WORKING paper

BANQUE DE FRANCE
EUROSYSTÈME

Navigating Shocks: The Ripple Effects of Shipping Route Closures

Carsten Philipp Brockhaus, Julian Hinz & Charles Serfaty[1]

August 2026, WP 1057

ABSTRACT

Who bears the cost when a maritime chokepoint is disrupted? Combining ship-level AIS trajectories, port-call logs, freight indices, customs data and Turkish shipment-level data, we document the trade effects of the 2023 Red Sea crisis: seaborne trade between affected country pairs fell, rerouted around the Cape of Good Hope, and recovered within months as freight costs surged, while some exporters shifted persistently toward air freight. We build a quantitative trade model in which freight costs are endogenous to modal and route choice through congestion, a global shipping capacity constraint, and monopoly toll-setting by canal authorities. Calibrated to observed route choices and canal revenues, the model implies that a permanent Red Sea closure lowers global real income only mildly. Because the Suez Canal Authority, pricing to maximize revenue, already captures much of the route's surplus, the loss falls on the toll collector: Egypt forgoes 3.0 percent of real income, almost all of it canal rent, while the large trading economies each lose less than 0.02 percent. A transit fee on the Strait of Hormuz, a chokepoint without maritime substitutes, instead concentrates losses on the Gulf economies themselves. Whether a chokepoint has substitutes, and who prices it, jointly determine who pays for its disruption.

Keywords: Shipping Routes; Chokepoint Disruptions; Red Sea Crisis; Endogenous Trade Costs; Monopoly Tolls.

JEL classification: F14, F17, R41, F62.

---

[1] Carsten Philipp Brockhaus: Kiel Institute for the World Economy and Kiel University; Julian Hinz: Bielefeld University and Kiel Institute for the World Economy; Charles Serfaty: Banque de France.
We thank participants at the European Trade Study Group, the PSE Shipping Workshop, the Aarhus–Kiel Workshop, and the IFC–IMIP–World Bank Workshop, as well as seminar audiences in Göttingen and Aachen, for valuable comments.

NON-TECHNICAL SUMMARY

Maritime trade depends on a small number of narrow passages, but chokepoints that look similar on a map can be economically very different. Suez and Panama offer long but feasible detours; Hormuz places the Gulf economies close to a maritime dead end; and some passages are priced infrastructures whose operators collect tolls. This paper asks how substitute routes, interactions between chokepoints and the pricing of passage determine the effects of disruption on trade, freight costs and real income.

The 2023 Red Sea crisis provides direct evidence. The authors combine ship-level position and port-call records with freight indices, maritime-distance measures, monthly customs data for the European Union and the United States, and shipment-level Turkish customs records. Following the first Houthi attacks in November 2023, transits through the Red Sea collapsed by two thirds, traffic around the Cape of Good Hope increased and travel times lengthened. Sea trade between country pairs dependent on Suez fell by more than 10 percent relative to unaffected pairs, but recovered within months as carriers reorganized routes. Turkish firms show the same two-speed adjustment: maritime exports to via-Suez partners fell by 5 percent starting three weeks after the first attack, then recovered within a quarter. The air share of their exports subsequently increased by 0.7 percentage point on average, reaching 1.4 points above the comparison group after ten weeks.

To interpret these facts, the paper develops a multi-country, multi-sector quantitative trade model with endogenous transport costs. Shippers choose across modes and routes; rerouting creates congestion on alternative passages; longer voyages tighten global shipping capacity; and canal authorities set monopoly tolls. The model is disciplined by observed route choices, sectoral transport patterns and canal revenues. It reproduces the Suez Canal Authority's $10.3 billion of toll revenue in 2023.

The first counterfactual closes the Red Sea route permanently. Seaborne flows rise by about 18 percent around the Cape of Good Hope and 20 percent through Panama. Global real-income losses remain mild because alternative routes exist and freight represents only a modest share of the delivered value of most goods. The large European and Asian trading economies each lose less than 0.02 percent. Egypt bears by far the largest loss, 3.0 percent of real income: 2.8 percentage points come directly from the disappearance of Suez toll rents, while only 0.3 point reflects the residual trade-cost and general-equilibrium effects.

A comparison with a no-toll baseline clarifies the incidence. Without Suez and Panama rents, Egypt's loss falls to 0.2 percent, while losses become larger for economies whose cargo must travel farther: Qatar's loss rises from 0.2 to 1.1 percent, for example. Removing the rents leaves the global cost of closure broadly similar but spreads the burden away from the collector and towards route users. The toll therefore cushions shippers against the detour while making the Canal Authority the residual bearer of disruption risk.

The contrast with the Strait of Hormuz is stark: Gulf economies have no maritime detour. At the intermediate transit fee considered in the paper, which raises freight costs on exposed trade by 35 percent, real income falls by 3.9 percent in Qatar, 2.7 percent in Kuwait, and 2.5 percent in Iraq, while Hormuz-transiting trade declines by 60 percent. Although toll revenue is split equally in our exercise, Oman gains 12.2 percent and Iran only 1.5 percent because their exposure to the disruption differs sharply. Large third countries experience only limited losses. This incidence pattern partly reflects the model's high long-run elasticities of substitution: importers can shift demand toward unaffected suppliers at relatively low cost, pushing much of the burden onto Gulf producers through lower pre-fee export prices and reduced sales. Competing oil and gas exporters consequently expand production, such as Brunei that experiences 0.8 percent gain in real income. These results should be interpreted as long-run equilibrium effects; short-run costs to importers could be larger when substitution across suppliers is slower or more constrained.

Vulnerability depends not only on transit volumes, but also on the availability of alternative suppliers and routes and on who controls and prices the passage.

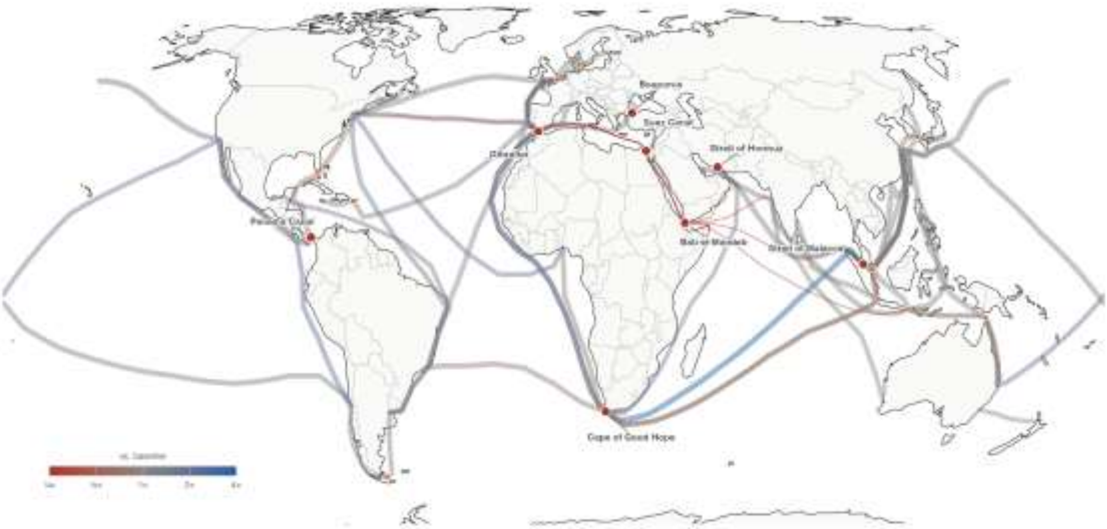

Seaborne trade flows under a permanent Red Sea closure

Line width is proportional to the log of post-shock shipment value; colours show the change relative to the baseline (red: decline; blue: increase). The closure shifts traffic towards the Cape of Good Hope and the Panama Canal.

