaps

	Income of taken over released	The crossing η^*
As the account stands	0.40	1.38
War coasts out of the account entirely	0.37	1.02
Dropped cities back at the last reported income	0.43	none
at 0.5 times the last reported income	0.43	none
at 2 times the last reported income	0.43	3.25
at 4 times the last reported income	0.43	1.76

Notes. The first column is the population-weighted income of the people taken on over that of the people released. The quantities are enumerated rather than sampled, so no standard errors are carried. The refilled row returns every dropped city at the last income its country reported to the World Development Indicators, Yemen in 2018 and Eritrea in 2011, dated rather than modeled, and a crossing of none means the two readings never meet, the released side staying poorer at every aversion. The multiplier rows scale those last-reported figures so the crossing can be read as a function of what the unmeasured coasts are assumed to earn. The territories no source prices stay out. Coverage and unit of analysis are as in Table 8.

Table S49: Egypt and the canal bank, consolidated

	Coefficient	Released-to-taken ratio	Crossing η^*
Everything in	0.035	3.59	1.38
Egypt out of the estimation and the account	0.034	2.31	0.60
The canal-bank cities out	0.036	3.46	1.33
the lost side alone on that cut	0.046	–	–

Notes. Each row reads the paper's three headline objects under one sample cut. The coefficient is the switch under the specification of Table 3, estimated without the named cities, and the ratio and the crossing recompute the incidence account and the welfare crossing on the same cut. The canal-bank row drops Port Said, Al Adabiyah and Ismailia, the boxes on the canal itself whose columns carry queueing and convoying as well as passage. The Egypt row removes the account's largest poor released population, and the income assigned to the released coastline is priced in Table S48. Standard errors and units as in the source tables. A dash marks the two quantities the lost-side cut does not define.

Table S50: The share of each side's weighted sum in its five largest cities

	$\eta = 1.5$	$\eta = 2$
Share of the weighted taken side in its five largest cities	51 percent	58 percent
Share of the weighted released side in its five largest cities	30 percent	36 percent

Notes. The share of each side's income-weighted person log points carried by its five largest city contributions. Coverage and unit of analysis are as in Table 8. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S51: The bands the decay length puts on the headline objects

	At 150 km	Over the fitted length
Efficient ratio of the arcs	2.59	2.44, range 2.23-2.62
The crossing η^*	1.38	1.34, range 1.27-1.41
Net exposure removed, millions a year	3.6	2.5-4.7
The indifference value, dollars	1,279	974-1,864

Notes. Two thousand draws of the decay length from its fitted sampling distribution, the cities held fixed and each city's arc weights interpolated across the decay grid, and of the response coefficient from its own, which scales the net and the indifference value linearly. Unit of analysis is the draw, the interval is the draws' own percentile range, and no clustered error enters because the coastline is enumerated, while the estimated response cancels out of the ratio and rescales the crossing's axis without moving it, so the decay length is the one estimated input either object inherits. Coverage is as in Tables 6 and 8.

Table S52: The construction ingredients varied one at a time and jointly

What varies	Range
The efficient ratio, 2.59 at the baseline	
The decay length, over the sweep's interior	2.44-2.79
each corridor on its own fitted pair	2.70
The endpoint set, merged, deepened, unweighted, gravity-weighted or all-hull	2.14-3.45
The grid, sixteen neighbors rather than eight	2.58
The weighting window, each pre-year alone	2.57-2.61
The standoff, over the settings the gate rule admits	2.11-3.03
The base a charge is written on	1.96-3.73
The kernel's isotropy, wind weighted	2.16
The kernel replaced by a plume model in the ERA5 fields	2.74
The damage's curvature, inverse to the column	1.76
The sample, cities carrying an income figure	2.26
Joint span of every reading above	1.76-3.73
The crossing, 1.38 at the baseline	
Across the coefficient readings	1.11-2.11
Exposure in level units rather than log points	2.34
War coasts out of the account entirely	1.02
Dropped cities refilled at stale incomes	none

Notes. Each row varies one ingredient and reads the resulting range out of the table that computes it. The grid row re-solves every lane on a sixteen-neighbor grid admitting the knight's moves. The plume row replaces the kernel with a sector-Gaussian plume model with a first-order life of eight hours in transit. The refill row is the one completion under which no crossing exists. Coverage and unit of analysis follow the source tables. The ranges are enumerations of computed readings, not sampling intervals, so no standard errors are carried.

S5.5 The Arc Price

Table S53: Every reading of the arc ratio

Reading	Ratio	What it changes
The baseline, at the length the corridor weights are built on	2.59	none
Charged per transit at the sizes observed	1.96	the base a charge is written on
with emissions as deadweight to the 0.65	2.16	the base a charge is written on
Charged per tonne of fuel	3.73	the base a charge is written on
At the length the water itself fits	2.44	the decay length
Each corridor on its own fitted length and amplitude	2.70	the decay length and the concentration
Netting out water both routings cross	2.41	the water counted
Pricing each lane by its closest point	2.53	the kernel
Weighted by the hours the wind carries the plume ashore	2.16	the kernel's isotropy
With the canal emitting at its mandated convoy speed	2.19	what a mile of convoy emits
On the cities carrying an income figure	2.26	the sample
Weighted by the people over sixty-five	2.49	who counts as a person at risk
Jointly by the over-65 share and inverse income	1.62	who counts as a person at risk
Letting the response differ by city density	2.61	the response
Letting the response differ by city size	2.59	the response
Under a damage inverse to the baseline column	1.76	the damage
Netting out the toll-setter's own nationals	2.35	the population counted
With the Cape fleet's observed speed-up priced	2.35	the emission rate

Notes. One baseline and a named departure from it in every row, so a reader meeting a ratio anywhere in the paper can see which ingredient produced it. The baseline is the exposure a transit delivers, summed over every coastal city with a population count, at the 150 kilometer length the corridor weights are built on, below the length the water fits, and per transit of a given hull. Each figure is read out of the table or the log that produces it. Coverage and unit of analysis are as in Table 10. The coastline is enumerated rather than sampled, so the table carries no standard errors.

Table S54: The ratio when the canal emits at its mandated convoy speed

Exponent on speed in fuel per mile	0.0	0.5	1.0	1.5	2.0	2.5	3.0
What the convoy stretch emits per mile	1.000	0.707	0.500	0.354	0.250	0.177	0.125
Efficient price of the Red Sea arc, Cape = 1	2.59	2.43	2.32	2.24	2.19	2.15	2.12

Notes. Convoy speed in the canal and the Gulf of Suez is half of sea speed for the hulls that transit, so what a mile of that stretch emits relative to open water is the speed ratio raised to the exponent in the first row. Fuel per hour is conventionally cubic in speed and fuel per mile quadratic, and an exponent of zero recovers the baseline of Table 10 exactly. Coverage and unit of analysis are as in Table 10. The coastline is enumerated rather than sampled, so the table carries no standard errors.

Table S55: The composition of the Cape's extra transits

Transits a day at the Cape, before	47.4
Transits a day at the Cape, since	87.7
The gain	40.3
diverted from Bab el-Mandeb	26.3
the growth the unaffected passages show	-1.3
pushed from Panama, at most	1.3
unattributed	14.0

Notes. Daily means before and since December 2023. The growth row applies the median transit growth of the passages the diversion never touched to the Cape's own base. The Panama row is that canal's own daily loss, capped at the share a Cape routing can absorb once the Magellan gain on the same counted series is netted out. Coverage is the chokepoint panel. Unit of analysis is the chokepoint day. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S56: Transits gained at the Cape and lost at Bab el-Mandeb, by hull class

Hull class	Gained at the Cape	Lost at Bab el-Mandeb	Difference
Container	14.0	11.1	2.9
Dry bulk	15.0	4.9	10.1
Tanker	6.9	7.5	-0.6
General cargo	2.1	0.4	1.7
Roll-on roll-off	2.4	2.5	-0.1
All hulls	40.3	26.3	14.0

Notes. Daily mean transits by hull class before and since December 2023, from the counting points themselves. Panama traffic rerouted around Africa is bounded above by that canal's counted loss net of the Magellan gain, which is where the one-transit-a-day cap the attribution carries comes from. Coverage is the daily transit series, January 2019 to the end of the window. Unit of observation is the passage. The counts are enumerated, so the table carries no standard errors.

Table S57: The people a passage reaches, at four decay scales

Passage	50 km	150 km	300 km	500 km
Suez Canal	5.4	29.2	55.7	87.2
Sunda Strait	4.6	25.2	47.0	69.3
Taiwan Strait	2.1	18.9	60.0	131.5
Bosporus Strait	10.2	18.0	31.4	60.4
Bohai Strait	1.1	16.8	66.0	143.4
Dover Strait	1.9	16.3	38.2	64.3
Korea Strait	0.9	9.8	38.0	96.6
Malacca Strait	1.2	9.2	22.0	45.1
Mindoro Strait	0.3	7.5	20.6	46.4
Gibraltar Strait	1.6	6.3	15.7	30.4
Lombok Strait	1.4	5.6	16.0	33.7
Bab el-Mandeb Strait	0.5	4.6	12.8	28.3
Oresund Strait	1.1	4.0	15.8	40.0
Luzon Strait	0.0	2.9	24.0	75.2
Rank correlation with the 150 km scale	0.898	1.000	0.959	0.887

Notes. Population in millions at the gate point, weighted by the same exponential the corridor weights use, at scales from a third of the headline to more than three times it. These are point-kernel headcounts and sit beside Table 11's along-the-passage columns as two readings of one geometry, the point and the line. The centers and the passages are enumerated rather than sampled, so the table carries no standard errors. Coverage is every urban center in the settlement layer and the 28 passages the transit product covers, of which the 14 heaviest are shown. Unit of analysis is the passage.

Table S58: The schedule the toll traces

Toll, nm	Lanes	Traffic gone	Exposure	Income weighted	Revenue	Trade on the canal	Trade on the lane	Income	Carbon
0	0	0	100.0	100.0	0	100	100	$11,901	0.0
1,000	0	0	100.0	100.0	31	99	99	$11,901	0.0
2,000	0	0	100.0	100.0	62	98	98	$11,901	0.0
3,000	6	3	97.8	98.5	91	94	97	$11,871	0.5
3,300	6	3	97.8	98.5	100	94	97	$11,871	0.5
4,000	80	67	43.0	62.9	41	31	96	$10,262	14.5
5,000	108	83	37.7	61.6	26	16	96	$9,977	19.1
6,000	112	84	37.3	61.8	29	15	96	$9,922	19.5
8,000	138	99	39.4	66.7	2	1	96	$10,102	25.9

Notes. Lanes and traffic gone are the crossing lanes the toll drives to the detour and the share of transits they carry. Exposure and the income-weighted exposure beside it are the plume the surviving traffic still delivers, each indexed on the untolled column, as are the two trade columns. Revenue is the toll times the traffic that stays, indexed on its peak. Trade retained carries the arc price into delivered cost. Carbon is in million tonnes. Fuel prices a mile of toll at the fleet-average burn and 2024 bunker price. Carbon is extra distance sailed times ship deadweight and IMO fleet efficiency. Income is the population-weighted mean across exposed cities. Coverage is the 507 coastal cities of the incidence account and the crossing lanes. Unit of observation is the toll. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S59: The efficient differential against the toll the authority posts

Container transits through the canal in 2023	5,847
Canal net tonnage they carried, millions	657.0
per transit	112,371
What the posted rates charge that transit, rights	414,562
in dollars	551,995
Carbon one transit burns, tonnes	
through the canal	3,717
round the Cape	5,364
Exposure one transit delivers, person log points	
through the canal	399
round the Cape	154
The value of exposure at which the two arcs price equally	1,279
Efficient extra charge on a canal transit at 200 dollars	-263,948
against what the canal already charges, percent	-48
Efficient extra charge on a canal transit at 1,000 dollars	-68,171
against what the canal already charges, percent	-12
Efficient extra charge on a canal transit at 2,000 dollars	176,550
against what the canal already charges, percent	32
Efficient extra charge on a canal transit at 5,000 dollars	910,712
against what the canal already charges, percent	165

Notes. Carbon and exposure per transit follow Table 9 and Table 10. The transit due is the Suez Canal Authority's posted rates for container ships in the year the closure began, converted at that day's special drawing right. The total efficient extra charge is the carbon price times the difference in fuel burned plus the value of exposure times the difference in exposure delivered. The dollar figure named in each of those rows is the value of one person log point of exposure, with carbon held at 190 dollars a tonne throughout. Unit of analysis is one container transit. The rate card is enumerated rather than sampled, so the table carries no standard errors.

Table S60: The diversion's ledger

	A year	What the entry is
Exposure removed, million person log points	3.6	externality, before attribution
attributed at the Bab el-Mandeb share	2.4	externality
removed beside unchanged lanes by withdrawn capacity	1.2	externality, uncounted above
Carbon added, million tonnes	15.8	externality
worth a year at 190 dollars a tonne, billions	3.0	
Fuel the detour burns, million tonnes	5.1	private resource
worth a year at 600 dollars a tonne, billions	3.0	
Ship time the detour adds, thousand ship-days	93	private resource
worth a year at 30 to 50 thousand dollars a day, billions	2.8-4.7	
Trade moved off Suez-dependent routes, billions	65	displacement within importers
Canal receipts forgone, billions	6	fiscal transfer
Exposure value balancing the externality account, dollars	1,279	
the same once the carriers' own fuel bill is added, dollars	2,575	
and the time at the published charter range, dollars	3,764-4,556	

Notes. Annual flows are attributed at the Bab el-Mandeb share except where a row is labeled before attribution. The withdrawn-capacity row prices the control corridor's own coefficient across the 144 cities with control-lane exposure above a twentieth of a standard deviation and enters no balance point. The balance points divide the priced costs by the exposure removed. Coverage is the counted diversion and the coastal cities of the incidence account. Unit of observation is the flow. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S61: The canal authority's transits and tonnage after the closure

	Canal authority net tonnage	Satellite AIS deadweight
Transits	-40.3 percent	-35.6 percent
Tonnage crossing	-60.5 percent	-50.7 percent
Tonnage per ship	-34.1 percent	-23.5 percent

Notes. Change from the 2019-2023 mean to the 2024-2025 mean. Net tonnage is Suez Canal Net Tonnage, taken from Table 1 of the canal authority's annual navigation reports, and it is the base the toll is levied on. Deadweight capacity and transits are from satellite AIS at the same chokepoint. Coverage is the 51 years the authority reports and the 2019-2026 window of the AIS series. Unit of analysis is the year. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S62: The fuel and days the closure added to a voyage

	Value	
Route through the canal, traffic weighted	9,329	nautical miles
Route around the Cape, traffic weighted	13,463	nautical miles
Extra distance per voyage, traffic weighted	4,133	nautical miles
Transits a day at the Cape before the closure	47.4	ships
Transits a day after	87.7	ships
Additional transits a day	40.3	ships
Days since the closure	969	days
of which lost at Bab el-Mandeb	26.3	ships a day
Additional sailing	161	million ship-nautical-miles
attributed to the Red Sea closure	105	million ship-nautical-miles
Additional sailing, capacity weighted	10.2	10^{12} deadweight-nautical-miles
Extra fuel a voyage burns on the detour	611	tonnes
Extra days a voyage takes at 16 knots	10.8	days

Notes. The extra distance is the difference between the two solved states, averaged over the crossing lanes with the container weights. Transit counts and deadweight capacity are as reported at the Cape of Good Hope. Coverage is December 2023 to the end of the transit series in July 2026. Unit of analysis is the chokepoint by day. The speed and the fuel burn are the ones the rest of the paper uses, and no price enters. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S63: The representative voyage under the European scheme

	Through the canal	Around the Cape
Distance, nautical miles	9,329	13,463
Days at 16 knots	24.3	35.1
Fuel burned at the diverted hull's rate	1,379	1,990
Carbon dioxide, tonnes	4,295	6,198
Covered share, one end in the Union	50 percent	50 percent
Cost at the 2024 allowance price, euros	139,147	200,808

Notes. The representative Asia to North Europe voyage on each routing, at the traffic-weighted solved distances, a service speed of 16 knots, the burn rate of the container hull that actually diverted, 611 tonnes over its 4,133 mile average detour, 3.114 tonnes of carbon dioxide a tonne of fuel, and the 2024 average allowance price of 64.8 euros. A voyage with one end outside the Union surrenders allowances on half its emissions. Unit of analysis is the voyage. The figures are computed rather than estimated, so no standard errors are carried.

Table S64: The identification checks on one page

What the reading changes	Estimate		In full
A different reading of the lane	0.029***	(0.006)	S65
No city trend	0.018***	(0.005)	S69
Trend fitted before the closure and projected	0.036***	(0.005)	S69
City and month effects alone	0.008***	(0.003)	S71
The canal's own country out	0.034***	(0.005)	S71
The port crisis coast out	0.029***	(0.005)	S80
New capacity on the gained side out	0.034***	(0.005)	S83
The fast and the weather entered	0.033***	(0.006)	S84
Off the trading scheme's coasts	0.046***	(0.009)	S87
Scheme-boundary transshipment hubs out	0.036***	(0.006)	S88
Land pixels only	0.033***	(0.006)	S91
Every city the sensor covers, no thinning	0.035***	(0.003)	S73
City by calendar month absorbed	0.026***	(0.004)	S26
The two pandemic years out	0.033***	(0.005)	S13
Errors on a one-degree cell	0.035***	(0.006)	S73
The most exposed tenth out	0.041***	(0.007)	S104
Exposure replaced by its rank	0.033***	(0.005)	S104
Two periods, one mean each side	0.018***	(0.005)	S76
The corridor from calls of every hull class	0.030***	(0.005)	S105
The corridor graded by which end emptied	0.023***	(0.005)	S105
Within one country, Spain	0.018***	(0.004)	S71
Within one country, Morocco	0.022**	(0.009)	S71

Notes. Each row is the headline switch of Table 3 re-estimated against one threat, with the appendix table that carries it in full in the last column. Every number is read out of the table it comes from rather than typed here. Standard errors in parentheses, clustered on the city except in the row that widens them to a one-degree cell. Coverage and unit of analysis are as in Table 3 except where the row says otherwise.