Naviguer face aux chocs : les effets en cascade des fermetures de routes maritimes

RÉSUMÉ

Cet article étudie les conséquences locales et mondiales des perturbations aux points de passage maritimes stratégiques. À partir de trajectoires AIS, d'escales portuaires, d'indices de fret, de données douanières et de données microéconomiques turques, nous montrons les effets la crise de la mer Rouge : recul temporaire du commerce maritime entre partenaires exposés avant un recouvrement en quelques mois, contournement par le cap de Bonne-Espérance et hausse des coûts de fret. Nous construisons ensuite un modèle quantitatif où les coûts de transport dépendent du choix des routes, de la congestion, des capacités mondiales de transport maritime et des péages de canal. Une fermeture permanente de la mer Rouge réduit peu le revenu réel mondial : les grandes économies perdent moins de 0,02 % de revenu, tandis que l'Égypte perd 3,0 %, presque entièrement du fait de la disparition des recettes de Suez. À l'inverse, un péage dans le détroit d'Ormuz, sans itinéraire maritime de substitution pour les économies du Golfe qui y sont soumises, concentre les pertes sur ces dernières. La possibilité de contourner le détroit et la tarification des passages déterminent donc qui supporte le coût des perturbations.

Mots-clés : routes maritimes ; goulets d'étranglement ; crise de la mer Rouge ; coûts du commerce endogènes ; péages.

---

1 Introduction

As the *Ever Given* accident in the Suez Canal in 2021, the Houthi attacks in the Red Sea in 2023 and 2024, the drought-related restrictions in the Panama Canal in 2023, and the transit fees mooted for the Strait of Hormuz in 2026 have all made clear, a large share of international trade depends on a small number of maritime chokepoints. The lesson from these episodes could be simple: closing a narrow passage raises trade costs and disrupts supply chains. However, not all chokepoints are alike. Some, such as Malacca, are part of dense route networks. Others, such as Hormuz or the Danish straits, constrain trade in regions that are close to maritime dead-ends. Still others, such as Suez and Panama, have long but feasible detours. Some are interdependent: a shock near Bab-el-Mandeb operates economically as a shock to the Suez Canal, because ships cannot use one without the other on the Europe–Asia route. And some chokepoints are not just geographic features but priced infrastructures, whose owners charge for passage.

This paper studies how these differences shape the trade and welfare consequences of maritime chokepoint shocks. We reconstruct maritime exposure at a granular level by combining bilateral trade data, information on the modal composition of trade, ship-level AIS trajectories and port-call logs, and route distances computed on a maritime network. These data allow us to gauge the impact of the Red Sea crisis on trade flows and shipping patterns. We then build a quantitative trade model in the spirit of Caliendo and Parro (2015) in which the bilateral trade wedge is endogenous to route choice, with three additional cost layers: congestion on alternative passages, a global shipping capacity constraint, and monopoly toll-setting by canal authorities. The model lets us compare shocks that look similar on a map but differ economically — disruptions to routes with close maritime substitutes, to routes without good substitutes, to chokepoints that are complements, and to canal tolls.

The Red Sea crisis is an instructive laboratory. Following the first Houthi attacks in November 2023 (Figure 1), seaborne trade between country pairs dependent on the Suez passage declined by more than 10 percent relative to unaffected pairs, while air trade showed no comparable drop. The AIS data reveal the mechanics (Figure 2): transits through the Red Sea collapsed by two thirds, traffic around the Cape of Good Hope picked up, and travel durations between affected country pairs shifted markedly to the right. Within months, however, trade between affected pairs had recovered to pre-crisis levels relative to unaffected pairs. Rerouting absorbed the shock; the freight market bore it.

The quantitative model rationalizes this pattern and traces its welfare incidence. We calibrate the route-choice elasticity that disciplines monopoly pricing to observed canal revenues: the revenue-maximizing Suez fee reproduces the Suez Canal Authority's \$10.25 billion of toll revenue in 2023, and — as an untargeted check — the implied Panama toll

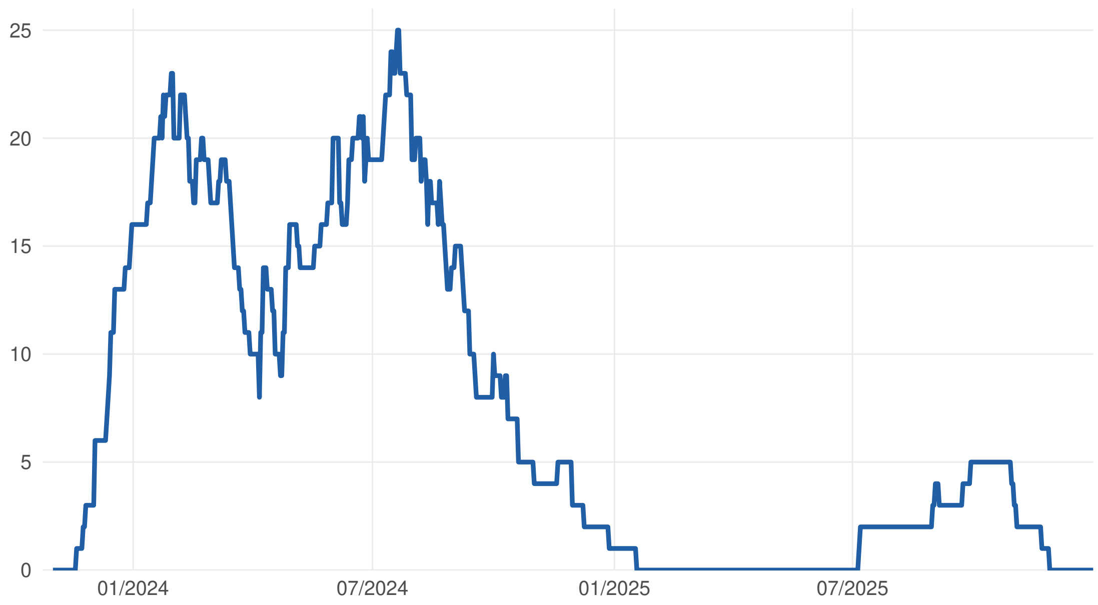

Figure 1: Number of Houthi attacks on commercial vessels.

Notes: Documented attacks over the preceding 60 days, whether or not they led to damage. Source: `https://en.wikipedia.org/wiki/Houthi_attacks_on_commercial_vessels`, revisions as of January and July 2025.

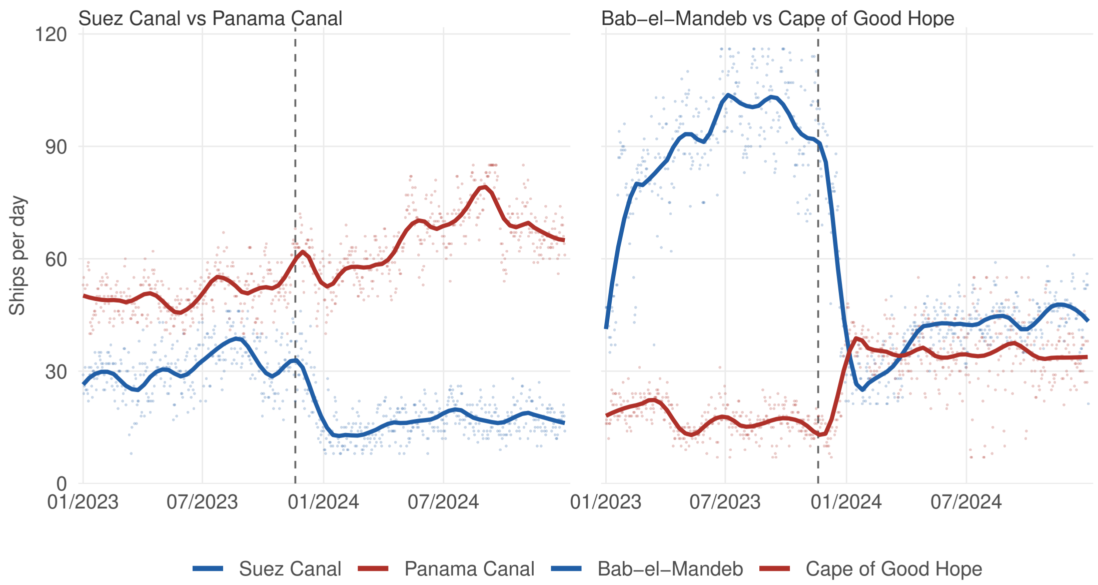

Figure 2: Daily container-ship counts at maritime chokepoints and their alternatives.

revenue of \$3.62 billion is close to the observed \$3.18 billion. Under a permanent Red Sea closure, global losses are then mild: cargo reroutes around the Cape of Good Hope at a freight penalty that is small relative to the value of the goods carried, and the large European and Asian economies each lose less than 0.02 percent of real income. The bulk of the loss falls on Egypt, at 3.0 percent of real income, almost entirely through the collapse of canal-toll revenue. Because a canal authority that charges revenue-maximizing fees already captures a large part of the surplus created by the shorter route, the disruption of such a route is, in welfare terms, a shock to the toll collector rather than to the shippers. The same logic does not extend to chokepoints without alternatives. A transit fee on the Strait of Hormuz — where, for the Gulf economies, no maritime detour exists — falls almost entirely on the Gulf itself: at the intermediate fee we consider, Qatar loses 3.9 percent of real income, Kuwait 2.7 percent, and Iraq 2.5 percent, while third countries substitute towards other suppliers and lose almost nothing. Even the two collectors fare differently: Oman, whose own trade bypasses the strait, receives a nearly pure transfer, while Iran levies much of the fee on its own trade.