S5.6 Specification of the Estimating Equation

Table S65: The estimate under three readings of what a lane is

	(1) Line integral	(2) Nearest approach	(3) Integral, wind weighted
Exposure × Post	0.035***	0.029***	0.031***
	(0.005)	(0.006)	(0.007)
Mean of dep. var.	-10.745	-10.745	-10.745
Observations	38,160	38,160	38,160
Adjusted R^2	0.944	0.944	0.944

Notes. The line integral sums the kernel over every path step and is the construction the paper uses. The nearest-approach column prices a lane by its closest point, and the wind column scales the integral by the climatological share of hours the air moves from the lane toward the city, doubled so a coin-flip wind recovers the isotropic scale. Coverage, unit of analysis and clustering are as in Table 3.

Table S66: The estimate across the specification grid

Specifications run	2,880
Median estimate	0.0393
10th and 90th of the estimates	0.0161-0.0533
Smallest and largest	0.0087-0.0721
Share with the sign of the reported estimate	1.000
Share also distinguishable from zero	0.941

Notes. The distribution of the estimate over the combinations drawn in Figure S12. Coverage is the 2,880 specifications of Figure S12, with the clustering choice already among the dimensions. Unit of observation is the specification. Each estimate carries its own city-clustered error.

Table S67: The three specification choices, every combination

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Exposure × Post	0.019***	0.025***	0.017***	0.023***	0.033***	0.038***	0.026***	0.033***
	(0.005)	(0.004)	(0.005)	(0.004)	(0.005)	(0.006)	(0.004)	(0.005)
City trend	No	No	No	No	Yes	Yes	Yes	Yes
Month × exposure	No	No	Yes	Yes	No	No	Yes	Yes
Population weights	No	Yes	No	Yes	No	Yes	No	Yes
Mean of dep. var.	-10.733	-10.733	-10.733	-10.733	-10.733	-10.733	-10.733	-10.733
Observations	27,984	27,984	27,984	27,984	27,984	27,984	27,984	27,984
Adjusted R^2	0.944	0.958	0.944	0.958	0.944	0.958	0.944	0.958

Notes. All eight columns are estimated on the pandemic-excluded window, so nothing differs across them but the three rows, and every column runs on the same cities, since the weight is a population count every city in the sample carries. City and country by month effects are in every column. Coverage is otherwise as in Table 3. Unit of analysis is the city by month. Standard errors are clustered on the city.

Table S68: Curvature, the dual coasts, the retrieval's own breaks

	(1)	(2)	(3)	(4)	(5)
Exposure × Post	0.046***	0.030***	0.035***	0.028***	0.032**
	(0.008)	(0.007)	(0.007)	(0.008)	(0.016)
Neighbors' exposure × Post					0.003
					(0.019)
Variant	Spain and Morocco out	From 9/2019	From 1/2020	City quadratic trends	Neighbors entered
Mean of dep. var.	-10.745	-10.740	-10.740	-10.745	-10.724
Observations	33,660	34,768	33,072	38,160	36,990
Adjusted R^2	0.946	0.944	0.944	0.944	0.943

Notes. The outcome is the log nitrogen dioxide column over the city box. Column (1) drops Spain and Morocco, the two coasts that sit beside both corridors at once. Columns (2) and (3) open the window after the instrument's along-track pixel shrank in August 2019 and after the calendar turned to 2020. Column (4) adds a city-specific quadratic beside the linear trend. Column (5) enters the mean exposure of the other coastal cities within 300 kilometers of each city, interacted with the closure, beside the city's own. Coverage, unit of analysis and clustering are as in Table 3.

Table S69: Every reading of the city trend

Reading	Estimate	What identifies it
No city trend	0.018	City and country by month effects only
City and month effects alone	0.008	No country by month effects either
The event study, which carries no trend	0.046	Mean of the post-period half-year bins
Trend fitted on the pre-period and projected	0.036	No post-treatment month informs the line
Trend fitted on the whole window	0.035	The reading the paper reports

Notes. Each row varies only how the city-specific trend enters, with the balance test that motivates the trend in Table S96. The row fitted on the pre-period alone lets no post-treatment month inform the line the break is read against. Coverage, unit of analysis and clustering are as in Table 3, except that the event-study row is read on the window that removes 2020 and 2021 and against a base half-year one year before the attacks rather than against the whole pre-period, so little of the differential trend accumulates into it.

Table S70: The trend the correction removes, over fifteen earlier years

	City log-column trend, 2005-2018, a year
Exposure, standardized	-0.0031
	(0.0020)
Fixed effects	Country
Mean of dep. var.	0.0064
Observations	417
Adjusted R^2	0.233

Notes. The outcome is each city's own linear trend in the log OMI nitrogen dioxide column fitted over 2005-2018, four months per year at the season's corners, from the daily quarter-degree grids averaged over the same city boxes the panel uses. The regressor is the paper's standardized exposure and the fixed effects are countries, so the coefficient is the within-country association between exposure and a fifteen-year trend that predates the instrument, the closure and the pandemic. Coverage is the estimation sample's cities with at least twenty OMI months. Unit of analysis is the city. Standard errors are clustered at the country level.

Table S71: The fixed effects, and the country the canal runs through

	(1)	(2)	(3)	(4)	(5)	(6)
Exposure × Post	0.008***	0.018***	0.035***	0.034***	0.018***	0.022**
	(0.003)	(0.005)	(0.005)	(0.005)	(0.004)	(0.009)
Country × month FE	No	Yes	Yes	Yes	n.a.	n.a.
City trend	No	No	Yes	Yes	Yes	Yes
Sample	Coastal	Coastal	Coastal	No Egypt	Spain only	Morocco only
Mean of dep. var.	-10.754	-10.745	-10.745	-10.772	-10.694	-10.783
Observations	38,970	38,160	38,160	36,360	2,250	2,250
Adjusted R^2	0.837	0.944	0.944	0.942	0.910	0.880
Bias-adjusted estimate, selection at delta one			0.049			

Notes. The bias-adjusted row applies the delta-one convention of Oster (2019) at a maximal R-squared of one to Column (3), and spans the table because it is one number rather than a column entry. Column (1) holds only the city and the month. Column (2) adds country by month, Column (3) the city trend. Column (4) drops Egypt. Columns (5) and (6) keep one country at a time, every country with at least four coastal cities on each side. Country by month is not identified inside one country, so month effects replace it in those two columns. Coverage, unit of analysis and clustering are as in Table 3.

Table S72: The population the box rule counts twice

	Millions
Cells the boxes hold, on the global grid	648.9
after each cell is given to one city	472.9
counted more than once	176.0
as a share of the box total	27.1 percent
The account's own weights, box rule	346.6
after the exclusive assignment	275.0
counted more than once	71.6
as a share of the box total	20.7 percent

Notes. A city's population is the gridded count inside a box of 0.35 degrees on each side of its center. The upper block reads the global one-kilometer mosaic over those boxes, first summing every cell inside and then giving each cell to the nearest city center alone. The lower block applies the same assignment to the weights the incidence account carries. Coverage is the 507 coastal cities the account enumerates. Unit of analysis is the city. The counts are enumerated rather than sampled, so the table carries no standard errors.

Table S73: Every city the sensor covers

	(1)	(2)	(3)	(4)
Exposure × Post	0.035***	0.035***	0.035***	0.036***
	(0.005)	(0.003)	(0.006)	(0.004)
Cities entering	Thinned	Every city	Every city	Only the removed
Clustered on	City	City	One-degree cell	City
Mean of dep. var.	-10.745	-10.575	-10.575	-10.451
Observations	38,160	95,760	95,760	56,790
Adjusted R^2	0.944	0.960	0.960	0.965

Notes. The thinning rule keeps one city of any pair whose retrieval boxes overlap, which the incidence account needs so that no air is counted into two populations. A regression needs no such rule, since overlapping boxes make neighbors' errors correlated rather than double counted. Coverage is every city the sensor covers within 100 km of the coast, outside the war coasts. Column (2) therefore adds the coastal cities the rule removes, whose columns were downloaded for this, column (3) clusters on a one-degree cell so that every box that can overlap another sits in one cluster, and column (4) estimates on the removed cities alone. Unit of analysis is the city by month.

Table S74: The dose counted at each city's own gate

	(1)	(2)	(3)	(4)	(5)
Exposure × Post	0.035***				
	(0.005)				
Exposure × Corridor traffic gone		0.052***			
		(0.006)			
Exposure × Own-gate traffic gone			0.047***	0.069***	0.059***
			(0.007)	(0.018)	(0.008)
Shock entered as	Switch	One corridor count	Own-gate count	Own-gate Mediterranean	Own-gate diversion coasts
Mean of dep. var.	-10.745	-10.745	-10.745	-10.489	-10.945
Observations	38,160	38,160	38,160	12,870	15,930
Adjusted R^2	0.944	0.944	0.944	0.929	0.932

Notes. A Mediterranean city sees what comes through Suez and a Red Sea city sees what comes through Bab el-Mandeb, and the two counts fell by different amounts, so column (3) builds the dose from the gate governing each city's own water rather than from one corridor-wide share. Columns (4) and (5) read that dose on the Mediterranean and on the coasts the diverted ships took. Coverage, unit of analysis and clustering are as in Table 3.

S5.7 Inference and Clustering

Table S75: Standard errors under dependence in space and in time

	(1)	(2)	(3)	(4)	(5)	(6)
Exposure × Post	0.035***	0.035***	0.035***	0.035***	0.035***	0.035***
	(0.005)	(0.007)	(0.008)	(0.008)	(0.012)	(0.007)
Standard errors	City	Conley, 100 km	Conley, 250 km	Conley, 500 km	DK, 6 lags	DK, 12 lags
Mean of dep. var.	-10.745	-10.745	-10.745	-10.745	-10.745	-10.745
Observations	38,160	38,160	38,160	38,160	38,160	38,160
Adjusted R^2	0.944	0.944	0.944	0.944	0.944	0.944

Notes. Columns (2) to (4) use Conley (1999) errors, which let the covariance fall with distance to a stated cutoff, taken below, at and well beyond the kernel's own decay length. The effects are swept out first, so the coefficient is that of Table 3 and the columns differ only in the error. Columns (5) and (6) use the covariance of Driscoll and Kraay (1998), written DK in the row, which allows arbitrary cross-sectional dependence in every month and a kernel in time, at six and twelve monthly lags. Coverage and unit of analysis are as in Table 3.

Table S76: Coarser errors and the two-period collapse

	(1)	(2)	(3)	(4)
Exposure × Post	0.035***	0.035***	0.035***	0.018***
	(0.005)	(0.008)	(0.008)	(0.005)
Errors	City	5 degree cells	10 degree cells	Two periods
Mean of dep. var.	-10.745	-10.745	-10.745	-10.739
Observations	38,160	38,160	38,160	848
Adjusted R^2	0.944	0.944	0.944	0.997

Notes. Columns (2) and (3) cluster on cells of five and ten degrees. Column (4) collapses every city to one pre-closure and one post-closure mean before estimating, which is the answer to serial correlation in a persistent outcome. A collapse carries no within-window trend, so its estimate sits where the no-trend cells of Table S67 sit. Coverage and unit of analysis are otherwise as in Table 3.

Table S77: Inference against randomly rearranged coastlines

The estimate	0.035
Spread of the estimate across randomized coastlines	0.006
Draws	1,000
Draws whose estimate is at least as large	<0.001
Draws whose ratio to its own error is at least as large	<0.001
Rotations of the exposure vector along each coastline	1,000
spread of the estimate across them	0.011
rotations at least as large	<0.001

Notes. Each draw permutes the exposure values among the coastal cities of the same country and re-estimates equation (8) unchanged with errors clustered on the city, so no draw of the thousand reaches the estimate. Seventeen of the 53 countries hold fewer than three coastal cities and their labels barely move under permutation, which sharpens nothing, since the estimate rests on the within-country contrasts of the multi-city coasts. The lower block rotates the whole exposure vector along each country's coastline instead, which keeps contiguous runs contiguous and so preserves the spatial structure a permutation destroys, and no rotation of the thousand reaches the estimate either. Coverage, unit of analysis and specification are as in Table 3.

Table S78: The dose, read as a binary contrast and against rotated paths

	(1)	(2)
Top tercile × Post	0.051***	
	(0.009)	
Exposure × Share of the corridor gone		0.052***
		(0.006)
Reading	Binary, top against bottom tercile	Continuous dose, as reported
Rotations of the dose path, as large as the estimate		1 of 90
Mean of dep. var.	-10.823	-10.745
Observations	25,290	38,160
Adjusted R^2	0.932	0.944

Notes. Column (1) reads the design as top tercile of exposure against bottom, the bound the continuous-treatment critique asks for. Column (2) is the continuous dose as reported. The rotation row re-estimates the dose column at every circular offset of the dose path through the ninety-month window, and counts those returning a coefficient as large as the estimate. Coverage, unit of analysis and clustering are as in Table 3.

S5.8 Named Confounders

Table S79: The named threats, one at a time

	(1)	(2)	(3)	(4)
Exposure × Post	0.035***	0.034***	0.035***	0.036***
	(0.005)	(0.006)	(0.005)	(0.005)
Exposure × Ocean index		0.001		
		(0.002)		
Port scale × Post			-0.006	
			(0.006)	
Variant	As reported	Ocean state in	Port scale in	Pre-fitted trends
Mean of dep. var.	-10.745	-10.745	-10.749	-0.011
Observations	38,160	38,160	37,620	38,160
Adjusted R^2	0.944	0.944	0.944	0.707

Notes. Column (2) adds the Oceanic Niño Index interacted with exposure. Column (3) adds the log of pre-period port calls interacted with the post indicator. Column (4) fits each city's trend on the pre-period alone with the pandemic years excluded, projects it forward, and estimates on the residual with no within-window trend at all. Coverage, unit of analysis and clustering are as in Table 3.

Table S80: The South African port crisis

	(1)	(2)	(3)	(4)
Exposure × Post	0.035***	0.029***	0.033***	0.035***
	(0.005)	(0.005)	(0.005)	(0.006)
Sample	All	South Africa out	Durban out	Cape Town out
Mean of dep. var.	-10.745	-10.736	-10.745	-10.745
Observations	38,160	36,630	38,070	38,070
Adjusted R^2	0.944	0.948	0.944	0.944

Notes. South African ports queued through the closing months of 2023 for reasons of their own, and a ship at anchor emits beside the city. The columns drop the country and its two largest port cities one at a time. Coverage, unit of analysis and clustering are as in Table 3.

Table S81: Each side without its ports

	(1)	(2)	(3)	(4)	(5)	(6)
Exposure × Post	0.045***	0.039**	0.049***	0.038**	0.053*	0.034
	(0.009)	(0.015)	(0.012)	(0.016)	(0.027)	(0.021)
Side of the shock	Lost	Lost	Lost	Gained	Gained	Gained
Cities entering	Every city	No port	Port cities	Every city	No port	Port cities
Mean of dep. var.	-10.401	-10.440	-10.352	-11.100	-11.027	-11.192
Observations	19,530	7,650	11,250	18,540	10,620	7,110
Adjusted R^2	0.942	0.928	0.948	0.889	0.859	0.923

Notes. Equation (8) run on each side of the reallocation and then on the cities of that side that own a port and those that do not. The split is exhaustive, but the two halves do not add back to the side, because a country left with one city inside a half is absorbed whole by the country by month effects and drops out of that column. Coverage, unit of analysis and clustering are as in Table 3.

Table S82: Egypt, the corridor and the resorts

	(1)	(2)	(3)
Exposure × Post	0.035***	0.032***	0.034***
	(0.005)	(0.005)	(0.005)
Sample	All coasts	Egypt's Red Sea resorts out	Egypt out
Mean of dep. var.	-10.745	-10.747	-10.772
Observations	38,160	37,710	36,360
Adjusted R^2	0.944	0.944	0.942

Notes. Egypt holds the most exposed cities on the corridor that emptied and carries a second shock on the same date, since Red Sea tourism collapsed from December 2023. The resorts that lost it sit on the Red Sea coast rather than the Mediterranean one, so Column (2) drops the Egyptian cities facing the Red Sea and Column (3) drops Egypt altogether. Coverage, unit of analysis and clustering are as in Table 3.

Table S83: The capacity that came on stream where the ships arrived

	(1)	(2)	(3)
Exposure × Post	0.035***	0.034***	0.037**
	(0.005)	(0.005)	(0.016)
Sample	All coasts	Start-up cities out	Gained side, start-ups out
Mean of dep. var.	-10.745	-10.740	-11.099
Observations	38,160	37,710	18,090
Adjusted R^2	0.944	0.944	0.886

Notes. Hydrocarbon capacity came on stream on the corridor that filled over the same window, the Dangote refinery at Lekki from January 2024, first oil at Sangomar off Senegal from June 2024 and drilling in the Orange Basin off Walvis Bay. Each would raise the column at one city for a reason that is not the ships, so the five cities carrying them or sitting inside their plume are dropped. Coverage, unit of analysis and clustering are as in Table 3.