This distinction matters for policy. Treating all chokepoints as identical overstates the welfare cost of shocks to priced canals with feasible detours and understates the exposure of regions behind passages without substitutes. The contribution of the paper is hence twofold. Conceptually, we show that the welfare cost of a maritime shock depends on the type of chokepoint — the existence of substitute routes, the interaction between chokepoints, and the pricing rule of the infrastructure owner. Quantitatively, we use reconstructed maritime trade flows, route-level congestion estimates, and observed canal revenues to discipline each of these mechanisms within a unified framework. Firm-level customs micro-data sharpen the quantitative side of this: for Turkey, a country whose eastern-Mediterranean position placed it beside the Suez approach, we show that the quantity of seaborne exports to disrupted partners fell only transitorily, whereas the modal shift toward air freight outlasted the shock.

The remainder of the paper is structured as follows: Section 2 relates the paper to the literature. Section 3 documents stylized facts of the Red Sea crisis. Section 4 presents the model, Section 5 its calibration, and Section 6 the counterfactual simulations. Section 7 concludes.

2 Related Literature

The paper is related to two strands of literature. The first strand studies distance, freight, and the persistent role of transport costs in international trade. The persistent relevance of distance despite technological progress in transport is documented in long-run trade data and in modern gravity applications Head and Mayer (2014); Fouquin and Hugot (2016); Mayer and Zignago (2011); Conte et al. (2022). This holds despite historical work showing that declining trade costs and the maritime transport revolution were important features of globalization. For the 19th century, Pascali (2017) finds positive impacts of the maritime freight revolution through an instrumental-variable approach, although the more direct approach of Jacks and Pendakur (2010) finds no specific impact of lower maritime freight costs on bilateral trade. For the containerization revolution, Hummels (2007) provides an overview of its effects, which compensated for higher costs from oil shocks and raised the speed and consistency of travel, and Bernhofen et al. (2016) provides quantitative evidence for its positive impact on international trade. The use of time and distance variation in effective geography to estimate trade effects goes back to Feyrer (2019). More generally, transportation is not a passive iceberg cost: the geography and organization of the transport sector shape trade responses, as emphasized by Brancaccio et al. (2020). A growing literature endogenizes transport costs in general equilibrium: Asturias (2020) models an oligopolistic shipping industry within a quantitative trade model, Wong (2022) shows that round-trip economics make freight rates respond to trade imbalances, Kalouptsidi (2014) provides the structural time-to-build foundation for short-run shipping capacity constraints, and Hummels et al. (2009) document that market power in international shipping reduces trade — the antecedent of our monopoly toll channel. Cristoforoni et al. (2025) extend this margin to two-sided market power: freight prices are bargained between oligopolistic carriers and large importers, and the endogenous response of transport costs dampens the welfare cost of tariffs. On land, Fajgelbaum and Schaal (2020) and Allen and Arkolakis (2022) study congestion and endogenous routing on transport networks in spatial equilibrium. Time in transit is itself a trade cost: Hummels and Schaur (2013) put each additional day at sea at 0.6–2.1 percent of the cargo's value, and Djankov et al. (2010) find that a day of delay reduces trade by more than one percent. Ganapati and Wong (2023) document how falling transport costs and lengthening supply chains have made world trade more dependent on reliable transportation—and more exposed to its disruption. Closest to our transport block, Tolva (2026) estimates the substitution between transport modes from the closure of Russian airspace, Ko et al. (2025) provide sectoral estimates of the same margin, and Fuchs and Wong (2026) embed mode choice with congestion at intermodal terminals in a quantitative spatial model. In contemporaneous work, Massoni (2025) links Brazilian bill-of-lading data to tropical-cyclone tracks and embeds port-level congestion in a spatial production-network model with endogenous routing, arguing that maritime bottlenecks arise at nodes—ports and terminals—rather than on links. Our focus is complementary: the disruptions we study are link-level, at canal chokepoints, and the cost of passage there is not a congestion outcome but is set endogenously by a revenue-maximizing canal authority. That paper's firm-level finding that adaptation runs through rerouting—for roughly two months, without severing buyer–supplier relationships—corroborates the routing-not-sourcing pattern we document for Turkish trade in Section 3.3. Ganapati et al. (2024) document that a large share of trade is routed indirectly through entrepôt hubs — the routing margin our route-choice block operates on.

The second strand studies large transport infrastructures and maritime chokepoints. The Suez Canal closure from 1967 to 1975 is used as a natural experiment for the effect of distance on trade and income in Feyrer (2021). Historical work on the Suez and Panama canals shows that openings and closures altered trade costs and trade patterns (Hugot and Umana Dajud, 2016; Jacks et al., 2024). The Panama Canal also had long-run effects on the economic geography of the United States: Maurer and Rauch (2023) shows that counties gaining market access through the canal experienced stronger demographic growth. Heiland et al. (2025) estimate the positive trade effects of the Panama Canal expansion for country pairs whose shortest maritime path already ran through the canal, and Chowdhry et al. (2026) use the same expansion to estimate the elasticity of substitution between ship types in a quantitative trade model with a rich transportation sector. The 2023–24 Red Sea crisis itself has begun to attract attention: Notteboom et al. (2024) provide a maritime-economics account of the ensuing network reconfiguration, Attinasi et al. (2024) gauge the fallout for global trade and inflation and find muted effects given spare shipping capacity, and Verschuur et al. (2025) quantify the trade exposed at maritime chokepoints on the global port network. Closest to our paper, Dunn and Leibovici (2025) build a dynamic general-equilibrium model of the global shipping capacity market and apply it to the Red Sea disruptions. Relative to their aggregate-capacity focus, we model route choice at the corridor level, the complementarity between chokepoints, and the toll-setting of the canal authorities. In contemporaneous work, Fuchs et al. (2026) study chokepoint disruptions in a spatial equilibrium model with endogenous routing, deriving country-level measures of strategic vulnerability from asymmetric disruption exposure and the build-versus-operate structure of the global container fleet. Their disruptions and policy fees enter as exogenous cost wedges, whereas we endogenize the cost of passage itself. Our contribution is hence to integrate chokepoints as heterogeneous objects within the route network, and to endogenize the behavior of the canal authorities.

3 Stylized Facts of a Shipping Route Closure

3.1 Shipping impact

Data We observe 6,032 IMO-registered container vessels from August 2021 to August 2024 at regular intervals. The AIS data link each vessel to arrival, departure, and berth times, letting us reconstruct realized routes and route-specific deadweight tonnage. Because the panel covers container vessels, the route shares and the route-choice elasticity we estimate from it reflect container routing; tanker and dry-bulk cargo—which dominates some chokepoints, notably the Strait of Hormuz—reroutes on less elastic, more contract-bound terms, a limitation we return to in the counterfactuals.

The two tables below summarize the fleet and the port network it spans: Table 1 reports the port-call and sea-leg statistics, and Table 2 the structure of the port-call network before and after the crisis. The vessels in our sample make 1.2 million port calls at 1,593 ports

Table 1: Summary statistics for AIS port calls and sea legs

Variable	N	Mean	SD	Min	P25	Med	P75	Max
Port calls per vessel	5,985	205.05	156.77	1	117	166	251	1,376
Dwell time (days)	1,227,205	1.25	3.39	0.04	0.44	0.75	1.23	521.66
Vessel deadweight tons	5,606	57,044	51,515	998	18,830	37,621	81,098	272,405
Vessel capacity (TEU)	5,442	4,670	4,721	70	1,284	2,732	6,500	24,116
Sea-leg duration (days)	1,220,568	3.39	5.26	0	0.63	1.81	3.88	89.83
Port pair voyages	30,647	39.83	180.67	1	1	3	17	7,767

Notes: Port calls are collapsed to consecutive cluster-port visits. Sea legs longer than 90 days are treated as data gaps and excluded.

Table 2: Port-call network summary before and after the Red Sea crisis

Network	Ports	N links	Density	Mean degree	Mean path	Diameter
Entire network	1,593	30,647	0.0121	38.50	5.267	8
Before Red Sea crisis	1,478	26,413	0.0121	35.70	5.349	10
Following Red Sea crisis	1,297	17,495	0.0104	27.00	6.112	9

Notes: Network links are directed port-to-port sea legs. Sea legs longer than 90 days are treated as data gaps and excluded.

over the observation window, covering 30,647 directed port pairs. The fleet spans the full size distribution of the container market, from feeder vessels below 1,000 deadweight tons to carriers of 24,000 TEU, with a median capacity of 2,732 TEU. Port stays are short — the median dwell time is 0.75 days — and the median sea leg between two consecutive ports lasts 1.8 days. The structure of the network shifts visibly with the onset of the Red Sea crisis: the port-call network becomes sparser, with network density falling from 0.0121 to 0.0104 and the mean degree from 38 to 27, while the mean shortest path between ports lengthens from 5.3 to 6.1 legs — traffic concentrates on fewer, longer corridors.

Analysis The first Houthi attacks took place at the beginning of the war in Gaza triggered by the Hamas-led attacks of October 7, 2023. The first of them was the hijacking of the *Galaxy Leader* by Houthi fighters on November 19, 2023. In total, 95 attacks were documented between October 2023 and September 2025, when a second ceasefire was agreed, as Figure 1 shows. Yemeni Houthi rebels targeted vessels in the Red Sea that they suspected of providing support to Israel.

As a consequence, ship operators rerouted their vessels, increasingly opting for the long journey around the Cape of Good Hope to reach Europe. This is reflected in the AIS position data in Figure 2: the number of ships traversing the Red Sea and the Suez Canal declines dramatically after the onset of the attacks, while traffic picks up around the Cape of Good Hope. Traffic at other bottlenecks such as the Panama Canal or the Strait of Taiwan remains roughly constant during this time.

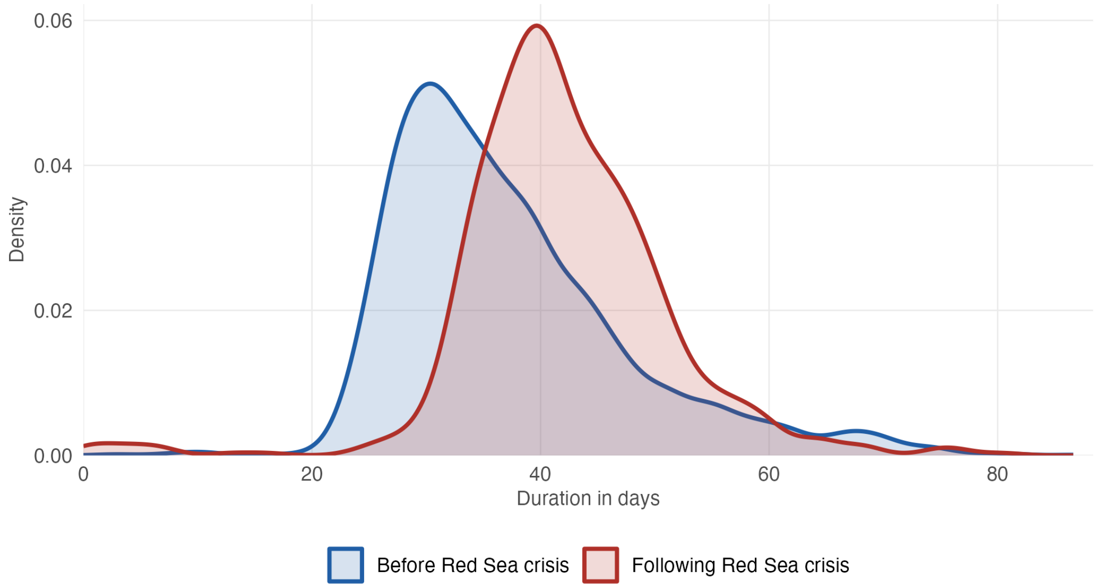

Figure 3: Travel duration between China and the Netherlands: unconditional distribution.

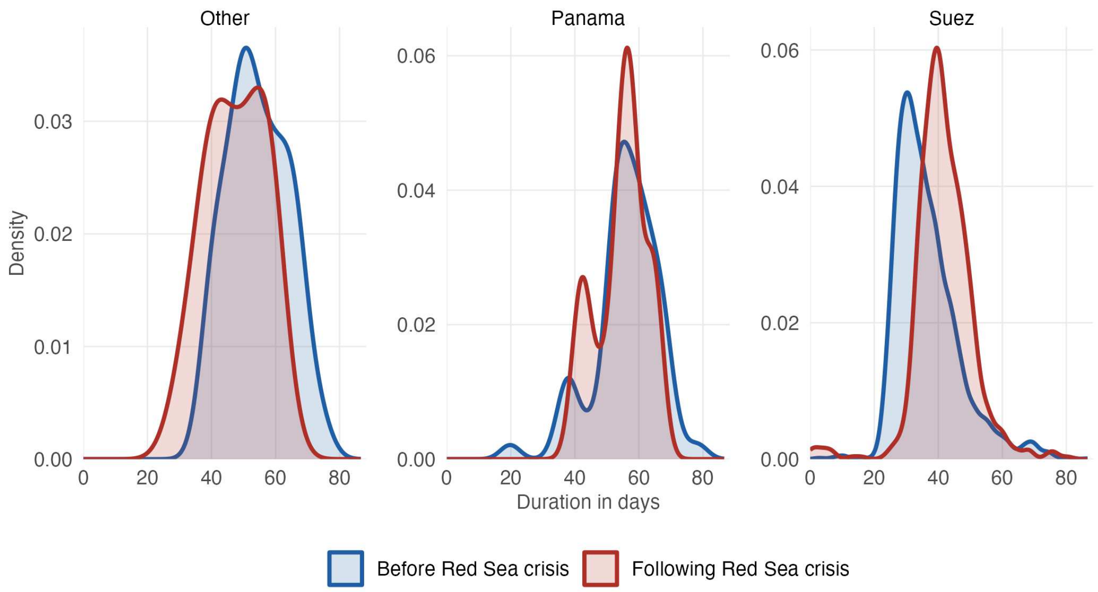

Figure 4: Travel duration between China and the Netherlands, by passage taken.

Transit durations. The Red Sea crisis had differential effects on transit durations between country pairs. In the case of a directly affected country pair — China and the Netherlands — the distribution of travel durations shifts distinctly to the right following the crisis (Figure 3). Conditioning on the passage taken sharpens the picture: routes that previously relied on the Suez Canal show a strong rightward shift in travel durations (Figure 4). These shifts are economically meaningful: standard estimates put the cost of an additional day in transit at 0.6–2.1 percent of the cargo's value (Hummels and Schaur, 2013).

3.2 Trade impact

Data We use monthly customs data for the United States and the European Union. For the United States, we rely on U.S. Census Bureau foreign trade data, which report imports and exports by partner country, HS product, port or district, and transport mode. For the European Union, we use Eurostat Comext, which reports monthly trade in goods by reporter, partner, product, flow, value, quantity, and mode of transport. In both sources, we work at the HS6-month-reporter-partner-flow level and retain information on the transport mode.