Table S84: The moving month and the weather

	(1)	(2)	(3)	(4)
Exposure × Post	0.035***	0.033***	0.034***	0.033***
	(0.005)	(0.005)	(0.006)	(0.006)
Exposure × Ramadan month		0.036***		0.035***
		(0.008)		(0.008)
Control	None	Ramadan × exposure	Temperature and rain	Both
Mean of dep. var.	-10.745	-10.745	-10.742	-10.742
Observations	38,160	38,160	37,980	37,980
Adjusted R^2	0.944	0.944	0.944	0.944

Notes. The Ramadan month is the calendar month holding the majority of each year's fast, from the published lunar calendar, interacted with exposure. Temperature and rain are the reanalysis land monthly means over each city box. Coverage, unit of analysis and clustering are as in Table 3.

Table S85: Ramadan on the pre-closure months, before the traffic moved

	log NO_2, pre-closure months
Exposure × Ramadan	0.016*
	(0.010)
Mean of dep. var.	-10.746
Observations	14,840
Adjusted R^2	0.942

Notes. The outcome is the log nitrogen dioxide column on the pre-closure months only, January 2019 to November 2023 with the pandemic years removed. The regressor interacts exposure with the months holding a majority of Ramadan days, May 2019 and April of 2022 and 2023. City and country by month effects and the calendar month interacted with exposure included, the same seasonal absorption every sub-annual design in the paper carries. Coverage and unit of analysis are as in Table 3. Standard errors are clustered at the city level.

Table S86: Other national shocks against the corridor

	(1)	(2)	(3)
Exposure × Post	0.034***	0.030***	0.030***
	(0.005)	(0.006)	(0.006)
Exposure × Eid al-Adha month	-0.017		-0.017
	(0.015)		(0.015)
Exposure × Egypt after the pound, 3/2024		0.101***	0.102***
		(0.018)	(0.018)
Shock entered	Eid al-Adha	The pound	Both
Mean of dep. var.	-10.745	-10.745	-10.745
Observations	38,160	38,160	38,160
Adjusted R^2	0.944	0.944	0.944

Notes. The outcome is the log nitrogen dioxide column over the city box. The Eid al-Adha month is the calendar month holding the feast in each year, which advances through the calendar the same way Ramadan does. The pound row interacts exposure with an indicator for Egyptian cities from March 2024, the month the Egyptian pound was floated, so any break specific to exposed Egyptian cities at that date is absorbed rather than attributed to the corridor. Coverage, unit of analysis and clustering are as in Table 3.

Table S87: The closure against the carbon scheme that began a month later

	(1)	(2)	(3)
Exposure × Post	0.046***	0.046***	0.051***
	(0.009)	(0.009)	(0.012)
Exposure × Post × European Union coast	-0.025***		
	(0.010)		
Sample	All	Union coasts out	Lost corridor, no Union coast
Mean of dep. var.	-10.745	-10.741	-10.322
Observations	38,160	32,580	14,940
Adjusted R^2	0.944	0.946	0.943

Notes. The scheme charges voyages that touch Union ports from January 2024, so its footprint is the Union coast and the closure's is the corridor. Column (1) interacts exposure with a Union-coast indicator, Column (2) drops Union coasts, and Column (3) keeps only lost-corridor cities on coasts the scheme does not touch. Coverage, unit of analysis and clustering are as in Table 3. Exposure varies as much within the scheme's coasts as outside them, a standard deviation of 0.97 against 0.96.

Table S88: The transshipment hubs at the scheme boundary

	(1) log NO_2	(2) log NO_2	(3) log calls	(4) log calls	(5) log tanker calls	(6) log tanker calls
Exposure × Post	0.035***	0.036***	0.054***	0.054***	0.032	0.032
	(0.005)	(0.006)	(0.019)	(0.019)	(0.020)	(0.020)
Ports removed	None	Boundary hubs	None	Boundary hubs	None	Boundary hubs
Mean of dep. var.	-10.745	-10.749	0.684	0.684	0.270	0.270
Observations	38,160	37,710	5,214	5,214	5,214	5,214
Adjusted R^2	0.944	0.944	0.978	0.978	0.948	0.948

Notes. The even columns remove the five transshipment hubs sitting just outside the boundary of the European trading scheme. Coverage for the first two columns is the coastal sample of Table 3. The last four columns are the gained-side cities owning a port, none of which is one of the five hubs, so the removal leaves those columns where they stood. Unit of analysis is the city by month and standard errors are clustered on the city throughout.

Table S89: The latitude band

	(1)	(2)
Exposure × Post	0.035***	0.073***
	(0.005)	(0.021)
Sample	As reported	South of 20 south
Mean of dep. var.	-10.745	-11.154
Observations	38,160	2,880
Adjusted R^2	0.944	0.850

Notes. Column (2) keeps cities south of twenty degrees south, the extratropical reach of the corridor that gained the traffic, where the air carries neither a lightning belt nor a biomass season. Coverage, unit of analysis and clustering are as in Table 3.

Table S90: The satellite column against instruments on the ground

Swept out	Slope	s.e.	Correlation	Observations
City effects	0.902	(0.046)	0.584	749
City and month-of-year effects	0.547	(0.059)	0.327	749
City and month effects	0.292	(0.061)	0.185	749
City and month-of-year effects, side that gained	0.507	(0.106)	0.350	178
City and month-of-year effects, side that lost	0.710	(0.071)	0.393	571

Notes. The outcome is the log monthly mean nitrogen dioxide at street level from the European monitoring network, read against the log satellite column in the same city-month. Coverage is the nine cities of the estimation sample carrying a monitor, in Spain, Portugal, Italy, Malta and Greece, over January 2019 to June 2026, read on both sides of the corridor. Unit of analysis is the city by month. Standard errors are conventional, each slope re-read on a restricted wild cluster bootstrap by city.

Table S91: The column read over land pixels alone

	(1)	(2)	(3)
Exposure × Post	0.035***	0.033***	0.032***
	(0.005)	(0.006)	(0.006)
Column read over	Whole box	Land pixels only	Land, passes held
Mean of dep. var.	-10.745	-10.714	-10.714
Observations	38,160	37,980	37,980
Adjusted R^2	0.944	0.947	0.947

Notes. The outcome in Columns (2) and (3) is the tropospheric column averaged over the land pixels of the box alone, the permanent-water class of the 10 m WorldCover land-cover map removed under the 7 km retrieval footprint. Column (3) adds the log count of valid retrievals in the city-month, since the two corridors sit under different skies and a drift in retrieval frequency would read as a trend. Coverage, unit of analysis and clustering are as in Table 3.

Table S92: The response on the level column, by Poisson

	(1)	(2)
	NO_2 column, level, by Poisson	
Exposure × Post	0.022***	0.044***
	(0.006)	(0.006)
City trends	No	Yes
Mean of dep. var.	2.62e-05	2.62e-05
Observations	38,160	38,160
Pseudo R^2	0.018	0.018

Notes. The outcome is the nitrogen dioxide column in levels and the estimator is Poisson pseudo-likelihood, so no log is taken and no cell could fall out of the sample by turning non-positive. The panel is balanced, 424 cities in each of 90 months with no missing and no non-positive city-month, so log selection has nothing to select on and this table is the demonstration. Coverage and unit of analysis are as in Table 3. Standard errors are clustered at the city level.

Table S93: No coast carries the estimate

	(1)	(2)	(3)	(4)	(5)	(6)
Exposure × Post	0.035***	0.033***	0.034***	0.045***	0.043***	0.029***
	(0.005)	(0.005)	(0.006)	(0.009)	(0.007)	(0.005)
Coast removed	None	Largest coast	Largest three	Largest five	Removal raising it most	Removal lowering it most
Mean of dep. var.	-10.745	-10.770	-10.823	-10.837	-10.748	-10.736
Observations	38,160	35,370	30,150	25,650	35,910	36,630
Adjusted R^2	0.944	0.944	0.941	0.943	0.945	0.948

Notes. Equation (8) with coastlines removed. Columns (2) to (4) drop the countries holding the most coastal cities in the sample, singly and together. Columns (5) and (6) drop the country whose removal moves the estimate furthest in each direction, found by estimating with each of the 53 countries out in turn. Country by month effects mean nothing national identifies the coefficient, and 0.272 of the variance of exposure survives country demeaning. Coverage, unit of analysis and clustering are as in Table 3.

Table S94: The design inside one eastern country

	(1)	(2)	(3)
	log NO_2, Indonesian coastal cities		
Exposure × Post	-0.022***	-0.004	
	(0.004)	(0.004)	
Sunda side × Post			0.001
			(0.008)
City trends	No	Yes	Yes
Mean of dep. var.	-11.168	-11.168	-11.179
Observations	22,836	22,836	21,186
Adjusted R^2	0.938	0.940	0.941

Notes. The outcome is the log nitrogen dioxide column over Indonesian coastal cities within 100 kilometers of the sea, with city and month effects. Exposure is the solved change in the 150 kilometer lane kernel between the two states, standardized within Indonesia. The sample keeps the Sunda-side gains and the neutral cities and excludes the Malacca-approach coast. Column (3) compares the Sunda side against the neutral cities. The pandemic years are removed. Unit of analysis is the city by month. Standard errors are clustered at the city level.

S5.9 Placebos and the Measurement of Exposure

Table S95: The smallest effect each null could have detected

Null	Estimate	s.e.	Upper limit	Detectable	Compared with	Verdict
Upwind passes	0.014	0.008	0.030	0.022	0.035	Bounds it below
The corridor placebo, four years early	0.006	0.008	0.022	0.022	0.035	Bounds it below
Unchanged lanes, entered beside the real ones	-0.001	0.004	0.007	0.011	0.035	Bounds it below
Cities 200-400 km inland	0.025	0.035	0.094	0.098	0.035	Underpowered
Formaldehyde	0.003	0.005	0.013	0.014	0.035	Bounds it below
Tanker calls where ships stop, gained side	0.022	0.013	0.047	0.036	0.029	Underpowered
Sulfur under the control area	0.069	0.056	0.179	0.157	0.178	Bounds it below
Dry bulk, the shallow chambers	0.038	0.061	0.158	0.171	0.070	Underpowered
The Panama placebo, two years early	0.016	0.018	0.051	0.050	0.070	Bounds it below
Tanker calls, count form	0.075	0.050	0.173	0.140	0.102	Underpowered
Cargo, logged	-0.016	0.130	0.239	0.364	0.102	Underpowered

Notes. Coverage is every null the paper reports. The upper limit column is the 95 percent confidence bound on the estimate. The detectable column is the effect a test at five percent with eighty percent power would have rejected zero for, which is 2.80 times the standard error, and the column beside it is the treatment effect the same design finds elsewhere. A null whose detectable effect exceeds the treatment found elsewhere reads Underpowered, and one that falls below it reads Bounds it below. Unit of analysis and clustering as in the table each row is drawn from.

Table S96: The corridor weights against the pre-period level and trend

	(1) Pre-period level	(2) Pre-period trend	(3) Pre-period trend
Red Sea corridor weight	0.077	0.007***	
	(0.083)	(0.002)	
Cape corridor weight	-0.058	0.001	
	(0.041)	(0.003)	
Exposure			-0.005**
			(0.002)
Mean of dep. var.	-10.758	0.018	0.018
Observations	424	424	424
Adjusted R^2	0.683	0.499	0.491

Notes. Each city's pre-period level and its own pre-period trend, regressed on the standardized corridor weights and on the exposure the design uses, with country effects. City effects absorb the level by construction, so the object the design needs balanced is the trend the estimating equation nets out. Coverage is the 424 cities of the panel sample of 433 whose pre-period supports a fitted trend, the identified sample every trend-carrying specification estimates on. Unit of analysis is the city. Standard errors are clustered on the country.

Table S97: The shock moved to two clean Decembers

	(1)	(2)
Exposure × Placebo post	0.006	0.002
	(0.008)	(0.005)
Placebo date	12/2019	12/2022
Mean of dep. var.	-10.758	-10.758
Observations	25,016	25,016
Adjusted R^2	0.944	0.944

Notes. Each column moves the shock to a December before the real one and discards every month from the real one onward, so the exposure measure keeps the coastal variation it is built on and the date is the only thing that moves. Coverage, unit of analysis and clustering are as in Table 3.

Table S98: The lanes whose traffic had no reason to move

	(1) log NO_2	(2) log NO_2	(3) log NO_2
Exposure × Post	0.035***		0.034***
	(0.005)		(0.007)
Unchanged-lane exposure × Post		-0.013***	-0.001
		(0.002)	(0.004)
Sample	Coastal	Coastal	Coastal
Mean of dep. var.	-10.745	-10.745	-10.745
Observations	38,160	38,160	38,160
Adjusted R^2	0.944	0.944	0.944

Notes. Unchanged-lane exposure is built exactly as corridor exposure is, from the thirty lanes joining ports on the same side of the canal, whose solved distance is identical in both states of the world. Coverage, unit of analysis and clustering are as in Table 3.

Table S99: The estimate on the lanes whose distance did not change

	(1)	(2)	(3)
Unchanged-lane exposure × Post	-0.013***	-0.019**	-0.014
	(0.002)	(0.009)	(0.011)
Exposure × Post			0.015
			(0.011)
Outcome	log NO_2 column	log port calls	log port calls
Mean of dep. var.	-10.745	0.380	0.380
Observations	38,160	27,984	27,984
Adjusted R^2	0.944	0.981	0.981

Notes. Exposure built from the thirty lanes joining ports on the same side of the canal, whose solved distance the closure did not change. The diversion absorbed fleet capacity and carriers canceled sailings on strings whose distance had not moved, so the measure carries the capacity margin rather than the plume. Column (1) covers January 2019 to June 2026 and the call columns leave out 2020 and 2021, as the activity tables do. Coverage, unit of analysis and clustering are as in Table 3.

Table S100: The core readings with the Hormuz months out

Reading	Coefficient	s.e.	Obs.	Adj. R^2
Switch, full window	0.035***	(0.005)	38,160	0.944
Switch, window ending February 2026	0.034***	(0.005)	36,464	0.944
Dose, full window	0.101***	(0.027)	38,160	0.944
Dose, window ending February 2026	0.101***	(0.027)	36,464	0.944
Trade term, full window	-0.007*	(0.004)	6,708,154	0.839
Trade term, window ending February 2026	-0.007*	(0.004)	6,572,007	0.839
Atlantic share, full window	-0.036***	(0.013)	2,576	0.921
Atlantic share, window ending February 2026	-0.034***	(0.013)	2,519	0.922

Notes. Y is the log of the nitrogen dioxide column for the switch and dose rows, over coastal cities within 100 kilometers of the shore outside the war coasts, the log of bilateral import value for the trade rows, and the Atlantic share of containerized imports for the origin rows. The dose divides the corridor share gone by one hundred, so its coefficient reads per whole corridor. Each pair runs the full window and the window ending February 2026. Errors are clustered by city, by pair and by origin.

Table S101: The dose on the window the rule assigns it

	(1)	(2)	(3)	(4)
Exposure × Share of the corridor gone	0.052***	0.051***		
	(0.006)	(0.006)		
Exposure × Post			0.035***	0.033***
			(0.005)	(0.005)
Window	2019-2026	Pandemic years out	2019-2026	Pandemic years out
Mean of dep. var.	-10.745	-10.733	-10.745	-10.733
Observations	38,160	27,984	38,160	27,984
Adjusted R^2	0.944	0.944	0.944	0.944

Notes. The paper's window rule excludes the two pandemic years from anything read at the half-year or finer, and the continuous dose is monthly, so it belongs on the excluded window. Transits fell in 2020 and urban nitrogen dioxide fell nearly everywhere in 2020, which is the co-movement the rule exists to remove. The switch is shown beside it on both windows. Coverage, unit of analysis and clustering are as in Table 3.

Table S102: The estimate as the post period's first month moves

	(1)	(2)	(3)	(4)	(5)	(6)
Exposure × Post	0.024***	0.027***	0.035***	0.041***	0.037***	0.032***
	(0.006)	(0.006)	(0.005)	(0.004)	(0.005)	(0.005)
Post begins	10/2023	11/2023	12/2023	1/2024	2/2024	3/2024
Mean of dep. var.	-10.745	-10.745	-10.745	-10.745	-10.745	-10.745
Observations	38,160	38,160	38,160	38,160	38,160	38,160
Adjusted R^2	0.944	0.944	0.944	0.944	0.944	0.944

Notes. Each column redefines the post period at the month in the row and leaves everything else as in Table 3. The first attacks on shipping came in November 2023, the carriers announced their suspensions through December, and the diversion was complete by February. Coverage, unit of analysis and clustering are as in Table 3.

Table S103: The estimate as the coastal band widens

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure × Post	0.036***	0.033***	0.034***	0.035***	0.032***	0.030***	0.029***
	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
Distance to the coast	25 km	50 km	75 km	100 km	150 km	200 km	300 km
Mean of dep. var.	-10.779	-10.755	-10.745	-10.745	-10.747	-10.767	-10.795
Observations	23,670	29,340	33,570	38,160	47,160	54,720	66,690
Adjusted R^2	0.953	0.951	0.947	0.944	0.939	0.935	0.925

Notes. Each column keeps cities within the distance in the row of the coast. Coverage, unit of analysis and clustering are as in Table 3, except for the coastal band each column names.

Table S104: The tail of the exposure measure

	(1)	(2)	(3)	(4)	(5)
Exposure × Post	0.035***	0.038***	0.041***		0.034***
	(0.005)	(0.007)	(0.007)		(0.006)
Exposure rank × Post				0.033***	
				(0.005)	
Exposure entered as	Above the 99th out	Above the 95th out	Above the 90th out	Rank, none out	As estimated
Outcome trimmed	No	No	No	No	At the 1st and 99th
Mean of dep. var.	-10.752	-10.760	-10.768	-10.745	-10.753
Observations	37,710	36,000	34,200	38,160	37,326
Adjusted R^2	0.944	0.943	0.944	0.944	0.938

Notes. The first three columns drop the cities above the percentile in the row. Column (4) replaces exposure by its rank, standardized, which throws the skew away entirely, and Column (5) keeps the measure and trims the outcome instead. Coverage, unit of analysis and clustering are as in Table 3.

Table S105: Three ways of building the corridor

	(1)	(2)	(3)
Exposure × Post	0.035***		
	(0.005)		
Exposure × Post, all hulls		0.030***	
		(0.005)	
Exposure × Post, graded by end			0.023***
			(0.005)
Corridor built from	Container calls, uniform	Calls of every hull class	Graded by which end emptied
Mean of dep. var.	-10.745	-10.745	-10.745
Observations	38,160	38,160	38,160
Adjusted R^2	0.944	0.944	0.944
Efficient ratio of the arcs	2.59	3.45	2.59

Notes. Column (2) ranks the endpoint ports on calls of every hull class rather than on container calls and re-solves every lane on that ranking. Column (3) lets the corridor empty by different amounts at its two ends, in the proportion the transits at Suez and at Bab el-Mandeb actually fell, allocated to a city by which end it sits nearer. The last row prices the two arcs on each column's own corridor at the 150 kilometer decay length. The graded rebuild reallocates the observed fall between the gates without touching the static weights, so its ratio is the baseline's by construction. Coverage, unit of analysis and clustering are as in Table 3.

Table S106: The endpoint set rebuilt five ways

	(1)	(2)	(3)	(4)	(5)
Exposure × Post	0.035***	0.038***	0.037***	0.036***	0.035***
	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)
Endpoints	Ports merged by agglomeration	Twenty ports a side	Equal lane weights	Calls over lane distance	Alexandria hub forced in
Efficient arc ratio	2.59	2.14	2.20	2.33	2.41
Mean of dep. var.	-10.745	-10.745	-10.745	-10.745	-10.745
Observations	38,160	38,160	38,160	38,160	38,160
Adjusted R^2	0.944	0.944	0.944	0.944	0.945

Notes. The outcome is the log nitrogen dioxide column under the specification of Table 3, with the exposure measure rebuilt from a different endpoint set each time and re-standardized. Column (1) merges the terminals of one agglomeration before the ranking. Column (2) deepens the ranking to the twenty busiest container ports on each side. Column (3) keeps the baseline lanes and weights each equally instead of by the product of calls at its ends. Column (4) divides the call product by the lane's own open-state distance. Column (5) forces the Alexandria agglomeration into the western endpoint set and re-solves every lane. The arc-ratio row reprices the two arcs on each construction. Coverage, unit of analysis and clustering are as in Table 3.

Table S107: The efficient exposure ratio under each pre-year's own weights

Weights built on	Efficient ratio
The pre-period, 2019 to November 2023	2.59
2019 alone	2.59
2020 alone	2.57
2021 alone	2.58
2022 alone	2.61
2023 alone	2.57

Notes. The lane weights are products of endpoint container calls, and each row rebuilds them from a single pre-year's daily calls with the lane set held fixed, so only the weights move. The calls are fetched from the same PortWatch layer the pre-period average uses. Coverage is the 507 coastal cities of the incidence account at the 150 kilometer decay length. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S108: The decay scale of the plume weight

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure × Post	0.018***	0.025***	0.029***	0.035***	0.038***	0.043***	0.053***
	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)	(0.007)	(0.008)
Decay scale, km	50	75	100	150	200	300	500
Mean of dep. var.	-10.745	-10.745	-10.745	-10.745	-10.745	-10.745	-10.745
Observations	38,160	38,160	38,160	38,160	38,160	38,160	38,160
Adjusted R^2	0.944	0.944	0.944	0.944	0.944	0.944	0.944

Notes. How fast the corridor weight falls with distance from the lane is the one free parameter left in the exposure. The weight is rebuilt at each scale on the grid and standardized each time, so a coefficient is always the effect of one standard deviation of that version. Coverage, unit of analysis and clustering are as in Table 3. The standard-deviation normalization is what carries the magnitude across columns rather than the raw weight.

Table S109: The exposure measure treated as estimated rather than given

	Estimate
At the 150 km length the corridor weights use	0.035
clustered standard error	0.005
standard error the decay length adds	0.003
the two together	0.006
Over the fitted length's own distribution, mean	0.040
its 95 percent range	0.033-0.045
at the fitted 225 km itself	0.040

Notes. The regressor is the output of a solve whose decay length is estimated off the water at 225 kilometers with a standard error of 47. Here the estimate is re-run at each of the seven scales the kernel is built on, a curve is fitted through them, and the length is drawn five thousand times from its own sampling distribution. The spread of the resulting estimates is combined with the clustered sampling error in quadrature. The length is the one construction input drawn. Coverage, unit of analysis and clustering are as in Table 3.

Table S110: The shock as a dose rather than a switch

	(1)	(2)	(3)
Plume, P	0.101***		
	(0.027)		
Cape corridor × transits		-0.071	-0.070
		(0.075)	(0.076)
Suez corridor × transits		0.196***	0.186***
		(0.020)	(0.020)
Sample	Coastal	Coastal	No Egypt
Mean of dep. var.	-10.745	-10.745	-10.772
Observations	38,160	38,160	36,360
Adjusted R^2	0.944	0.944	0.942

Notes. P is the corridor weight multiplied by the transits counted that month, divided by 100, so a unit of the regressor is a hundred weight-transits a month, and Columns (2) and (3) split it by corridor. Coverage, unit of analysis and clustering are as in Table 3.

Table S111: The decay length, read off the water

Decay length fitted to the water, km	225
standard error	47
refitted without the 325-425 km bins	207
Decay length the corridor weights use, km	150
Enhancement on the corridor, μmol m^{-2}	0.405
Ocean cells	41,650

Notes. The change in the tropospheric column is the mean over January 2024 to June 2026 less the mean over January 2019 to November 2023, at 25 km. Cells are binned by distance in 25 km steps and an exponential with a level absorbing the field-wide drift is fitted. Unit of observation is the 25 km ocean cell. Coverage is the ocean cells within 600 km of the closed-state path and clear of the open-state path. The fit carries the standard error of the exponential it estimates.

Table S112: The main estimate under the two-lengths kernel

	(1)	(2)
	log NO_2	
Exposure × Post, common kernel	0.035***	
	(0.005)	
Exposure × Post, corridor-specific kernels		0.022***
		(0.007)
City, country-month FE, city trends	Yes	Yes
Mean of dep. var.	-10.745	-10.745
Observations	38,160	38,160
Adjusted R^2	0.944	0.944

Notes. Y is the log of the nitrogen dioxide column over city-months of the estimation sample. Column (1) is the baseline, exposure built at the common 150 kilometer length. Column (2) rebuilds exposure with each state's paths read at its own fitted length, the closed state's water at 225 kilometers and the open state's at 31, each exposure standardized by its own standard deviation over the coastal band. Both columns absorb city, country by month and city trend effects. Standard errors are clustered at the city level. Unit of analysis is the city-month.

Table S113: The mixing depth, measured rather than acknowledged

	(1)	(2)	(3)	(4)
Exposure × Post	0.035***	0.035***	0.034***	
	(0.005)	(0.005)	(0.006)	
Exposure × Post × log mixing depth			0.026	
			(0.022)	
Exposure × Post, shallow-layer half				0.032***
				(0.008)
Exposure × Post, deep-layer half				0.037***
				(0.006)
log mixing depth		0.025	0.028	0.025
		(0.022)	(0.022)	(0.022)
Mixing depth	Out	Control	Interacted	Split at the median
Mean of dep. var.	-10.745	-10.745	-10.745	-10.745
Observations	38,160	38,160	38,160	38,160
Adjusted R^2	0.944	0.944	0.944	0.944

Notes. The outcome is the log nitrogen dioxide column over the city box. The mixing depth is the ERA5 monthly boundary-layer height averaged over the same box, entered in logs, demeaned where it interacts with the treatment, and split at the median of each city's own mean in the last column. Coverage, unit of analysis and clustering are as in Table 3.

S5.10 The Trade Cost of a Chokepoint

Table S114: The Red Sea trade elasticity, under its own checks

Panel A. The fixed-effect ladder

	(1)	(2)	(3)	(4)	(5)
Distance × Post	-0.006*	-0.006*	-0.007*	-0.007	-0.010***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)
Month FE by	Month	Exp., imp.	Exp., imp.	Exp., imp.	Exp., imp.
Product × month FE	No	No	Yes	Yes	Yes
Clustered on	Pair	Pair	Pair	Exporter	Pair
Window	Full	Full	Full	Full	Ending 2025Q1
Mean of dep. var.	11.190	11.190	11.190	11.190	11.198
Observations	6,708,155	6,708,154	6,708,154	6,708,154	5,812,197
Adjusted R^2	0.836	0.838	0.839	0.839	0.842

Panel B. The pre-period placebo

	log trade value
Added distance × Placebo post	0.001
	(0.003)
Sample	Pre-period
Mean of dep. var.	11.217
Observations	4,640,214
Adjusted R^2	0.846

Notes. Panel A estimates the Red Sea term alone, the added distance the closure imposes on a pair in thousands of nautical miles interacted with December 2023, across the fixed-effect ladder named in its rows. Column (5) ends the window in the first quarter of 2025, before the American tariff round of that spring. Panel B interacts the same distance with a placebo break in the middle of the pre-period, estimated on pre-closure months only. Coverage is the global bilateral chapter-month panel of Table S121. Unit of analysis is the importer by exporter by chapter by month. Standard errors are clustered at the pair level except where the ladder states the exporter.

Table S115: The fixed effects, the sample and the clustering

	(1)	(2)	(3)	(4)
Dependence, 1,000 nm × Draft shortfall	-0.073***	-0.070***	-0.071**	-0.070**
	(0.027)	(0.026)	(0.030)	(0.035)
Product × month FE	No	Yes	Yes	Yes
Importers	All	All	United States dropped	All
Clustered on	Pair	Pair	Pair	Exporter
Mean of dep. var.	10.993	10.993	10.877	10.993
Observations	8,339,180	8,339,180	7,918,960	8,339,180
Adjusted R^2	0.840	0.841	0.835	0.841

Notes. Coverage and unit of analysis are as in Table S121. Errors are clustered as the table's own row states, on the country pair except in the column that clusters on the exporter. Column (3) drops the United States, which is the largest canal-dependent importer in the panel, so the estimate is not one country's.

Table S116: The placebo that moves the draft path two years early

	(1)	(2)
Dependence, 1,000 nm × Draft shortfall	-0.070***	
	(0.026)	
Dependence × Draft shortfall, placebo		0.016
		(0.018)
Draft path dated	Actual	Two years early
Mean of dep. var.	10.993	11.019
Observations	8,339,180	5,813,745
Adjusted R^2	0.841	0.848

Notes. The placebo gives every pair the same draft path dated two years earlier, so geography is held fixed and only the timing moves. The placebo window ends before the drought begins, so it holds fewer months and fewer cells than the baseline by construction, on the baseline's own pairs and chapters. Coverage, unit of analysis and clustering are as in Table S121.

Table S117: Canal dependence under three scalings and under Poisson

	(1)	(2)	(3)	(4)
Dependence, 1,000 nm × Draft shortfall	-0.070***			-0.089*
	(0.026)			(0.048)
Reliance share × Draft shortfall		-0.955***		
		(0.343)		
log dependence × Draft shortfall			-0.210***	
			(0.079)	
Dependence scaled as	Distance, 1,000 nm	Route share saved	log distance	Distance, 1,000 nm
Estimator	Log OLS	Log OLS	Log OLS	Poisson
Mean of dep. var.	10.993	10.993	10.993	8,106,191
Observations	8,339,180	8,339,180	8,339,180	8,339,180
Adjusted R^2	0.841	0.841	0.841	0.978

Notes. Column (2) scales dependence as the share of the shorter route the canal saves rather than as its length, and Column (3) takes logs. Column (4) estimates the baseline scaling by Poisson pseudo-likelihood on the level of trade, so its mean and its fit statistic are read on the level rather than the log. Coverage, unit of analysis and clustering are as in Table S121.

Table S118: The binding shortfall and the estimation window

	(1)	(2)	(3)	(4)
Dependence, 1,000 nm × Transit shortfall	-0.024***		-0.003	
	(0.008)		(0.008)	
Dependence, 1,000 nm × Draft shortfall		-0.070***	-0.066**	-0.047**
		(0.026)	(0.032)	(0.021)
Shortfall measured on	Transits	Draft	Both	Draft
Window	2019-2024	2019-2024	2019-2024	2019-2026
Mean of dep. var.	10.993	10.993	10.993	10.987
Observations	8,339,180	8,339,180	8,339,180	10,031,826
Adjusted R^2	0.841	0.841	0.841	0.837

Notes. Coverage, unit of analysis and clustering are as in Table S121, except Column (4), which runs to June 2026. Monthly Comtrade revises upward in its trailing months, so the final six panel months, all after the restoration, are the least revised.

Table S119: Trends, weights and the errors

	(1)	(2)	(3)	(4)
Dependence, 1,000 nm × Draft shortfall	-0.070***	-0.035***	-0.066**	-0.070**
	(0.026)	(0.013)	(0.026)	(0.027)
Variant	As reported	Pair trends	Value weighted	Pair and month errors
Mean of dep. var.	10.993	10.993	11.002	10.993
Observations	8,339,180	8,339,180	8,320,124	8,339,180
Adjusted R^2	0.841	0.842	0.847	0.841

Notes. Column (2) adds a linear trend for every country pair, so a bilateral drift like the American decoupling from China cannot run through the treatment. Column (3) weights by the log of the cell's pre-period trade, which puts the estimate where the money is without a handful of billion-dollar cells carrying the whole weight. Column (4) clusters on the pair and the month at once. The importer-by-month effects of equation (12) absorb each importer's total demand, so the elasticity prices substitution across suppliers within an importer, which is why the account is denominated in trade displaced rather than destroyed. Coverage and unit of analysis are as in Table S121.

Table S120: The ocean that caused the drought

	(1)	(2)	(3)
Draft cut × canal dependence	-0.070***	-0.071*	-0.084*
	(0.026)	(0.043)	(0.045)
Dependence × Ocean index		0.000	
		(0.002)	
Dependence × Ocean index, warm side only			0.002
			(0.002)
Ocean index	Out	In	Warm side only
Mean of dep. var.	10.993	10.993	10.993
Observations	8,339,180	8,339,180	8,339,180
Adjusted R^2	0.841	0.841	0.841

Notes. The ocean index is the Oceanic Niño Index, published monthly by the Climate Prediction Center, interacted with canal dependence. Column (3) keeps only its warm side, since it is the warm phase that empties the lake. Coverage, unit of analysis and clustering are as in Table S121.

Table S121: Trade on canal dependence and the draft the canal allowed

	(1)	(2)
	log trade value	
Dependence, 1,000 nm × Draft shortfall	-0.069***	-0.070***
	(0.026)	(0.026)
Suez dependence, 1,000 nm × Post	-0.010***	
	(0.004)	
The Red Sea closure	In	Out
Window	2019-2024	2019-2024
Mean of dep. var.	10.993	10.993
Observations	8,339,180	8,339,180
Adjusted R^2	0.841	0.841

Notes. Canal dependence is the sailing distance a pair loses if the canal is unusable, in thousands of nautical miles. The shortfall is one minus the authorized draft over its unconstrained limit. Column (1) enters the Red Sea closure's own pair exposure beside the draft, and Column (2), the specification the paper reports, drops the closure term. The panel is imports and import values include freight. Coverage is monthly imports by 39 reporters from 162 partners across 96 chapters and one residual code. Unit of analysis is the importer by exporter by product by month. Standard errors are clustered on the country pair.

Table S122: The lock the goods cross in

	(1)	(2)	(3)
Dependence, 1,000 nm × Draft shortfall	-0.078***	0.038	-0.080
	(0.027)	(0.061)	(0.083)
Goods that cross in	Containerized	Dry bulk	Liquid bulk
Mean of dep. var.	10.846	11.808	12.490
Observations	7,352,880	569,477	414,885
Adjusted R^2	0.846	0.780	0.802

Notes. The draft that fell is the draft of the Neopanamax chambers, which are the chambers the container ships use, while the dry bulk crosses in the older and shallower Panamax chambers, whose limit the drought did not move. Chapters are assigned to the hull they travel in, containerized in Column (1), dry bulk in Column (2) and liquid bulk in Column (3). Coverage, unit of analysis and clustering are as in Table S121. The chambers rather than the cargo are what the split is on, and detectable effects at eighty percent power are in Table S95.

Table S123: The chamber split read as one contrast

	log trade value
Dependence × Shortfall, containerized	-0.078***
	(0.027)
Dependence × Shortfall, dry bulk	0.038
	(0.061)
Difference, containerized less dry bulk	-0.115*
	(0.061)
95 percent interval on the difference	[-0.236, 0.005]
Containerized, strict dry-bulk mapping	-0.078***
	(0.027)
Dry bulk, strict mapping	0.086
	(0.069)

Notes. Y is the log of bilateral import value over importer by exporter by chapter by month cells through December 2024, split by the lock chamber the goods cross in. Chapters map to chambers by hull class, grain, ores, cement, fertilizer, woodpulp and steel to the Panamax dry bulk chambers and fats, fuels and chemicals to the Neopanamax tankers, with the remainder containerized, and the strict rows reassign woodpulp and steel. The cells and their difference come from one stacked regression with chamber-specific fixed effects and errors clustered by pair.