The second ingredient is a measure of bilateral maritime distance, which we reconstruct so as to capture precisely the additional distance that a Red Sea closure imposes on each country pair. We compute shortest maritime routes with the SeaRoute algorithm of Gaffuri and Eurostat (2022), which runs a Dijkstra shortest-path search on a global maritime-network graph—built from the Oak Ridge National Laboratory shipping-lane data and enriched with the Suez, Panama, and Bab-el-Mandeb passages—whose edges are weighted by their great-circle length. For each ordered pair of countries we compute this distance twice, once with all passages open and once with Bab-el-Mandeb closed, so that the difference isolates the distance shock of a Red Sea disruption. We then aggregate port-to-port distances to the country level using port weights from IMF PortWatch, which reports, for each port p, the share x_p of a country's maritime exports leaving from it and the share $m_{p'}$ of its maritime imports arriving through it. Rather than assuming that all trade between i and j passes through their largest ports, we let the weight on a pair of ports fall with the maritime distance between them, so that the country-to-country distance is

$$d_{i,j} = \frac{\sum_{p\in\mathcal{P}(i),\, p'\in\mathcal{P}(j)} \frac{x_p\, m_{p'}}{d^{\alpha}_{p,p'}}\, d_{p,p'}}{\sum_{p\in\mathcal{P}(i),\, p'\in\mathcal{P}(j)} \frac{x_p\, m_{p'}}{d^{\alpha}_{p,p'}}}, \tag{1}$$

where $\mathcal{P}(i)$ is the set of ports in country i and α governs the elasticity of trade to distance. We set $\alpha = 1$, in line with the gravitational distance elasticity estimated by Head and Mayer (2014). Landlocked countries, for which no maritime distance is defined, are excluded from the analysis.

Aggregate patterns Combining import data for the United States with import and export data for the European Union, we construct a partial index of world trade by mode of transport and trace its evolution in the aftermath of the Houthi attacks. Figure 5 shows the evolution of maritime, air, and total trade values in the combined Census and Comext data, distinguishing between country pairs for which a Red Sea closure increases bilateral distance and those for which it does not.[1] Sea trade between pairs dependent on the Red

[1]Landlocked countries are excluded from the analysis. In this descriptive graph, a pair of countries is affected by the Red Sea closure if the bilateral distance between the two countries increases by more than 10 kilometers.

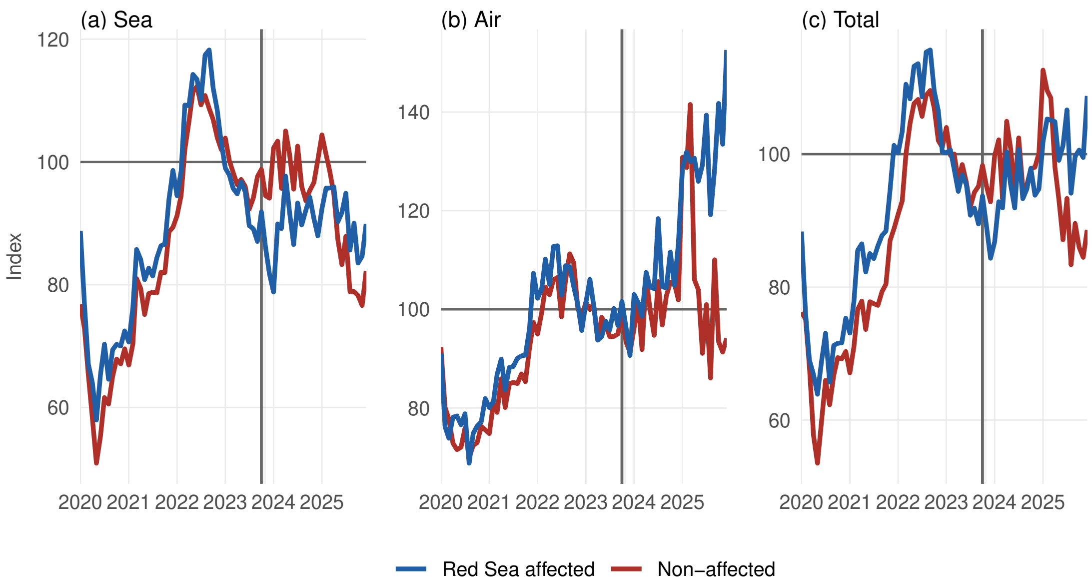

Figure 5: Aggregate trade from Comext and Census data, by transport mode and Red Sea exposure.

Notes: Values in euros, indexed to 100 for the twelve months before October 2023. Series are deseasonalized using data from 2016 to 2022.

Sea passage declined by more than 10% in December 2023 relative to September 2023, with no comparable decline for air trade. Total trade, which also includes other modes of transport, shows a smaller decline in the months that followed. Trade returned to pre-crisis levels, at least relative to unaffected pairs. We read this as suggestive evidence that the Red Sea shock was short-lived, and confirm it in a difference-in-differences regression below.

Difference-in-differences evidence We confirm these aggregate patterns in an event study that exploits the continuous variation in maritime exposure across country pairs and products. Let i denote the exporting country, j the importing country, s the HS6 product, and t the month. Event time τ is dated relative to the beginning of the conflict, so that October 2023 is $\tau = 0$ and September 2023 ($\tau = -1$) is the omitted reference month. Exposure combines two margins. The first is the additional distance a Red Sea closure imposes on a country pair, $D_{ij} \equiv (d^c_{ij} - d^o_{ij})/9000$, where d^o_{ij} and d^c_{ij} are the maritime distances of (1) with the Red Sea route open and closed; the normalization by $9{,}000$ kilometers—roughly the detour faced by China–Europe trade—sets one unit of exposure to that benchmark. The second is the pre-shock maritime intensity of the product-route, the predicted sea share π^{sea}_{ijs} measured on 2022 transport-mode patterns. The interaction $X_{ijs} = D_{ij} \times \pi^{sea}_{ijs}$ is our treatment variable: it is large only for goods normally shipped by sea on routes whose maritime distance rises sharply when the Red Sea closes, and close to zero otherwise.

We estimate

$$\begin{aligned} y_{ijst} = &\sum_{\tau \neq -1} \beta_\tau \mathbb{I}\{t = \tau\} \left(D_{ij} \times \pi_{ijs}^{sea}\right) \\ &+ \sum_{\tau \neq -1} \theta_\tau \mathbb{I}\{t = \tau\} D_{ij} \\ &+ \sum_{\tau \neq -1} \lambda_\tau \mathbb{I}\{t = \tau\} \pi_{ijs}^{sea} \\ &+ \alpha_{ijs} + \alpha_{jst} + \alpha_{it} + \varepsilon_{ijst}, \end{aligned} \tag{2}$$

where the dependent variable y_{ijst} is the log value, log quantity, or log unit value of trade, or the sea share in percentage points. The fixed effects are demanding: α_{ijs} absorbs time-invariant differences across product-level bilateral relationships, α_{jst} absorbs shocks specific to an importer–product–month, and α_{it} absorbs shocks common to an exporter in a given month. Standard errors are clustered at the reporter–partner level.[2] The specification also lets each exposure margin follow its own dynamic path, through separate event-time interactions of $\mathbb{I}\{t = \tau\}$ with D_{ij} and with π_{ijs}^{sea}, so that routes with a large distance shock and goods with a high predicted sea share are each allowed their own trend regardless of the other.

The coefficients of interest, $(\beta_\tau)_{\tau \neq -1}$, trace the differential monthly effect of a one-unit increase in $D_{ij} \times \pi_{ijs}^{sea}$—that is, of a $9{,}000$-kilometer distance shock for a good shipped entirely by sea before the crisis. We report $\tilde{\beta}_\tau = 100 \times (\exp(\beta_\tau) - 1)$, the effect in percent. Figure 6 shows no clear pre-trend and a single significant decline in December 2023, when ships began rerouting around the Cape and delaying their voyages; trade then returns to its pre-crisis path, with no rebound.

The demanding fixed effects absorb much of the general-equilibrium adjustment, consistent with a local shock whose consequences spill beyond the affected pairs to the rest of the world. Read together with the aggregate patterns, the event study indicates that the effect of the Red Sea disruption on trade between affected pairs was concentrated in the first months and had largely faded within half a year.

3.3 Firm-level evidence

A central node relegated to a backwater The aggregate and event-study evidence above pools thousands of country pairs, and so speaks to the average consequence of the closure rather than to the experience of any single place. We now zoom in on a single node whose position makes it especially informative. Turkey sits in the eastern Mediterranean, immediately beside the northern approach to the Suez Canal, astride the trunk lane that

[2] We omit exporter–product–month effects α_{ist}: with mostly European importers at similar sea distance from Asia, they would largely absorb the causal effect of the Red Sea closure operating through a common exporter.