Table S124: Trade lost against trade delayed

	(1)	(2)	(3)
The month itself	-0.070***		
	(0.026)		
One month after		-0.069**	
		(0.027)	
Mean over the five months			-0.075***
			(0.029)
Shortfall entered at	The month itself	Shifted a month	Five-month window
Mean of dep. var.	10.993	10.993	10.993
Observations	8,339,180	8,339,180	8,339,180
Adjusted R^2	0.841	0.841	0.841

Notes. A ship that cannot get a slot waits or bids for one, and waiting moves a shipment rather than destroying it, which a panel matched to the same month reads as a loss now with the offset never counted. Column (2) shifts the shortfall by a month, since customs clearance is later than sailing by the length of the voyage. Column (3) enters the shortfall as its mean over the two months before the match through the two after, so the window's whole response is one coefficient and an offset from waiting would pull it toward zero. Coverage, unit of analysis and clustering are as in Table S121.

Table S125: The elasticity on the way down and on the way back

	(1)	(2)
Dependence, 1,000 nm × Draft shortfall	-0.070***	
	(0.026)	
on the falling segment		-0.073***
		(0.027)
on the recovering segment		-0.066**
		(0.029)
Shortfall entered	One elasticity	By segment
Mean of dep. var.	10.993	10.993
Observations	8,339,180	8,339,180
Adjusted R^2	0.841	0.841

Notes. The falling segment runs from the first advisory of March 2023 through January 2024 and the recovering one from February 2024, when the lake let the draft back up. Coverage, unit of analysis and clustering are as in Table S121.

Table S126: The canal at the other end of the world

	(1)	(2)
Draft cut × canal dependence	-0.070***	-0.069***
	(0.026)	(0.026)
Suez dependence × Post		-0.010***
		(0.004)
The other canal	Out	In
Mean of dep. var.	10.993	10.993
Observations	8,339,180	8,339,180
Adjusted R^2	0.841	0.841

Notes. Suez dependence is the extra distance a pair's route takes when that canal is shut, solved on the same grid as the Panama measure and in the same units, interacted with December 2023. The two canals were constrained over the same months and the pairs that run between the US East Coast and Asia can use either. Coverage, unit of analysis and clustering are as in Table S121.

Table S127: The cells with no trade in them

	(1) log value	(2) Value	(3) Value	(4) Any trade
Draft cut × canal dependence	-0.070***	-0.089*	-0.080	0.000
	(0.026)	(0.048)	(0.049)	(0.004)
Estimator	Log OLS	Poisson	Poisson	Linear probability
Cells with no trade	Dropped	Dropped	Kept	Kept
Mean of dep. var.	10.993	8,106,191	3,853,254	0.478
Observations	8,339,180	8,339,180	17,544,527	17,544,527
Adjusted R^2	0.841	0.978	0.979	0.664

Notes. A month in which a country reports no imports of a chapter from a partner carries no row in the source file, and taking logs drops it. The panel here is filled out instead, with every importer, exporter and chapter triple that trades at least once, crossed with every month of the window, with the absent cells set to zero. Fifty-two percent of the cells are zero. Column (4) asks whether the cell trades at all. Coverage, unit of analysis and clustering are as in Table S121. The fit statistic in the Poisson columns is the pseudo R squared.

Table S128: The product dimension read twelve times finer

	(1) log trade value	(2) log trade value
Dependence, 1,000 nm × Draft shortfall	-0.038*	-0.064**
	(0.023)	(0.025)
Products read at	Four digits	Two digits
Mean of dep. var.	9.786	10.894
Observations	34,512,116	7,780,955
Adjusted R^2	0.809	0.838

Notes. Chapters number 97 and headings 1,234, so a product by month effect at the finer level absorbs twelve times as much. The comparison is made on the same rows, since the four-digit panel is summed to chapters and the identical specification run on both. Coverage is all 39 importers, 34.5 million cells at four digits and 7.8 million at two, January 2019 to December 2024. Unit of analysis is the importer by exporter by heading by month at four digits and by chapter by month at two. Standard errors are clustered on the country pair.

Table S129: The cost of the draft restriction

Months the draft was below its limit	18
Mean shortfall over those months	9.9 percent
Trade on these pairs in the year before	13,204 bn
Trade displaced over the episode	162 bn
95 percent interval	43-277 bn
under pair-specific trends	83 bn
As a share of trade over the same months	0.82 percent

Notes. The loss on a pair is one minus the exponent of the estimated elasticity times its canal dependence times the mean shortfall, applied to that pair's trade over the twelve months to February 2023 and scaled to the episode. The interval propagates the standard error on the elasticity. The share row divides the displaced trade by trade over the same eighteen months, not by one year of it. Coverage is the reporters the panel holds. Unit of analysis is the country pair, on which the elasticity's errors are clustered.

Table S130: The Atlantic share of an origin's American imports

	(1) Atlantic share	(2) Weighted	(3) Before the tariffs
East of Suez × Closure	-0.036***	-0.036***	-0.026**
	(0.013)	(0.010)	(0.012)
Window	2019-2026	2019-2026	Through March 2025
Mean of dep. var.	0.668	0.668	0.672
Observations	2,576	2,576	2,171
Adjusted R^2	0.921	0.972	0.925

Notes. The outcome is the Atlantic and Gulf share of the containerized vessel imports the United States takes from each origin in each month, from the Census port-level series. The comparison origins are Northeast Asia, and the closure month is December 2023. Column (2) weights each origin by its containerized flow. Column (3) ends the window in March 2025. Coverage is the origins shipping at least a million dollars a month by container on average, 2019-2026. Unit of analysis is the origin by month. Standard errors are clustered by origin.

S5.11 Robustness of the Ratio

Table S131: The efficient exposure ratio across the decay scale

Decay scale, km	50	75	100	150	200	225	300	500	1000	2000	5000
Efficient price of the Red Sea arc, Cape = 1	2.99	2.79	2.70	2.59	2.49	2.44	2.28	1.87	1.33	1.03	0.82

Notes. The plume weight falls with distance from the lane at a scale the pollution regressions sweep rather than assert, and the ratio is recomputed at each scale on the same sample. Coverage and unit of analysis are as in Table 10. The quantities are enumerated rather than sampled, so no standard errors are carried. The limit of the schedule as the decay length grows is the ratio of the two route lengths, which is what a charge proportional to fuel sets.

Table S132: The decay length, fitted on each corridor's own water

	Cape corridor	Red Sea corridor
Decay length on the corridor's own water, km	225	31
standard error	47	7
Change fitted	rise	fall
size at the corridor, μmol m^{-2}	0.405	2.862
Ocean cells	41,650	8,673
Efficient ratio, both corridors on the common kernel	2.59	
each on its own fitted length and amplitude	2.70	
Recovered on synthetic 225 km data, median		226
5th-95th percentile across draws		224-229
Amplitude factor, per-transit concentration	10.8	
Reach factor, kernel weight at own lengths	0.096	
product, the own-kernel over the common ratio	1.04	

Notes. Each column fits an exponential to the change in the nitrogen dioxide column over ocean cells within 600 kilometers of one corridor and more than 600 from the other. The recovered-length rows refit a 225 kilometer decay laid on the Red Sea's own cells. The closing rows reprice the two arcs at the common 150 kilometer length and at each corridor's own. Coverage is the ocean cells of the difference map. Unit of analysis is the 25 kilometer cell. The standard error of each fitted length is from the nonlinear fit and sits on its own row beneath it. The ratio and factor rows span both columns, since each is a property of the pair rather than of one corridor.

Table S133: The arc price with the shared water removed

Water counted	Red Sea arc	Cape arc	Ratio
On the kernel the paper is built on			
every stretch of each routing	100.0	38.7	2.59
only the water each routing alone crosses	65.4	27.2	2.41
Pricing each lane by its closest point			
every stretch of each routing	30.5	12.1	2.53
only the water each routing alone crosses	20.6	8.9	2.31

Notes. A point on one routing counts as the other routing's water as well when the other runs within 25 km of it. Exposure delivered is indexed on the Red Sea arc counted over every stretch. The first row of the upper panel is the construction of Table 10, and the lower panel prices each lane by its closest approach, so a city two lanes pass counts once for each, while the arc table sums the solved corridor as one bundle. Coverage is the 507 coastal cities the incidence table uses. Unit of analysis is the city. The quantities are enumerated rather than sampled, so no standard errors are carried.

Table S134: The efficient exposure ratio at every stand-off setting

Penalty	Scale, km	Efficient ratio	Used
0.0	30	2.11	
0.5	30	2.46	
1.0	30	2.59	✓
2.0	30	2.96	
3.0	30	3.03	
0.5	60	3.06	
1.0	60	3.41	
2.0	60	3.71	
3.0	60	4.09	

Notes. The corridor weights are rebuilt from a fresh solve at each setting of the stand-off penalty and its scale, and the efficient exposure ratio is recomputed on the same populations. Coverage and unit of analysis are as in Table 10. The quantities are enumerated rather than sampled, so no standard errors are carried. The checked row is the setting the calibration selects.

Table S135: The unfavorable margins, one at a time and stacked

Unfavorable margin	Ratio
The baseline, per transit of a given hull	2.59
Wind weighting alone	2.16
Shared water netted alone	2.41
Saturating damage alone	1.76
The three margins at once, per hull	1.17
the same stack per tonne of fuel	1.70
per transit with the fleet and speed margins too	0.75

Notes. Each row reprices the arc ratio under a margin that runs against it, the kernel weights rescaled by the climatological share of hours the wind carries the plume ashore, the water both routings cross removed, and the saturating damage bound. The stacked rows apply all three at once, read per transit and per tonne of fuel, and the last row multiplies in the per-transit observed-fleet and convoy speed factors. Populations are the register's. The quantities are enumerated rather than sampled, so no standard errors are carried.

S6 The City Register

Every coastal city the incidence account enumerates, with its country, its center latitude and longitude, the distance from that center to the shoreline, its exclusive population, its standardized exposure, the side of the reallocation the account assigns it, whether a port sits inside the call radius, its 2022 income per head, and whether it enters the estimation sample, with a dash where the country reports no income figure. The account reproduces from this register alone, and the replication script ships with the data. A city reading No in the last column sits on one of the ten war coasts, which the account carries and the regressions exclude. Populations are the urban-center agglomerations of the settlement layer, so a municipal figure can differ from the printed one where the box merges neighbors under the 78 kilometer rule. The register ships as a CSV beside this printed table, with the corridor weight matrix in both states at the city-by-arc level. The register, the weight matrix and every solved path ship openly, while the Comtrade extraction ships as build scripts against the reader's own subscription.

City	Country	Lat	Lon	Shore, km	People, m	Exposure, sd	Side	Port	$/head	In
Buco-Zau	AGO	-4.76	12.54	52.3	0.050	0.0000	Neither	No	3,599	Yes
Malongo	AGO	-5.39	12.20	0.9	0.839	0.0001	Neither	Yes	3,599	Yes
Soyo Angola LNG Terminal	AGO	-6.12	12.34	1.1	0.251	0.0001	Neither	Yes	3,599	Yes
N'zeto	AGO	-7.23	12.87	1.7	0.036	0.0001	Neither	No	3,599	Yes
Caxito	AGO	-8.58	13.66	28.4	0.065	0.0002	Neither	No	3,599	Yes
Luanda	AGO	-8.78	13.29	1.1	13.359	0.0002	Neither	Yes	3,599	Yes
Gabela	AGO	-10.85	14.38	59.5	0.307	0.0004	Neither	No	3,599	Yes
Seles	AGO	-11.41	14.30	54.0	0.077	0.0005	Neither	No	3,599	Yes
Sumbe	AGO	-11.22	13.85	2.3	0.179	0.0006	Neither	No	3,599	Yes
Lobito	AGO	-12.34	13.57	0.5	0.487	0.0012	Neither	Yes	3,599	Yes
Namibe	AGO	-15.16	12.14	1.6	0.253	0.0101	Neither	Yes	3,599	Yes
Tombua	AGO	-15.80	11.84	1.1	0.022	0.0159	Neither	No	3,599	Yes
Jebel Dhanna	ARE	24.17	52.59	1.5	0.108	-0.0025	Neither	Yes	50,760	Yes
Al Aamerah	ARE	24.22	55.55	99.2	0.315	-0.0022	Neither	No	50,760	Yes
ICAD I	ARE	24.36	54.50	3.9	0.509	-0.0020	Neither	Yes	50,760	Yes
Zirku Island	ARE	24.89	53.07	0.8	0.001	-0.0016	Neither	Yes	50,760	Yes
Khalifa Port	ARE	24.81	54.69	1.2	0.132	-0.0015	Neither	Yes	50,760	Yes
Jebel Ali	ARE	25.00	55.08	2.0	1.154	-0.0013	Neither	Yes	50,760	Yes
Fujairah	ARE	25.20	56.36	1.2	0.379	-0.0011	Neither	Yes	50,760	Yes
Ajman	ARE	25.42	55.45	1.2	1.611	-0.0010	Neither	Yes	50,760	Yes
Mina Saqr	ARE	25.97	56.06	2.5	0.118	-0.0007	Neither	Yes	50,760	Yes
Ikpinle	BEN	6.91	2.62	60.3	0.413	0.0001	Neither	No	1,266	Yes
Cotonou	BEN	6.35	2.42	1.9	2.162	0.0001	Neither	Yes	1,266	Yes
Bohicon	BEN	7.18	2.04	98.1	0.719	0.0001	Neither	No	1,266	Yes
Come	BEN	6.41	1.88	15.3	0.527	0.0002	Neither	No	1,266	Yes
Sitrah	BHR	26.19	50.67	1.6	1.985	-0.0017	Neither	Yes	30,471	Yes
Aboisso	CIV	5.47	-3.20	12.4	0.246	0.0072	Neither	No	2,333	Yes
Adzope	CIV	6.11	-3.86	81.6	0.191	0.0099	Neither	No	2,333	Yes
Abidjan	CIV	5.28	-4.01	0.8	4.645	0.0129	Neither	Yes	2,333	Yes
Agboville	CIV	5.93	-4.21	69.5	0.173	0.0131	Neither	No	2,333	Yes
Dabou	CIV	5.33	-4.38	3.4	0.033	0.0164	Neither	No	2,333	Yes
Divo	CIV	5.83	-5.36	67.0	0.334	0.0291	Neither	No	2,333	Yes
Meagui	CIV	5.41	-6.56	66.9	0.339	0.0711	Gained	No	2,333	Yes
San Pedro	CIV	4.74	-6.62	1.4	0.303	0.0839	Gained	Yes	2,333	Yes
Nkongsamba	CMR	4.96	9.94	93.9	0.182	0.0000	Neither	No	1,605	Yes
Limboh	CMR	4.01	9.13	1.5	0.329	0.0000	Neither	Yes	1,605	Yes
Douala	CMR	4.05	9.68	2.2	3.222	0.0000	Neither	Yes	1,605	Yes
Edea	CMR	3.79	10.13	39.1	0.010	0.0000	Neither	No	1,605	Yes