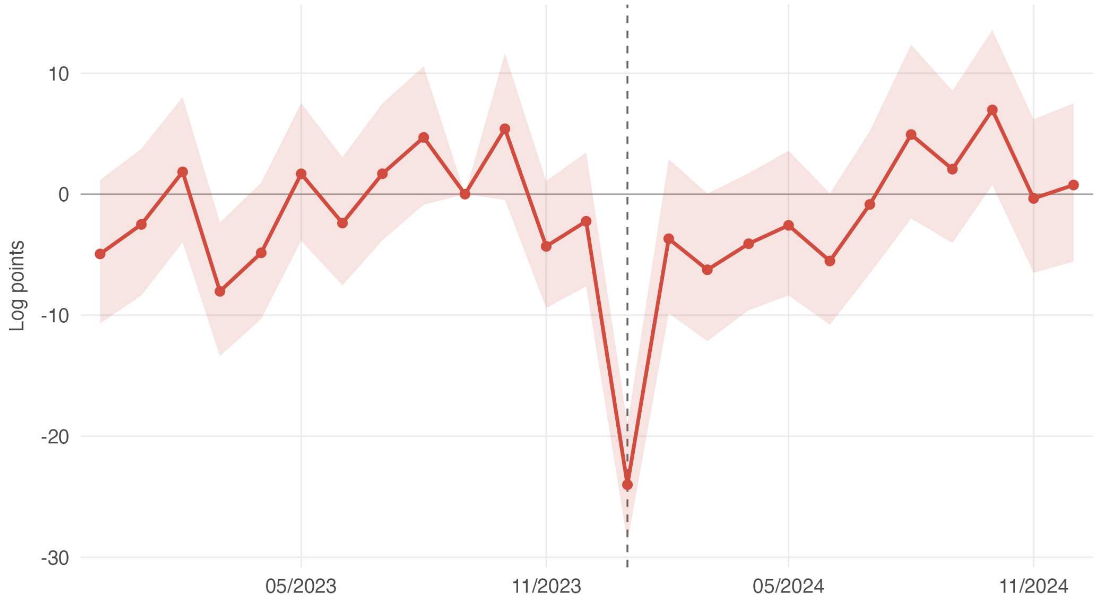

Figure 6: Event-study coefficients $\tilde{\beta}_\tau$, trade value as dependent variable (equation (2)).
Notes: Bands are 95% confidence intervals.

carries Asia–Europe container traffic. When ships reroute around the Cape of Good Hope, that trunk lane no longer passes its doorstep, and a country that had been on the main road finds itself on a spur. Firm-level customs records let us ask whether this abrupt loss of centrality left a footprint in the trade of individual Turkish firms, and whether that footprint faded on the same timescale as the aggregate shock.

Data We use shipment-level customs records for Turkey from Panjiva, which report, for each consignment, the firm, the partner country, the transport mode, the customs value, and the date. We collapse these records to a firm $\times$ partner-country $\times$ week panel of exports, splitting values by transport mode, over a window of twelve weeks on either side of the dropped seizure week. A partner is coded as via-Suez if reaching it from Turkey by sea required crossing the Bab-el-Mandeb strait before the closure, so that the Cape reroute lengthened the voyage; partners reached westward serve as the comparison group. The panel is cleaned for day-level coverage gaps and winsorized for unit-value outliers, and singleton firm–partner pairs are dropped.[3]

Specification For each outcome y we estimate a weekly event study,

$$y_{fjt} = \sum_{\tau \neq -1} \beta_\tau \, \mathbb{I}\{t = \tau\} \times \text{ViaSuez}_j + \alpha_{fj} + \alpha_t + \varepsilon_{fjt}, \tag{3}$$

[3]The seizure of the *Galaxy Leader* on 19 November 2023 straddles two of our weekly bins, so the week containing it is partially treated and is dropped from all specifications. Event time τ is measured in weeks relative to the first attack, and the week of 6 November 2023 is the omitted reference.

where f indexes the firm, j the partner country, and t the week. The firm × partner fixed effects α_{fj} absorb the level of each trading relationship, and the week fixed effects α_t absorb aggregate shocks common to all partners. The coefficients β_τ trace the differential weekly path of via-Suez relationships relative to the comparison group, and standard errors are clustered by partner country.

A transitory response Figure 7a plots the event study for the log value of maritime exports, on the sample of firm–partner pairs with positive maritime shipments. The pre-attack coefficients are jointly insignificant, consistent with parallel trends before the shock. Via-Suez maritime export value falls three weeks after the attack, by 5.1 percent relative to the comparison group (standard error 2.5 percent), and again at four weeks, by 5.4 percent (standard error 2.9 percent). The gap then closes: a partial rebound appears around six weeks, and by the end of the window the differential is statistically and economically indistinguishable from zero (-0.2 percent, standard error 3.5 percent). Because the trough and the recovery roughly offset, a single pooled difference-in-differences over the whole post-period is close to zero, and we read the maritime-value evidence as a drop-and-recovery rather than as a level shift. This is the firm-level counterpart of the short-lived aggregate response documented above: the quantity of seaborne trade dipped and then returned to its pre-crisis path within a quarter.

A persistent shift toward air The response of *how* goods move outlasts the response of how much moves. Figure 7b plots the same event study for the air share of all-mode export value. The air share of via-Suez exports dips slightly in the attack week itself, by 1.3 percentage points (standard error 0.6), and again two weeks later by 1.6 percentage points (standard error 0.7), as consignments in transit are disrupted. It then turns and rises, standing 1.4 percentage points above the comparison group by ten weeks after the attack (standard error 0.5), and the pooled difference-in-differences is a 0.65 percentage-point increase in the via-Suez air share (standard error 0.25, $p = 0.011$). Unlike the quantity response, this modal shift does not close within the window: firms trading with via-Suez partners kept a larger fraction of their exports in the air even as maritime volumes recovered. The pre-attack air-share coefficients are jointly insignificant on the eleven-week window used for the figure, though the test sits close to the conventional threshold, so we read the persistent-rise result as suggestive rather than decisive.[4]

Table 3 collects the pooled difference-in-differences for both outcomes. Note the contrast between the columns: the quantity margin recovers to zero, while the modal margin indicates a small but precisely estimated shift toward air.

[4]The joint pre-trend test has $p = 0.10$ on the ten pre-attack coefficients plotted; adding the earliest week lowers it to 0.09. The point estimates and their significance pattern are unchanged when the sensitive partners Israel, Yemen, Sudan and Russia are excluded, and the pre-trend improves in those samples.

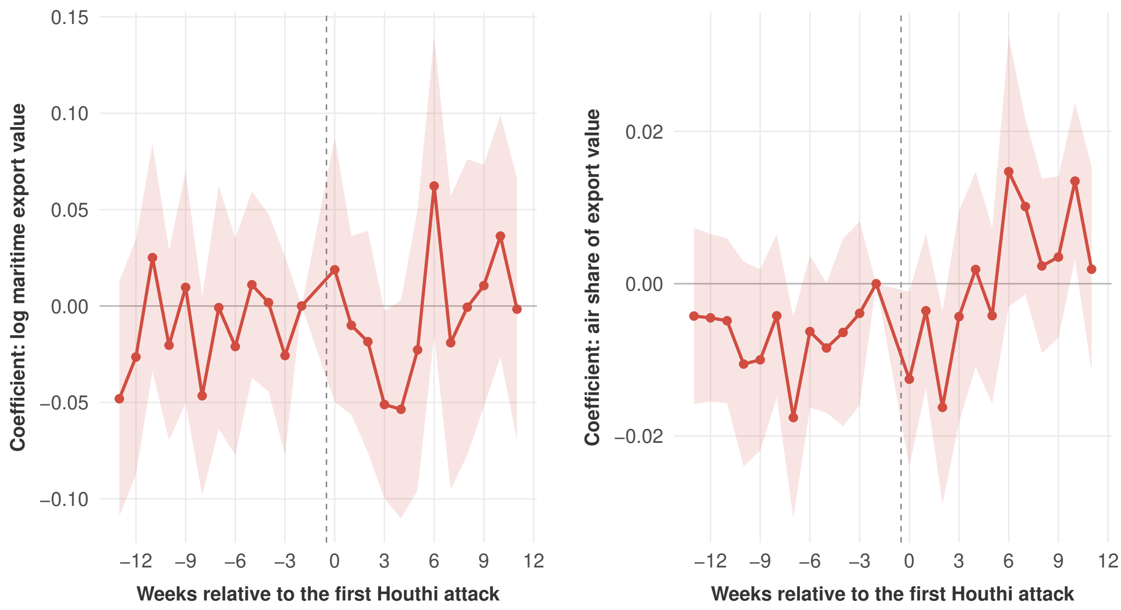

(a) Maritime export value

(b) Air share of export value

Figure 7: Firm-level weekly event-study coefficients β_τ around the Red Sea closure (equation (3)).

Notes: Weekly event study on Panjiva Turkish customs export records, twelve weeks on either side of the dropped seizure week; the first Houthi attack falls in the week of 20 November 2023. Coefficients are the via-Suez × week interactions, normalised to the omitted reference week of 6 November 2023; the partially treated seizure week of 13 November 2023 is dropped. Firm × partner and week fixed effects; 95% confidence bands from standard errors clustered by partner country. Panel (a) plots the log value of maritime exports on firm–partner cells with positive maritime shipments; panel (b) plots the air share of all-mode export value.

Table 3: Firm-level difference-in-differences around the Red Sea closure.

	(1) Log maritime export value	(2) Air share of export value
Via Suez × Post	0.0048	0.0065**
	(0.0129)	(0.0025)
Observations	405,777	1,023,191
Within R^2	0.000	0.000
Firm × partner FE	Yes	Yes
Week FE	Yes	Yes

Notes: Firm × partner × week export panel from Panjiva Turkish customs records, twelve weeks on either side of the dropped seizure week; the seizure week of 13 November 2023 is dropped; post is defined from the week of 20 November 2023. Column (1): log maritime export value on firm–partner cells with positive maritime value. Column (2): air share of all-mode export value. Standard errors clustered by partner country in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

From firms to the model The firm-level record thus tells a two-speed story. The volume of seaborne exports facing the longer route dipped and recovered within a quarter, but the mode by which goods left the country shifted toward air and stayed there. Read against the model in the subsequent section, the transitory quantity response reflects the recovery of the sea route once carriers reorganized around the Cape, while the persistent modal shift is the footprint of a lasting change in relative transport margins. It is precisely these two margins — the switch between transport modes and the switch between routes within a mode, and their response to disruption — that the model endogenizes, and to which we now turn.

4 Model

We model the effect of route disruptions (e.g. the Red Sea/Suez disruption) through two building blocks, which together embed the three mechanisms highlighted in the introduction — congestion on alternative passages, a global shipping capacity constraint, and monopoly toll-setting:

1. **Monopoly rent extraction on a bottleneck route.** A monopolist controls a short route and sets a fee that shifts traffic between the short route and a longer alternative.
2. **Endogenous freight costs via congestion and global capacity.** Route-level and global shipping costs increase with route usage and aggregate shipping activity.

The two building blocks are combined through endogenous route choice and then embedded in a standard multi-sector trade framework with input–output linkages in the spirit of Caliendo and Parro (2015). The embedding extends naturally to the general CES production-network setting of Baqaee and Farhi (2024).

Notation. Countries are indexed by i, j, sectors by s, and routes by $r \in \mathcal{M}_{ij}$, the set of feasible routes connecting i and j. In the *monopolist problem*, there are two routes,

$$r \in \{S, L\},$$

where S denotes the short (bottleneck) route and L the longer alternative. We denote route costs by d_{ij}^S and d_{ij}^L, with $d_{ij}^S < d_{ij}^L$, and the monopolist fee by F.

In the route-only quantitative model, route-level iceberg costs are $d_{ij,sr}$, route congestion Ξ_r, and aggregate shipping activity Ψ. Sector-specific route-choice elasticities are η_s, route-congestion elasticities λ_s, and global-capacity elasticities γ_s.

In the mode-route model, modes are indexed by $m \in \mathcal{N}_{ij}$ and routes within mode by $r \in \mathcal{M}_{ijm}$. Route congestion is Ξ_{mr}, mode-specific capacity tightness Ψ_m, and global capacity tightness Ψ. Monopoly wedges are ϕ_r in the route-only model and ϕ_{mr} in the mode-route model.

4.1 Transit fee

Consider a monopolist that sets a transit fee for facilitating trade between partners i and j through a chokepoint. This short-route-versus-untolled-alternative structure mirrors the second-best pricing problem studied in the road-congestion literature (Verhoef et al., 1996). This pricing reduces the effective (monetary) cost of distance between i and j from d_{ij}^L to $d_{ij}^S < d_{ij}^L$. The elasticity of substitution between the two travel routes is denoted by μ. The cost of travel is assumed to be $d_{ij}^S + F$ for the shorter route—where F is the monopolist's fee—and d_{ij}^L for the longer alternative. The share of trade using the shorter route is therefore given by

$$\pi_{ij}(F) = \frac{(d_{ij}^S + F)^{-\mu}}{(d_{ij}^S + F)^{-\mu} + (d_{ij}^L)^{-\mu}}. \tag{4}$$

The monopolist's profit is defined as

$$T(F) = F \times \pi_{ij}(F) \times T_{ij},$$

where T_{ij} denotes the total volume of trade between i and j, which is assumed to be exogenous to the monopolist's pricing decision.

For analytical tractability, suppose $d_{ij}^S = 0$. This simplification does not alter the qualitative nature of the results.[5] The monopolist's optimal fee is then

$$F^{\star} = d_{ij}^L \times (\mu - 1)^{-\frac{1}{\mu}}.$$

The fee is unbounded when $\mu \leq 1$, which we exclude from consideration. When $\mu \to \infty$, transporters are nearly indifferent between routes, and the fee converges to the cost difference between the two routes. Hence, for high elasticities of substitution, total travel costs across routes are approximately equalized.

If the monopolist internalizes the effect of its fee on the overall trade volume, it would incorporate the aggregate trade elasticity into its optimal pricing decision. The next section takes up this more general case. In this simple framework, the monopolist captures the rents generated by the shorter route, keeping these efficiency gains from diffusing throughout the economy. Hence, when a geopolitical shock raises d_{ij}^S, monopoly profits should fall while aggregate trade costs remain relatively stable.

This simplified setting captures a long-run outcome in which total trade volumes are largely unaffected by disruptions such as the Houthi attacks. In the short run, however, the adjustment is more complex: the sudden closure of a route creates logistical frictions, delays, and a temporary shortage of available ships. Longer routes require a larger fleet

[5] It can be shown that when $d_{ij}^S > 0$, the optimal fee $F^{\star}$ converges to $d_{ij}^L - d_{ij}^S$ as $\mu \to \infty$.

globally—capital that was previously saved (and captured as monopoly rents) must now be reinvested in transport capacity.

4.2 Endogenous freight costs through congestion and capacity constraints

To move from the two-route problem to the quantitative model, we let route costs depend on route usage and aggregate shipping activity.

A central object in what follows is the *freight cost share* $\chi_s \in [0,1]$, defined as the share of delivered price accounted for by the route-level shipping cost at the calibration baseline. The cargo-value iceberg $d_{ij,sr}$ facing importers depends on the underlying freight wedge $f_{ij,sr}$ through

$$d_{ij,sr} \;=\; f_{ij,sr}^{\chi_s}, \tag{5}$$

so that a proportional increase in the freight wedge $\hat{f}_{ij,sr}$ raises the cargo-value iceberg only by a factor $\hat{f}_{ij,sr}^{\chi_s}$. When $\chi_s \to 1$ shipping is the only cost and the iceberg moves one-for-one with the freight wedge; when $\chi_s \to 0$ shipping is a negligible share of value and route shocks do not pass through. Section 5.3 discusses calibration of χ_s; Appendix B discusses an additive alternative.

We define the two underlying state variables directly as flow aggregates.

For each route r, route congestion is driven by route-specific traffic:

$$\Xi_r \;\equiv\; \sum_{i,j,s} \delta_{ij,sr} X_{ij,sr}, \tag{6}$$

where $X_{ij,sr}$ denotes shipments from i to j in sector s using route r, and $\delta_{ij,sr}$ denotes shipment-capacity weight by product, pair, and route (measured in ton-kilometers).