City	Country	Lat	Lon	Shore, km	People, m	Exposure, sd	Side	Port	$/head	In
Kumba	CMR	4.63	9.45	57.3	0.417	0.0000	Neither	No	1,605	Yes
Tshela	COD	-5.00	12.95	90.3	0.153	0.0000	Neither	No	688	Yes
Matadi	COD	-5.83	13.44	28.9	0.210	0.0000	Neither	Yes	688	Yes
Lukula	COD	-5.40	12.95	46.6	0.145	0.0000	Neither	No	688	Yes
Boma	COD	-5.83	13.06	3.5	0.227	0.0001	Neither	No	688	Yes
Moroni	COM	-11.71	43.25	2.7	0.334	0.0268	Neither	Yes	1,534	Yes
Fomboni	COM	-12.29	43.75	0.9	0.312	0.0420	Neither	No	1,534	Yes
Mutsamudu	COM	-12.17	44.39	1.5	0.058	0.0527	Gained	Yes	1,534	Yes
Porto Grande	CPV	16.89	-25.00	1.3	0.098	0.0603	Gained	Yes	4,323	Yes
Port Vale Knights	CPV	14.92	-24.50	2.5	0.037	0.0783	Gained	Yes	4,323	Yes
Porto Da Praia	CPV	14.91	-23.50	1.0	0.208	0.1505	Gained	Yes	4,323	Yes
Palmeira	CPV	16.75	-22.98	0.8	0.020	0.2183	Gained	Yes	4,323	Yes
Sal Rei	CPV	16.18	-22.92	1.3	0.017	0.2241	Gained	Yes	4,323	Yes
Pafos	CYP	34.77	32.42	1.7	0.098	-0.9090	Lost	No	33,231	Yes
Vassiliko	CYP	34.72	33.31	1.6	0.046	-0.6689	Lost	Yes	33,231	Yes
Dhekelia	CYP	34.98	33.74	1.6	0.204	-0.5018	Lost	Yes	33,231	Yes
Kyrenia	CYP	35.34	33.33	0.6	0.380	-0.4935	Lost	Yes	33,231	Yes
Djibouti	DJI	11.60	43.10	2.3	0.575	-1.8903	Lost	Yes	3,133	No
Ali Sabieh	DJI	11.15	42.71	36.5	0.103	-1.3125	Lost	No	3,133	No
Annaba	DZA	36.90	7.77	0.9	0.582	-1.4400	Lost	Yes	4,960	Yes
Skikda (Port Methanier)	DZA	36.87	6.95	1.9	0.446	-1.3797	Lost	Yes	4,960	Yes
Djen Djen	DZA	36.82	5.89	1.6	0.337	-1.3021	Lost	Yes	4,960	Yes
El Milia	DZA	36.75	6.27	17.9	0.260	-1.2731	Lost	Yes	4,960	Yes
Bejaia	DZA	36.75	5.08	1.5	0.444	-1.2328	Lost	Yes	4,960	Yes
Algiers	DZA	36.77	3.07	0.9	3.560	-1.2283	Lost	Yes	4,960	Yes
Tizi Ouzou	DZA	36.71	4.04	20.0	0.685	-1.1960	Lost	No	4,960	Yes
Guelma	DZA	36.46	7.44	56.4	0.273	-1.1484	Lost	No	4,960	Yes
Lakhdaria	DZA	36.57	3.60	25.1	0.618	-1.1134	Lost	No	4,960	Yes
El Affroun	DZA	36.46	2.63	15.1	0.880	-1.0678	Lost	No	4,960	Yes
Constantine	DZA	36.36	6.62	64.7	0.859	-1.0551	Lost	No	4,960	Yes
Ain Defla	DZA	36.25	1.95	34.4	0.367	-0.9822	Lost	No	4,960	Yes
M'Daourouch	DZA	36.08	7.82	86.3	0.340	-0.9672	Lost	No	4,960	Yes
Chlef	DZA	36.15	1.33	40.7	0.620	-0.9635	Lost	No	4,960	Yes
Mostaganem	DZA	35.93	0.09	1.8	0.430	-0.9577	Lost	Yes	4,960	Yes
Arzew	DZA	35.82	-0.27	0.8	0.077	-0.9437	Lost	Yes	4,960	Yes
Chelghoum Laid	DZA	36.17	6.18	77.3	0.301	-0.9378	Lost	No	4,960	Yes
Oran	DZA	35.71	-0.64	1.1	1.509	-0.9309	Lost	Yes	4,960	Yes
El Eulma	DZA	36.15	5.68	60.2	0.445	-0.9181	Lost	No	4,960	Yes
Oued Rhiou	DZA	35.96	0.92	44.6	0.328	-0.9043	Lost	No	4,960	Yes
Bordj Bou Arreridj	DZA	36.08	4.76	76.0	0.411	-0.8662	Lost	No	4,960	Yes
Relizane	DZA	35.74	0.56	48.1	0.226	-0.8402	Lost	No	4,960	Yes
Ain Oulmane	DZA	35.92	5.29	80.7	0.272	-0.7994	Lost	No	4,960	Yes
Ksar El Boukhari	DZA	35.88	2.75	80.8	0.121	-0.7894	Lost	No	4,960	Yes
Ghazaouet	DZA	35.10	-1.86	2.1	0.071	-0.7505	Lost	Yes	4,960	Yes
Mascara	DZA	35.40	0.14	48.0	0.429	-0.7368	Lost	No	4,960	Yes
Sidi Bel Abbes	DZA	35.20	-0.62	58.0	0.418	-0.7166	Lost	No	4,960	Yes
Remchi	DZA	35.06	-1.43	24.3	0.489	-0.7128	Lost	Yes	4,960	Yes
Al Ismailiya	EGY	30.59	32.27	53.7	0.604	-3.1293	Lost	No	4,233	Yes
Al Adabiyah	EGY	29.90	32.46	0.9	0.522	-3.1112	Lost	Yes	4,233	Yes
Port Said	EGY	31.25	32.31	2.5	0.572	-3.0942	Lost	Yes	4,233	Yes
Damietta	EGY	31.46	31.76	2.4	1.083	-3.0732	Lost	Yes	4,233	Yes
Zeit Bay	EGY	27.82	33.58	1.0	0.005	-2.9871	Lost	Yes	4,233	Yes
Ras Gharib	EGY	28.34	33.11	0.4	0.016	-2.9805	Lost	Yes	4,233	Yes
Faqous	EGY	30.72	31.78	48.3	1.695	-2.8880	Lost	No	4,233	Yes
Hurghada	EGY	27.22	33.83	1.9	0.172	-2.8191	Lost	No	4,233	Yes

City	Country	Lat	Lon	Shore, km	People, m	Exposure, sd	Side	Port	$/head	In
Tenth Of Ramadan	EGY	30.31	31.74	82.4	0.584	-2.7368	Lost	No	4,233	Yes
Mansoura	EGY	31.17	31.22	36.8	2.458	-2.6802	Lost	No	4,233	Yes
Safaga	EGY	26.73	33.94	1.4	0.026	-2.5925	Lost	Yes	4,233	Yes
Mit Ghamr	EGY	30.76	31.33	70.7	2.796	-2.5862	Lost	No	4,233	Yes
Sidi Salim	EGY	31.27	30.80	14.0	1.145	-2.4491	Lost	No	4,233	Yes
Tanta	EGY	30.66	30.93	83.0	3.163	-2.2471	Lost	No	4,233	Yes
Egyptian LNG Terminal	EGY	31.36	30.31	1.6	0.547	-2.1981	Lost	Yes	4,233	Yes
Delingat	EGY	30.83	30.55	51.6	1.379	-2.0692	Lost	No	4,233	Yes
Sadat	EGY	30.38	30.55	98.3	0.102	-1.8671	Lost	No	4,233	Yes
Arish	EGY	31.13	33.80	2.4	0.149	-1.8426	Lost	No	4,233	Yes
Alexandria	EGY	31.14	29.81	0.1	4.166	-1.8048	Lost	Yes	4,233	Yes
Marsa Matruh	EGY	31.35	27.25	3.0	0.169	-1.1518	Lost	No	4,233	Yes
Assab	ERI	13.02	42.74	0.6	0.021	-2.8992	Lost	No	–	No
Massawa	ERI	15.60	39.46	0.7	0.087	-1.5203	Lost	Yes	–	No
Ghinda	ERI	15.44	39.09	40.8	0.399	-1.2162	Lost	No	–	No
Dekemhare	ERI	15.07	39.04	67.7	0.311	-1.0637	Lost	No	–	No
Keren	ERI	15.78	38.46	91.0	0.170	-0.9941	Lost	No	–	No
Dakhla	ESH	23.71	-15.94	0.5	0.113	1.9573	Gained	No	–	Yes
Boujdour	ESH	26.13	-14.49	0.6	0.036	2.2745	Gained	No	–	Yes
Carboneras	ESP	36.97	-1.91	0.5	0.053	-1.0169	Lost	Yes	30,319	Yes
Ibiza	ESP	38.91	1.43	1.2	0.157	-1.0076	Lost	No	30,319	Yes
Cartagena	ESP	37.58	-0.95	1.9	0.234	-0.9885	Lost	Yes	30,319	Yes
Almeria	ESP	36.83	-2.47	0.8	0.374	-0.9711	Lost	Yes	30,319	Yes
Mahon	ESP	39.89	4.27	3.2	0.069	-0.9187	Lost	Yes	30,319	Yes
Palma de Mallorca	ESP	39.55	2.63	0.6	0.641	-0.9157	Lost	Yes	30,319	Yes
Torrevieja	ESP	37.97	-0.70	0.7	0.213	-0.8911	Lost	Yes	30,319	Yes
Motril	ESP	36.72	-3.52	0.5	0.112	-0.8847	Lost	Yes	30,319	Yes
Melilla	ESP	35.29	-2.93	0.8	0.452	-0.8284	Lost	Yes	30,319	Yes
Murcia	ESP	37.98	-1.14	35.9	0.588	-0.8047	Lost	Yes	30,319	Yes
Alicante	ESP	38.33	-0.49	1.6	0.613	-0.7949	Lost	Yes	30,319	Yes
Malaga	ESP	36.71	-4.42	1.1	0.928	-0.7891	Lost	Yes	30,319	Yes
Marbella	ESP	36.50	-4.96	2.1	0.202	-0.7657	Lost	No	30,319	Yes
Granada	ESP	37.16	-3.61	46.9	0.538	-0.7436	Lost	No	30,319	Yes
Alcoi / Alcoy	ESP	38.70	-0.47	28.4	0.156	-0.6668	Lost	No	30,319	Yes
Huelva	ESP	37.19	-6.91	0.6	0.264	-0.6572	Lost	Yes	30,319	Yes
Cadiz	ESP	36.52	-6.26	0.9	0.682	-0.6175	Lost	Yes	30,319	Yes
Seville	ESP	37.32	-5.99	46.5	1.357	-0.5812	Lost	Yes	30,319	Yes
Sagunto	ESP	39.64	-0.22	0.8	1.553	-0.4467	Lost	Yes	30,319	Yes
Santa Cruz De La Palma	ESP	28.68	-17.77	0.2	0.087	0.8921	Gained	Yes	30,319	Yes
Los Cristianos	ESP	28.07	-16.73	0.5	0.149	1.5589	Gained	No	30,319	Yes
Santa Cruz De Tenerife	ESP	28.48	-16.24	1.3	0.501	1.7298	Gained	Yes	30,319	Yes
Arrecife	ESP	28.97	-13.53	1.7	0.120	2.3025	Gained	Yes	30,319	Yes
Las Palmas (de Gran Canaria)	ESP	28.15	-15.41	0.7	0.676	2.3902	Gained	Yes	30,319	Yes
Puerto Del Rosario	ESP	28.50	-13.86	0.4	0.076	2.4697	Gained	Yes	30,319	Yes
Adigrat	ETH	14.28	39.47	93.6	0.502	-1.0264	Lost	No	982	Yes
Port Owendo	GAB	0.30	9.50	0.7	0.889	0.0000	Neither	Yes	8,409	Yes
Gamba	GAB	-2.82	10.00	4.5	0.003	0.0001	Neither	Yes	8,409	Yes
Akatsi	GHA	6.13	0.80	30.8	0.245	0.0004	Neither	No	2,230	Yes
Ho	GHA	6.60	0.47	92.2	0.475	0.0005	Neither	No	2,230	Yes
Odumase	GHA	6.13	0.01	48.1	0.252	0.0007	Neither	No	2,230	Yes
Tema	GHA	5.63	0.00	1.4	3.255	0.0008	Neither	Yes	2,230	Yes
Suhum	GHA	6.04	-0.45	60.5	0.561	0.0010	Neither	No	2,230	Yes
Agona Swedru	GHA	5.53	-0.69	22.5	0.841	0.0013	Neither	No	2,230	Yes
Assin Fosu	GHA	5.70	-1.28	61.6	0.337	0.0018	Neither	No	2,230	Yes
Elmina	GHA	5.09	-1.35	1.8	0.412	0.0022	Neither	No	2,230	Yes

City	Country	Lat	Lon	Shore, km	People, m	Exposure, sd	Side	Port	$/head	In
Sekondi	GHA	4.94	-1.70	0.5	0.975	0.0028	Neither	Yes	2,230	Yes
Tarkwa	GHA	5.30	-1.99	51.1	0.091	0.0032	Neither	No	2,230	Yes
Kindia	GIN	10.04	-12.86	81.6	0.276	0.5660	Gained	No	1,417	Yes
Sangaredi	GIN	11.11	-13.81	79.0	0.140	0.6420	Gained	No	1,417	Yes
Tanene	GIN	10.08	-13.70	14.1	0.101	0.8291	Gained	No	1,417	Yes
Conakry	GIN	9.43	-13.39	1.2	1.304	0.8979	Gained	Yes	1,417	Yes
Dapilon	GIN	10.84	-14.50	1.0	0.255	0.9544	Gained	Yes	1,417	Yes
Banjul	GMB	13.44	-16.57	1.2	1.313	1.8760	Gained	Yes	836	Yes
Gabu	GNB	12.28	-14.23	90.1	0.106	0.6173	Gained	No	882	Yes
Bafata	GNB	12.18	-14.65	44.1	0.125	0.7784	Gained	No	882	Yes
Bissau	GNB	11.86	-15.58	1.4	0.596	1.2814	Gained	Yes	882	Yes
Bata	GNQ	1.83	9.74	1.2	0.102	0.0000	Neither	Yes	7,589	Yes
Luba	GNQ	3.46	8.55	1.9	0.090	0.0000	Neither	Yes	7,589	Yes
Kali Limenes	GRC	34.93	24.80	1.6	0.020	-1.9224	Lost	Yes	20,885	Yes
Atherinolakos	GRC	35.00	26.14	0.5	0.019	-1.9046	Lost	Yes	20,885	Yes
Heraklion	GRC	35.34	25.15	0.6	0.205	-1.6334	Lost	Yes	20,885	Yes
Rethymnon	GRC	35.37	24.48	0.5	0.052	-1.6108	Lost	Yes	20,885	Yes
Souda (Chania)	GRC	35.51	24.07	0.6	0.110	-1.5187	Lost	Yes	20,885	Yes
Lavrion	GRC	37.70	24.06	0.2	0.080	-0.3720	Lost	Yes	20,885	Yes
Pachi	GRC	37.97	23.38	1.4	0.260	-0.2970	Lost	Yes	20,885	Yes
Aiyion	GRC	38.34	21.85	0.6	0.235	-0.2448	Lost	Yes	20,885	Yes
Ormos Aliveriou	GRC	38.38	24.06	0.5	0.046	-0.2383	Lost	Yes	20,885	Yes
San Nicolas	GRC	38.35	22.68	0.9	0.041	-0.2327	Lost	Yes	20,885	Yes
Chalkis	GRC	38.46	23.59	0.5	0.101	-0.2196	Lost	Yes	20,885	Yes
Tsingeli	GRC	39.16	22.85	0.6	0.148	-0.1394	Lost	Yes	20,885	Yes
Ioannina	GRC	39.66	20.85	52.5	0.125	-0.1133	Lost	No	20,885	Yes
Larissa	GRC	39.63	22.42	38.6	0.185	-0.1049	Lost	No	20,885	Yes
Abadan	IRN	30.33	48.28	34.5	0.249	-0.0015	Neither	Yes	4,721	Yes
Shadegan	IRN	30.65	48.66	32.7	0.097	-0.0011	Neither	No	4,721	Yes
Bandar Khomeini	IRN	30.44	49.07	2.9	0.236	-0.0010	Neither	Yes	4,721	Yes
Ahwaz	IRN	31.33	48.71	94.5	0.975	-0.0009	Neither	No	4,721	Yes
Jask	IRN	25.66	57.77	2.4	0.004	-0.0008	Neither	Yes	4,721	Yes
Kharg Island	IRN	29.24	50.34	0.5	0.025	-0.0007	Neither	Yes	4,721	Yes
Bahregan	IRN	29.83	50.25	1.6	0.037	-0.0006	Neither	Yes	4,721	Yes
Bushehr	IRN	28.99	50.83	0.8	0.208	-0.0006	Neither	Yes	4,721	Yes
Chabahar	IRN	25.30	60.61	1.6	0.142	-0.0006	Neither	Yes	4,721	Yes
Lavan	IRN	26.79	53.35	2.0	0.011	-0.0005	Neither	Yes	4,721	Yes
Behbahan	IRN	30.59	50.24	46.5	0.108	-0.0005	Neither	No	4,721	Yes
Bandar-E Pars Terminal	IRN	27.56	52.51	2.6	0.293	-0.0005	Neither	Yes	4,721	Yes
Borazjan	IRN	29.27	51.22	36.7	0.120	-0.0004	Neither	No	4,721	Yes
Seerik	IRN	26.53	57.10	1.5	0.023	-0.0004	Neither	No	4,721	Yes
Dogonbadan	IRN	30.36	50.80	69.4	0.058	-0.0004	Neither	No	4,721	Yes
Charam	IRN	30.75	50.74	89.4	0.065	-0.0004	Neither	No	4,721	Yes
Kazerun	IRN	29.62	51.65	93.9	0.132	-0.0003	Neither	No	4,721	Yes
Bandar Shahid Rajaee	IRN	27.14	56.06	2.7	0.396	-0.0003	Neither	Yes	4,721	Yes
Minab	IRN	27.13	57.08	18.2	0.174	-0.0003	Neither	No	4,721	Yes
Umm Qasr	IRQ	30.11	47.90	3.3	0.090	-0.0020	Neither	Yes	6,521	Yes
Al Basrah	IRQ	30.56	47.80	51.8	1.806	-0.0019	Neither	No	6,521	Yes
Basrah Oil Terminal	IRQ	29.68	48.82	36.3	0.000	-0.0014	Neither	Yes	6,521	Yes
Ashdod	ISR	31.84	34.67	2.0	1.312	-1.0558	Lost	Yes	54,947	No
Hadera	ISR	32.47	34.88	0.5	0.629	-0.7981	Lost	Yes	54,947	No
Nahariyya	ISR	33.00	35.10	1.7	0.813	-0.6088	Lost	Yes	54,947	No
Tiberias	ISR	32.79	35.53	43.7	0.364	-0.5308	Lost	No	54,947	No
Trapani	ITA	38.01	12.51	0.3	0.177	-1.4800	Lost	Yes	35,654	Yes
Sarroch Oil Terminal	ITA	39.09	9.01	1.0	0.375	-1.4188	Lost	Yes	35,654	Yes

City	Country	Lat	Lon	Shore, km	People, m	Exposure, sd	Side	Port	$/head	In
Gela	ITA	37.05	14.28	0.9	0.198	-1.2828	Lost	Yes	35,654	Yes
Palermo	ITA	38.13	13.37	0.8	0.877	-1.1002	Lost	Yes	35,654	Yes
Santa Panagia	ITA	37.11	15.26	0.2	0.220	-0.9886	Lost	Yes	35,654	Yes
Oristano	ITA	39.86	8.56	0.6	0.098	-0.9543	Lost	Yes	35,654	Yes
Catania	ITA	37.49	15.09	0.4	0.743	-0.8747	Lost	Yes	35,654	Yes
Milazzo	ITA	38.21	15.26	0.6	0.099	-0.6139	Lost	Yes	35,654	Yes
Reggio Calabria	ITA	38.11	15.66	2.3	0.352	-0.5815	Lost	Yes	35,654	Yes
Gioia Tauro	ITA	38.46	15.91	1.4	0.184	-0.4626	Lost	Yes	35,654	Yes
Cosenza	ITA	39.32	16.24	17.2	0.306	-0.2808	Lost	No	35,654	Yes
Aqaba	JOR	29.45	34.96	1.0	0.146	-1.6088	Lost	Yes	4,754	Yes
Jerash	JOR	32.28	35.90	97.5	1.404	-0.5178	Lost	No	4,754	Yes
Masalani	KEN	-1.69	40.17	81.9	0.118	0.0000	Neither	No	2,110	Yes
Lamu	KEN	-2.20	40.91	0.6	0.085	0.0001	Neither	Yes	2,110	Yes
Malindi	KEN	-3.23	40.11	2.4	0.408	0.0001	Neither	No	2,110	Yes
Mombasa	KEN	-4.05	39.63	2.5	1.725	0.0002	Neither	Yes	2,110	Yes
Al Jahra	KWT	29.34	47.68	4.1	0.542	-0.0029	Neither	Yes	39,982	Yes
Shuaiba	KWT	29.04	48.17	2.6	0.474	-0.0024	Neither	Yes	39,982	Yes
Sidon	LBN	33.60	35.40	1.3	2.559	-0.4295	Lost	Yes	3,654	No
Tripoli	LBN	34.45	35.83	0.8	0.970	-0.2550	Lost	Yes	3,654	No
Pleebo	LBR	4.59	-7.67	23.6	0.165	0.1750	Gained	No	745	Yes
Kakata	LBR	6.53	-10.35	38.2	0.105	0.5745	Gained	No	745	Yes
Buchanan	LBR	5.87	-10.04	0.5	0.103	0.5964	Gained	Yes	745	Yes
Monrovia	LBR	6.34	-10.79	1.0	1.376	0.7848	Gained	Yes	745	Yes
Derna	LBY	32.76	22.66	1.6	0.057	-1.5321	Lost	Yes	5,987	No
Bayda	LBY	32.76	21.76	17.4	0.192	-1.3099	Lost	No	5,987	No
Marsa Tobruk	LBY	32.06	23.98	0.5	0.087	-1.2530	Lost	Yes	5,987	No
Marj	LBY	32.49	20.83	17.1	0.142	-0.9871	Lost	No	5,987	No
Benghazi	LBY	32.11	20.05	2.2	0.446	-0.7309	Lost	Yes	5,987	No
Khoms	LBY	32.68	14.25	2.3	0.221	-0.6812	Lost	Yes	5,987	No
Tripoli	LBY	32.90	13.20	1.8	1.225	-0.6487	Lost	Yes	5,987	No
Misrata	LBY	32.34	15.22	1.7	0.251	-0.6166	Lost	Yes	5,987	No
Az Zawiyah	LBY	32.79	12.70	2.5	0.318	-0.5377	Lost	Yes	5,987	No
Mellitah	LBY	32.85	12.24	2.5	0.153	-0.4981	Lost	Yes	5,987	No
Ez Zueitina	LBY	30.90	20.08	0.4	0.056	-0.3624	Lost	Yes	5,987	No
Sirte	LBY	31.20	16.58	2.5	0.056	-0.3291	Lost	No	5,987	No
Ras Lanuf (Al Burayqah)	LBY	30.41	19.60	0.6	0.026	-0.2543	Lost	Yes	5,987	No
Al Hoceima	MAR	35.25	-3.92	0.5	0.176	-0.7280	Lost	Yes	3,463	Yes
Berkane	MAR	34.92	-2.33	21.1	0.241	-0.6807	Lost	No	3,463	Yes
Tangier-Mediterranean	MAR	35.88	-5.51	2.1	0.454	-0.6564	Lost	Yes	3,463	Yes
Oujda	MAR	34.68	-1.91	43.9	0.568	-0.5837	Lost	No	3,463	Yes
Taourirt	MAR	34.40	-2.90	80.1	0.125	-0.4760	Lost	No	3,463	Yes
Chefchaouen	MAR	35.17	-5.27	33.0	0.192	-0.4559	Lost	No	3,463	Yes
Ouazzane	MAR	34.80	-5.58	65.1	0.151	-0.2488	Lost	No	3,463	Yes
Larache	MAR	35.18	-6.15	3.1	0.278	-0.2279	Lost	No	3,463	Yes
Souk El Arbaa	MAR	34.68	-5.99	34.2	0.221	-0.1036	Lost	No	3,463	Yes
Sidi Slimane	MAR	34.26	-5.93	59.4	0.344	-0.0129	Neither	No	3,463	Yes
Tiflet	MAR	33.89	-6.33	47.5	0.142	0.1568	Gained	No	3,463	Yes
Kenitra	MAR	34.29	-6.62	3.2	0.986	0.1824	Gained	Yes	3,463	Yes
Benslimane	MAR	33.62	-7.12	21.8	0.284	0.4301	Gained	Yes	3,463	Yes
Settat	MAR	33.00	-7.62	61.6	0.298	0.5897	Gained	No	3,463	Yes
Casablanca	MAR	33.61	-7.58	1.6	3.854	0.5969	Gained	Yes	3,463	Yes
Ouled-Teima	MAR	30.39	-9.22	37.2	0.190	0.7007	Gained	No	3,463	Yes
Guelmim	MAR	28.99	-10.06	38.5	0.126	0.7249	Gained	No	3,463	Yes
Tiznit	MAR	29.69	-9.73	17.7	0.129	0.7388	Gained	No	3,463	Yes
Youssoufia	MAR	32.25	-8.53	69.4	0.188	0.8189	Gained	No	3,463	Yes

City	Country	Lat	Lon	Shore, km	People, m	Exposure, sd	Side	Port	$/head	In
Sidi Bennour	MAR	32.65	-8.43	48.4	0.277	0.8477	Gained	No	3,463	Yes
Agadir	MAR	30.44	-9.65	0.7	0.974	0.8550	Gained	Yes	3,463	Yes
Jorf Lasfar	MAR	33.11	-8.61	2.4	0.342	0.9850	Gained	Yes	3,463	Yes
Safi	MAR	32.31	-9.25	2.3	0.413	1.1167	Gained	Yes	3,463	Yes
Essaouira	MAR	31.51	-9.76	2.0	0.120	1.1489	Gained	No	3,463	Yes
Tan Tan	MAR	28.47	-11.35	1.1	0.040	1.1735	Gained	Yes	3,463	Yes
Manakara	MDG	-22.13	48.01	3.1	0.239	0.0030	Neither	No	504	Yes
Farafangana	MDG	-22.82	47.82	2.0	0.255	0.0031	Neither	No	504	Yes
Toamasina	MDG	-18.16	49.43	0.7	0.392	0.0052	Neither	Yes	504	Yes
Ehoala	MDG	-25.07	46.96	0.1	0.091	0.0062	Neither	Yes	504	Yes
Ambovombe	MDG	-25.18	46.09	8.0	0.301	0.0071	Neither	No	504	Yes
Tulear (Toliara)	MDG	-23.37	43.67	1.1	0.267	0.0149	Neither	Yes	504	Yes
Morondava	MDG	-20.29	44.29	2.5	0.125	0.0290	Neither	No	504	Yes
Ankazomborona	MDG	-16.08	46.68	25.6	0.100	0.0417	Neither	No	504	Yes
Antsiranana (Diego Suarez)	MDG	-12.27	49.29	0.9	0.162	0.0547	Gained	Yes	504	Yes
Majunga	MDG	-15.73	46.31	1.0	0.311	0.0553	Gained	Yes	504	Yes
Andoany	MDG	-13.40	48.21	0.9	0.145	0.0587	Gained	Yes	504	Yes
Valletta	MLT	35.88	14.50	2.9	0.368	-1.9502	Lost	Yes	35,659	Yes
Mocimboa da Praia	MOZ	-11.36	40.35	1.8	0.078	0.0059	Neither	No	578	Yes
Nhamatanda	MOZ	-19.27	34.21	49.9	0.263	0.0062	Neither	No	578	Yes
Mopeia	MOZ	-17.98	35.71	97.1	0.123	0.0083	Neither	No	578	Yes
Mocuba	MOZ	-16.84	36.98	98.2	0.204	0.0101	Neither	No	578	Yes
Beira	MOZ	-19.81	34.84	1.8	0.512	0.0104	Neither	Yes	578	Yes
Chiure-Sede	MOZ	-13.38	39.79	80.7	0.165	0.0108	Neither	No	578	Yes
Nicoadala	MOZ	-17.60	36.81	30.8	0.306	0.0124	Neither	No	578	Yes
Pemba	MOZ	-12.97	40.49	1.1	0.330	0.0127	Neither	Yes	578	Yes
Maganja	MOZ	-17.31	37.50	24.4	0.148	0.0157	Neither	No	578	Yes
Pebane	MOZ	-17.27	38.15	1.8	0.110	0.0212	Neither	No	578	Yes
Namialo	MOZ	-14.92	39.98	62.9	0.313	0.0223	Neither	No	578	Yes
Chokwe	MOZ	-24.53	32.99	87.0	0.157	0.0257	Neither	No	578	Yes
Nacala	MOZ	-14.54	40.65	1.3	0.396	0.0264	Neither	Yes	578	Yes
Vilanculos	MOZ	-22.00	35.30	2.5	0.086	0.0270	Neither	No	578	Yes
Moma	MOZ	-16.75	39.22	5.1	0.190	0.0296	Neither	No	578	Yes
Chibuto	MOZ	-24.70	33.55	48.6	0.172	0.0340	Neither	No	578	Yes
Maputo	MOZ	-25.96	32.52	0.8	2.501	0.0413	Neither	Yes	578	Yes
Xai-Xai	MOZ	-25.06	33.69	6.3	0.266	0.0417	Neither	No	578	Yes
Massinga	MOZ	-23.32	35.37	10.6	0.140	0.0418	Neither	No	578	Yes
Maxixe	MOZ	-23.85	35.34	1.8	0.333	0.0491	Neither	No	578	Yes
Nouakchott	MRT	17.99	-16.03	1.5	1.171	1.4033	Gained	Yes	1,961	Yes
Nouadhibou	MRT	20.85	-17.05	2.0	0.133	2.3620	Gained	Yes	1,961	Yes
Port Louis	MUS	-20.15	57.50	1.1	0.896	0.0074	Neither	Yes	10,247	Yes
Port de Longoni	MYT	-12.72	45.17	1.8	0.230	0.0852	Gained	Yes	–	Yes
Walvis Bay	NAM	-22.95	14.49	1.5	0.073	0.1413	Gained	Yes	4,350	Yes
Luderitz Bay	NAM	-26.64	15.15	1.5	0.010	0.5474	Gained	Yes	4,350	Yes
Calabar	NGA	5.01	8.33	11.8	0.701	0.0000	Neither	Yes	2,899	Yes
Nguzu Edda	NGA	5.76	7.83	93.7	1.277	0.0000	Neither	No	2,899	Yes
James Town	NGA	4.66	8.31	1.6	0.329	0.0000	Neither	Yes	2,899	Yes
Ekenobizi	NGA	5.57	7.39	96.7	2.701	0.0000	Neither	No	2,899	Yes
Ikot Ekpene	NGA	5.14	7.72	51.6	1.483	0.0000	Neither	No	2,899	Yes
Eket	NGA	4.66	7.95	13.4	0.974	0.0000	Neither	No	2,899	Yes
Aba	NGA	5.10	7.35	48.4	1.144	0.0000	Neither	No	2,899	Yes
Opobo	NGA	4.50	7.54	2.2	0.706	0.0000	Neither	No	2,899	Yes
Port Harcourt	NGA	4.84	7.00	12.3	1.856	0.0000	Neither	Yes	2,899	Yes
Bonny	NGA	4.42	7.16	1.2	0.315	0.0000	Neither	Yes	2,899	Yes
Agbor	NGA	6.25	6.19	99.7	0.638	0.0000	Neither	No	2,899	Yes

City	Country	Lat	Lon	Shore, km	People, m	Exposure, sd	Side	Port	$/head	In
Akabuka	NGA	5.24	6.65	45.6	0.697	0.0000	Neither	No	2,899	Yes
Abua	NGA	4.85	6.64	8.0	0.729	0.0000	Neither	No	2,899	Yes
Abraka	NGA	5.79	6.11	58.4	0.422	0.0000	Neither	No	2,899	Yes
Benin City	NGA	6.34	5.63	60.3	1.352	0.0000	Neither	No	2,899	Yes
Patani	NGA	5.23	6.19	70.2	0.617	0.0000	Neither	No	2,899	Yes
Sapele	NGA	5.92	5.65	36.4	0.394	0.0000	Neither	Yes	2,899	Yes
Warri	NGA	5.52	5.72	9.4	1.017	0.0000	Neither	Yes	2,899	Yes
Brass	NGA	4.31	6.25	1.4	0.179	0.0000	Neither	No	2,899	Yes
Ekeremor	NGA	5.07	5.77	31.7	0.442	0.0000	Neither	No	2,899	Yes
Ondo	NGA	7.10	4.84	94.3	1.023	0.0000	Neither	No	2,899	Yes
Okitipupa	NGA	6.50	4.78	34.2	0.884	0.0000	Neither	No	2,899	Yes
Ibadan	NGA	7.40	3.91	87.4	3.386	0.0000	Neither	No	2,899	Yes
Ijebu-Igbo	NGA	6.97	4.00	43.4	0.499	0.0000	Neither	No	2,899	Yes
Lekki	NGA	6.40	4.06	2.6	0.191	0.0000	Neither	Yes	2,899	Yes
Sagamu	NGA	6.84	3.64	28.3	0.322	0.0001	Neither	No	2,899	Yes
Lagos	NGA	6.45	3.33	4.0	10.367	0.0001	Neither	Yes	2,899	Yes
Aiyetoro	NGA	7.24	3.03	86.5	0.413	0.0001	Neither	No	2,899	Yes
Ilaro	NGA	6.89	3.01	56.1	0.183	0.0001	Neither	No	2,899	Yes
Ajara	NGA	6.44	2.93	5.2	0.273	0.0001	Neither	No	2,899	Yes
Salalah	OMN	16.94	54.01	0.6	0.147	-0.1934	Lost	Yes	23,224	Yes
Duqm	OMN	19.67	57.72	0.4	0.013	-0.0460	Neither	Yes	23,224	Yes
Qalhat LNG Terminal	OMN	22.65	59.42	1.6	0.050	-0.0052	Neither	Yes	23,224	Yes
Seeb	OMN	23.64	58.17	5.3	0.487	-0.0032	Neither	Yes	23,224	Yes
Port Sultan Qaboos	OMN	23.63	58.57	2.5	0.249	-0.0031	Neither	Yes	23,224	Yes
Port of Sohar	OMN	24.51	56.61	1.7	0.206	-0.0018	Neither	Yes	23,224	Yes
Gwadar	PAK	25.11	62.34	1.4	0.035	-0.0004	Neither	Yes	1,538	Yes
Turbat	PAK	26.00	63.05	82.6	0.135	-0.0002	Neither	No	1,538	Yes
Lisbon	PRT	38.76	-9.23	7.8	2.121	-1.0234	Lost	Yes	24,621	Yes
Sines	PRT	37.96	-8.83	3.4	0.039	-1.0002	Lost	Yes	24,621	Yes
Setubal	PRT	38.50	-8.84	1.7	0.098	-0.9524	Lost	Yes	24,621	Yes
Horta	PRT	38.53	-28.62	1.2	0.018	0.0013	Neither	Yes	24,621	Yes
Praia da Vitoria	PRT	38.71	-27.05	1.1	0.043	0.0032	Neither	Yes	24,621	Yes
Ponta Delgada	PRT	37.74	-25.67	0.8	0.109	0.0072	Neither	Yes	24,621	Yes
Funchal	PRT	32.64	-16.91	1.3	0.217	0.8770	Gained	Yes	24,621	Yes
Porto Santo	PRT	33.06	-16.31	0.7	0.004	1.1232	Gained	Yes	24,621	Yes
Khan Yunis	PSE	31.32	34.30	5.7	1.904	-1.4187	Lost	No	3,800	No
Hebron	PSE	31.53	35.08	51.5	1.531	-0.9402	Lost	No	3,800	No
Nablus	PSE	32.22	35.26	40.0	0.713	-0.7178	Lost	No	3,800	No
Doha-Umm Said	QAT	24.96	51.57	2.3	1.505	-0.0020	Neither	Yes	88,701	Yes
Al Ruwais	QAT	26.15	51.21	1.2	0.005	-0.0014	Neither	Yes	88,701	Yes
Ras Laffan	QAT	25.91	51.59	2.3	0.625	-0.0013	Neither	Yes	88,701	Yes
Port Est	REU	-20.93	55.31	1.5	0.385	0.0041	Neither	Yes	–	Yes
Jeddah	SAU	21.46	39.17	0.5	2.524	-2.2948	Lost	Yes	38,510	Yes
Shoaiba	SAU	20.65	39.54	1.4	0.073	-2.2552	Lost	Yes	38,510	Yes
Duba (Port of Neom)	SAU	27.56	35.55	4.0	0.013	-2.1723	Lost	Yes	38,510	Yes
Yanbu (Yanbu city)	SAU	24.08	38.07	0.3	0.149	-2.1040	Lost	Yes	38,510	Yes
Rabigh	SAU	22.64	39.05	2.4	0.060	-1.9814	Lost	Yes	38,510	Yes
Jazan	SAU	16.88	42.54	0.2	0.232	-1.7665	Lost	Yes	38,510	Yes
Shuqaiq	SAU	17.63	42.08	2.7	0.070	-1.7552	Lost	Yes	38,510	Yes
Mecca	SAU	21.42	39.83	67.7	1.240	-1.7514	Lost	No	38,510	Yes
Abha	SAU	18.22	42.52	77.4	0.410	-1.2253	Lost	No	38,510	Yes
Al Ahsa	SAU	25.39	49.61	64.5	0.274	-0.0036	Neither	No	38,510	Yes
Ras Al Mishab	SAU	27.89	48.82	1.6	0.019	-0.0024	Neither	Yes	38,510	Yes
Ras Al Khafji (town)	SAU	28.41	48.52	0.1	0.060	-0.0024	Neither	Yes	38,510	Yes
Dammam 2nd Industrial City	SAU	26.26	49.97	8.2	1.554	-0.0022	Neither	Yes	38,510	Yes

City	Country	Lat	Lon	Shore, km	People, m	Exposure, sd	Side	Port	$/head	In
Ras Al Khafji (southern anchorage)	SAU	27.54	49.19	1.9	0.026	-0.0022	Neither	Yes	38,510	Yes
Jubail	SAU	27.06	49.66	3.3	0.251	-0.0020	Neither	Yes	38,510	Yes
Ras Tanura	SAU	26.67	50.12	0.7	0.345	-0.0018	Neither	Yes	38,510	Yes
Port Sudan	SDN	19.60	37.23	1.6	0.608	-1.7738	Lost	Yes	1,046	No
Suakin	SDN	19.12	37.35	1.1	0.069	-1.6434	Lost	Yes	1,046	No
Jubayt	SDN	18.95	36.83	57.3	0.122	-1.2222	Lost	No	1,046	No
Koungheul	SEN	13.98	-14.81	75.2	0.166	0.7241	Gained	No	1,574	Yes
Kolda	SEN	12.89	-14.94	49.6	0.185	0.8310	Gained	No	1,574	Yes
Kaffrine	SEN	14.10	-15.55	63.9	0.270	1.0882	Gained	No	1,574	Yes
Richard-Toll	SEN	16.46	-15.69	83.0	0.170	1.1592	Gained	No	1,574	Yes
Sedhiou	SEN	12.71	-15.56	1.7	0.216	1.1650	Gained	No	1,574	Yes
Darou Mousty	SEN	15.04	-16.05	88.1	0.615	1.4093	Gained	No	1,574	Yes
Louga	SEN	15.62	-16.23	37.4	0.228	1.5463	Gained	No	1,574	Yes
Ziguinchor	SEN	12.56	-16.27	3.0	0.308	1.6767	Gained	No	1,574	Yes
Fatick	SEN	14.33	-16.41	15.5	0.307	1.7078	Gained	No	1,574	Yes
Bambey	SEN	14.69	-16.45	54.7	0.386	1.7402	Gained	No	1,574	Yes
Saint-Louis	SEN	16.03	-16.48	5.8	0.301	1.7539	Gained	No	1,574	Yes
Thies	SEN	14.79	-16.92	27.9	0.759	2.1780	Gained	No	1,574	Yes
M'bour	SEN	14.43	-16.99	1.1	0.410	2.2452	Gained	No	1,574	Yes
Makeni	SLE	8.89	-12.05	93.9	0.276	0.5768	Gained	No	860	Yes
Bo	SLE	7.96	-11.74	63.0	0.441	0.7024	Gained	No	860	Yes
Mile 91	SLE	8.46	-12.22	72.8	0.160	0.7322	Gained	No	860	Yes
Kambia	SLE	9.12	-12.92	26.1	0.272	0.8067	Gained	No	860	Yes
Freetown	SLE	8.49	-13.20	1.0	1.378	1.1502	Gained	Yes	860	Yes
Bossaso	SOM	11.29	49.18	1.6	0.067	-2.0812	Lost	Yes	573	No
Erigavo	SOM	10.61	47.37	52.6	0.032	-1.4933	Lost	No	573	No
Berbera	SOM	10.43	45.00	1.1	0.021	-1.3124	Lost	Yes	573	No
Hargeisa	SOM	9.56	44.06	99.8	0.346	-0.6801	Lost	No	573	No
Garacad	SOM	6.93	49.32	2.1	0.011	-0.1439	Lost	Yes	573	No
Ceel Dheer	SOM	3.85	47.18	18.0	0.027	-0.0103	Neither	No	573	No
Jowhar	SOM	2.78	45.50	63.0	0.237	-0.0049	Neither	No	573	No
Wanlaweyn	SOM	2.62	44.89	81.0	0.124	-0.0046	Neither	No	573	No
Qoryooley	SOM	1.79	44.53	22.6	0.312	-0.0019	Neither	No	573	No
Mogadishu	SOM	2.02	45.33	2.2	1.272	-0.0019	Neither	Yes	573	No
Barawa	SOM	1.12	44.03	2.0	0.043	-0.0009	Neither	No	573	No
Jilib	SOM	0.49	42.77	46.5	0.063	-0.0006	Neither	No	573	No
Jamame	SOM	0.07	42.75	16.7	0.126	-0.0003	Neither	No	573	No
Kismayo	SOM	-0.37	42.55	1.8	0.068	-0.0001	Neither	Yes	573	No
Sao Tome	STP	0.35	6.65	6.4	0.177	0.0001	Neither	Yes	2,390	Yes
Victoria	SYC	-4.62	55.46	0.1	0.077	0.0189	Neither	Yes	16,683	Yes
Nawa	SYR	32.89	36.04	86.5	0.599	-0.4011	Lost	No	1,057	No
Damascus	SYR	33.36	36.25	84.8	3.724	-0.3112	Lost	No	1,057	No
Al-Zabadani	SYR	33.72	36.11	53.3	0.707	-0.2940	Lost	No	1,057	No
Tartous	SYR	34.90	35.87	2.1	0.487	-0.2087	Lost	Yes	1,057	No
Al Qusair	SYR	34.51	36.58	54.4	0.997	-0.1756	Lost	No	1,057	No
Latakia	SYR	35.52	35.77	0.9	0.738	-0.1673	Lost	Yes	1,057	No
Halfaya	SYR	35.25	36.58	56.8	1.016	-0.1304	Lost	No	1,057	No
Maarrat al-Nu'man	SYR	35.65	36.68	73.2	0.648	-0.1059	Lost	No	1,057	No
Aleppo	SYR	36.20	37.13	95.2	3.220	-0.0681	Lost	No	1,057	No
Tabligbo	TGO	6.60	1.50	44.3	0.290	0.0002	Neither	No	1,072	Yes
Notse	TGO	6.95	1.17	90.4	0.243	0.0003	Neither	No	1,072	Yes
Lome	TGO	6.14	1.28	1.8	2.535	0.0003	Neither	Yes	1,072	Yes
Kelibia	TUN	36.84	11.10	1.9	0.133	-1.8787	Lost	No	3,707	Yes
Bizerte	TUN	37.26	9.87	0.5	0.313	-1.8430	Lost	Yes	3,707	Yes
Tunis	TUN	36.81	10.28	1.5	2.285	-1.6602	Lost	Yes	3,707	Yes

City	Country	Lat	Lon	Shore, km	People, m	Exposure, sd	Side	Port	$/head	In
Hammamet	TUN	36.43	10.68	1.1	0.293	-1.5613	Lost	No	3,707	Yes
Beja	TUN	36.73	9.19	41.2	0.178	-1.4484	Lost	No	3,707	Yes
Jendouba	TUN	36.50	8.78	51.1	0.197	-1.2710	Lost	No	3,707	Yes
Sousse	TUN	35.82	10.65	1.4	0.917	-1.2623	Lost	Yes	3,707	Yes
Kairouan	TUN	35.67	10.09	50.7	0.251	-1.0607	Lost	No	3,707	Yes
Sfax	TUN	34.73	10.77	1.5	0.619	-0.8401	Lost	Yes	3,707	Yes
Sidi Bouzid	TUN	35.04	9.48	98.4	0.182	-0.7110	Lost	No	3,707	Yes
Skhira	TUN	34.34	10.14	2.0	0.039	-0.5948	Lost	Yes	3,707	Yes
Zarzis	TUN	33.48	11.11	0.9	0.070	-0.5116	Lost	Yes	3,707	Yes
Gabes	TUN	33.92	10.09	1.6	0.198	-0.4785	Lost	Yes	3,707	Yes
Medenine	TUN	33.35	10.49	26.2	0.122	-0.3983	Lost	No	3,707	Yes
Tataouine	TUN	32.92	10.45	67.1	0.076	-0.3145	Lost	No	3,707	Yes
Fethiye	TUR	36.65	29.15	3.0	0.175	-0.7082	Lost	No	10,898	Yes
Marmaris	TUR	36.86	28.26	1.1	0.047	-0.7075	Lost	No	10,898	Yes
Gulluk	TUR	37.25	27.61	0.6	0.186	-0.6023	Lost	Yes	10,898	Yes
Mugla	TUR	37.21	28.36	19.0	0.132	-0.5623	Lost	No	10,898	Yes
Port Akdeniz	TUR	36.84	30.61	0.4	0.987	-0.4805	Lost	Yes	10,898	Yes
Kusadasi	TUR	37.85	27.26	1.7	0.181	-0.4311	Lost	No	10,898	Yes
Manavgat	TUR	36.79	31.40	2.2	0.133	-0.4093	Lost	No	10,898	Yes
Emisbeleni	TUR	36.60	31.87	1.4	0.181	-0.4040	Lost	No	10,898	Yes
Nazilli	TUR	37.91	28.33	92.6	0.186	-0.3659	Lost	No	10,898	Yes
Tire	TUR	38.09	27.73	43.8	0.242	-0.3509	Lost	No	10,898	Yes
Cesme	TUR	38.32	26.30	1.0	0.068	-0.3338	Lost	Yes	10,898	Yes
Burdur	TUR	37.72	30.28	99.4	0.172	-0.2985	Lost	No	10,898	Yes
Izmir	TUR	38.44	27.16	0.8	2.745	-0.2946	Lost	Yes	10,898	Yes
Turgutlu	TUR	38.49	27.70	46.6	0.159	-0.2708	Lost	No	10,898	Yes
Salihli	TUR	38.48	28.14	85.5	0.133	-0.2591	Lost	No	10,898	Yes
Tasucu	TUR	36.31	33.89	0.6	0.066	-0.2450	Lost	Yes	10,898	Yes
Aliaga	TUR	38.80	26.93	0.9	0.054	-0.2335	Lost	Yes	10,898	Yes
Soma	TUR	39.19	27.61	55.6	0.219	-0.1703	Lost	No	10,898	Yes
Mersin	TUR	36.81	34.65	1.0	0.751	-0.1423	Lost	Yes	10,898	Yes
Akcay	TUR	39.58	26.92	0.4	0.232	-0.1349	Lost	No	10,898	Yes
Eregli	TUR	37.51	34.05	94.2	0.115	-0.1214	Lost	No	10,898	Yes
Balikesir	TUR	39.65	27.89	74.7	0.350	-0.1203	Lost	No	10,898	Yes
Samandag	TUR	36.09	35.98	3.0	0.570	-0.1185	Lost	No	10,898	Yes
Yenice	TUR	36.98	35.05	28.3	0.960	-0.1120	Lost	No	10,898	Yes
Tavsanli	TUR	39.55	29.50	99.0	0.098	-0.1056	Lost	No	10,898	Yes
Kirikhan	TUR	36.50	36.36	19.2	0.390	-0.0835	Lost	Yes	10,898	Yes
Toros Gubre	TUR	36.93	35.99	1.3	0.354	-0.0793	Lost	Yes	10,898	Yes
Bozuyuk	TUR	39.90	30.05	90.9	0.143	-0.0757	Lost	No	10,898	Yes
Kozan	TUR	37.46	35.82	61.7	0.130	-0.0652	Lost	No	10,898	Yes
Duzici	TUR	37.24	36.45	49.9	0.225	-0.0562	Lost	No	10,898	Yes
Kilis	TUR	36.72	37.11	80.3	0.337	-0.0545	Lost	No	10,898	Yes
Tanga	TZA	-5.06	39.12	1.6	0.308	0.0002	Neither	Yes	1,208	Yes
Zanzibar	TZA	-6.16	39.19	1.7	0.885	0.0003	Neither	Yes	1,208	Yes
Picha Ya Ndege	TZA	-6.75	38.92	31.1	0.168	0.0004	Neither	No	1,208	Yes
Dar Es Salaam	TZA	-6.84	39.30	0.8	5.538	0.0005	Neither	Yes	1,208	Yes
Mtwara	TZA	-10.26	40.20	1.1	0.191	0.0034	Neither	Yes	1,208	Yes
Mokha	YEM	13.30	43.23	1.8	0.038	-3.1013	Lost	Yes	–	No
Aden	YEM	12.78	44.95	2.6	0.870	-2.9010	Lost	Yes	–	No
Al Misrakh	YEM	13.46	44.01	78.9	1.681	-2.6724	Lost	No	–	No
Al Jarrahi	YEM	14.13	43.39	30.5	0.402	-2.5971	Lost	No	–	No
Hudaydah (Hodeidah)	YEM	14.83	42.94	0.9	0.678	-2.5482	Lost	Yes	–	No
‘Adan Barakah	YEM	13.84	43.80	57.1	0.394	-2.5338	Lost	No	–	No
Salif	YEM	15.32	42.67	1.1	0.040	-2.5079	Lost	Yes	–	No

City	Country	Lat	Lon	Shore, km	People, m	Exposure, sd	Side	Port	$/head	In
Ja‘ar	YEM	13.22	45.31	19.2	0.063	-2.4845	Lost	No	–	No
Bayt al Faqih	YEM	14.52	43.32	31.7	0.387	-2.4289	Lost	No	–	No
Ibb	YEM	13.95	44.15	97.5	1.610	-2.2761	Lost	No	–	No
Ad Dala	YEM	13.70	44.73	91.3	0.426	-2.2225	Lost	No	–	No
Az Zuhrah	YEM	15.73	43.01	26.5	0.302	-2.0072	Lost	No	–	No
Madinat Lab’us	YEM	13.85	45.24	72.2	0.379	-1.9763	Lost	No	–	No
Lawdar	YEM	13.88	45.86	54.4	0.225	-1.8378	Lost	No	–	No
Suq Shamar	YEM	15.90	43.38	58.0	0.994	-1.6628	Lost	No	–	No
Rudum Terminal	YEM	13.99	47.91	2.4	0.014	-1.6605	Lost	Yes	–	No
Harad	YEM	16.41	43.06	28.2	0.322	-1.6311	Lost	No	–	No
Mukalla	YEM	14.52	49.15	0.8	0.224	-1.2302	Lost	Yes	–	No
Ash Shihr	YEM	14.69	49.54	2.5	0.112	-1.1024	Lost	Yes	–	No
KaMaqhekeza	ZAF	-25.68	31.83	73.0	0.308	0.0290	Neither	No	6,534	Yes
Richards Bay	ZAF	-28.79	32.04	2.1	0.406	0.1372	Gained	Yes	6,534	Yes
Stanger / KwaDukuza	ZAF	-29.33	31.28	7.2	0.383	0.1571	Gained	No	6,534	Yes
Pietermaritzburg	ZAF	-29.63	30.36	68.0	0.752	0.1617	Gained	No	6,534	Yes
Durban	ZAF	-29.88	31.02	1.6	2.752	0.2021	Gained	Yes	6,534	Yes
Mthatha	ZAF	-31.60	28.77	54.9	0.412	0.4428	Gained	No	6,534	Yes
East London	ZAF	-33.03	27.90	0.2	0.487	0.9974	Gained	Yes	6,534	Yes
Makhanda	ZAF	-33.30	26.55	43.7	0.068	1.1435	Gained	No	6,534	Yes
Oudtshoorn	ZAF	-33.60	22.22	51.7	0.079	1.4422	Gained	No	6,534	Yes
Ngqura	ZAF	-33.81	25.69	1.9	0.893	1.5137	Gained	Yes	6,534	Yes
Fraaisig	ZAF	-34.04	23.09	3.3	0.101	1.7553	Gained	No	6,534	Yes
Mossel Bay	ZAF	-34.18	22.15	0.1	0.096	1.9312	Gained	Yes	6,534	Yes
Worcester	ZAF	-33.65	19.45	76.9	0.163	2.0049	Gained	No	6,534	Yes
Malmesbury	ZAF	-33.46	18.72	34.9	0.102	2.1881	Gained	No	6,534	Yes
Saldanha Bay	ZAF	-33.00	18.01	1.1	0.120	2.2409	Gained	Yes	6,534	Yes
Stellenbosch	ZAF	-33.93	18.86	19.0	0.983	2.4549	Gained	Yes	6,534	Yes
Cape Town	ZAF	-33.92	18.45	1.4	2.474	2.6471	Gained	Yes	6,534	Yes