Global capacity tightness is driven by aggregate shipping demand, allowing for route- and shipment-specific capacity weights $\delta_{ij,sr}$:

$$\Psi \;\equiv\; \sum_{i,j,s,r} \delta_{ij,sr}\, X_{ij,sr}. \tag{7}$$

Route-level *freight wedges* are

$$f_{ij,sr} \;=\; \bar{f}_{ij,sr} \cdot \left(\Xi_r/\bar{\Xi}_r\right)^{\lambda_s} \cdot \left(\Psi/\bar{\Psi}\right)^{\gamma_s} \cdot \phi_r, \tag{8}$$

where $\bar{f}_{ij,sr} \geq 1$, $\bar{\Xi}_r$, and $\bar{\Psi}$ are the baseline freight wedge, congestion, and capacity levels; $\lambda_s \geq 0$ and $\gamma_s \geq 0$ govern the sector-specific elasticities of route congestion and global capacity tightness; and ϕ_r is the monopoly transit fee wedge, with $\phi_r > 1$ on the monopolized chokepoint route and $\phi_r = 1$ otherwise. The cargo-value iceberg combines the freight wedge with the freight cost share via (5), so

$$d_{ij,sr} \;=\; \bar{d}_{ij,sr} \cdot \left(\Xi_r/\bar{\Xi}_r\right)^{\chi_s\lambda_s} \cdot \left(\Psi/\bar{\Psi}\right)^{\chi_s\gamma_s} \cdot \phi_r^{\chi_s}, \tag{9}$$

where $\bar{d}_{ij,sr} = \bar{f}_{ij,sr}^{\chi_s}$.

This setup captures the short-run mechanism highlighted in the discussion above: when trade is rerouted away from the bottleneck route, congestion rises on alternatives and global shipping demand raises capacity rents. Congestion in (8) thus accrues on links — it is the alternative passage that becomes more expensive as traffic reroutes — whereas Massoni (2025) models the complementary node margin, congestion at ports and terminals, which we hold fixed.

4.3 Route choice

Importers (or shippers) choose among available routes $r \in \mathcal{M}_{ij}$. We assume i.i.d. Gumbel taste shocks at the shipper level and deterministic utility proportional to $-\log d_{ij,sr}$. The resulting choice probabilities are the standard logit form of McFadden (1974). This yields a CES route aggregator with sector-specific route-substitution elasticity $\eta_s > 0$:

$$d_{ij,s} = \left[\sum_{r\in\mathcal{M}_{ij}} d_{ij,sr}^{-\eta_s}\right]^{-1/\eta_s}, \tag{10}$$

$$\pi_{ij,sr} = \frac{d_{ij,sr}^{-\eta_s}}{\sum_{\rho\in\mathcal{M}_{ij}} d_{ij,s\rho}^{-\eta_s}}. \tag{11}$$

Here $d_{ij,s}$ is the route-composite iceberg cost for shipping sector-s goods from i to j, and $\pi_{ij,sr}$ is the share of $i \to j$ sector-s shipments using route r. This discrete-choice structure over routes is the common backbone of the recent endogenous-routing literature: Fuchs and Wong (2026) and, in contemporaneous work on chokepoint disruptions, Fuchs et al. (2026) employ recursive-routing formulations of the same logic, and Massoni (2025) an analogous choice over port-to-port routes.

The two-route model of Section 4.1 shares this logit route-choice structure, with $\mathcal{M}_{ij} = \{S, L\}$ and a common elasticity $\eta_s = \eta$; there the fee enters additively as an illustration, whereas the quantitative model applies the multiplicative toll wedge of (9), which preserves the standard exact-hat structure of the trade model. The additive-in-freight alternative $d_{ij,sr} = 1 + \chi_s(f_{ij,sr} - 1)$ is presented in Appendix B.

4.4 Embedding in a standard trade model

We now embed the route block into a standard trade model: a multi-country, multi-sector Ricardian framework with input–output linkages in the spirit of Caliendo and Parro (2015), as used in the quantification exercise. The key idea is simple: the route-composite trade cost $d_{ij,s}$ replaces the usual exogenous bilateral trade wedge, but is itself endogenous through route choice, congestion, capacity, and monopoly pricing.

Economic environment. Let (i, s) denote a country-sector node. Production in each node uses labor and intermediate inputs from other nodes according to a standard nested CES production network structure. We do not restate the full trade model here; instead, we modify the bilateral trade wedge.

For each sector s, bilateral sourcing from origin i to destination j faces the endogenous route-composite iceberg cost $d_{ij,s}$ from (10). This wedge enters the standard origin-choice equation exactly as an exogenous bilateral trade cost would.

Implementation *à la* Eaton and Kortum (2002). To be concrete, in a multi-sector Eaton and Kortum (2002) model the equations look as follows:

$$
\begin{aligned}
P_{j,s} &= \left[\sum_{i=1}^{N} T_{i,s}\big(w_i d_{ij,s}\big)^{-\theta_s}\right]^{-1/\theta_s}, \\
\pi_{ij,s} &= T_{i,s}\big(w_i d_{ij,s}\big)^{-\theta_s} P_{j,s}^{\theta_s}, \\
X_{ij,s} &= \pi_{ij,s}\, E_{j,s},
\end{aligned}
$$

where $E_{j,s}$ is sectoral expenditure in destination j, determined by the production network and final demand system. Route-level shipments are then

$$X_{ij,sr} = \pi_{ij,sr}\, X_{ij,s}.$$

These route flows feed back into (6)–(9).

Monopoly route and transit fee setting. Let $\mathcal{R}^M \subseteq \mathcal{M}$ denote the set of monopoly-controlled routes. In the route-only model, these are routes r; in the mode-route extension below, they are mode-route pairs (m, r). For now, write the monopoly-controlled object generically as $q \in \mathcal{R}^M$, with toll wedge $\phi_q \geq 1$.

Each monopolist chooses its own toll wedge taking the competitive equilibrium as given:

$$\phi_q^\star \in \arg\max_{\phi_q \geq 1} \Pi_q(\phi_q; \phi_{-q}) \qquad \text{for each } q \in \mathcal{R}^M. \tag{12}$$

When multiple monopoly-controlled routes are present, two pricing conventions close the model. Under *Nash pricing*, (12) holds at the other authorities' equilibrium wedges $\phi^\star_{-q}$, so that tolls form a Nash equilibrium of the simultaneous fee-setting game. Under *pricing in isolation*, each authority solves (12) with the other wedges held at their reference levels—a non-strategic benchmark that coincides with the first round of best-response dynamics and, when strategic interaction across authorities is weak, closely approximates the Nash outcome. An authority may also control several routes at once (for instance, a strait through which multiple onward passages run); it then sets a single wedge on all of them, and routes tolled by more than one authority carry the product of the corresponding wedges. Monopoly revenue is the freight-share-weighted markup applied to spending routed through the tolled cells. On a route tolled by the set of authorities $A(q)$, the combined wedge is $\Phi_q = \prod_{a\in A(q)} \phi_a$, the total take is the markup at the combined wedge, and it is divided between the authorities in proportion to their log-contribution to it:

$$\Pi_a \;=\; \sum_{q:\, a\in A(q)} \; \sum_{i,j,s} \frac{\ln \phi_a}{\ln \Phi_q} \, \frac{\Phi_q^{\chi_s} - 1}{\Phi_q^{\chi_s}} \, X_{ij,sq}, \tag{13}$$

which for a route with a single authority collapses to $\sum_{i,j,s}(\phi_a^{\chi_s} - 1)/\phi_a^{\chi_s} \cdot X_{ij,sq}$, and further to the textbook expression $(\phi_a - 1)/\phi_a \cdot \sum X_{ij,sq}$ when $\chi_s = 1$; the markup scales down toward zero when shipping is a small share of delivered cost. Charging each authority's markup separately on shared routes would over-extract, since $(\phi^\chi - 1)/\phi^\chi + (\varphi^\chi - 1)/\varphi^\chi > ((\phi\varphi)^\chi - 1)/(\phi\varphi)^\chi$; the log-share division keeps the joint take consistent with the combined wedge and is exact whenever routes are not shared. Here $X_{ij,sq}$ denotes spending routed through monopoly-controlled route q. In the route-only model, $q = r$ and $X_{ij,sq} = X_{ij,sr}$; in the mode-route extension, $q = (m, r)$ and $X_{ij,sq} = X_{ij,smr}$. Monopoly revenue can be rebated to the collector country (or split between several collectors) in equilibrium, as in the quantitative implementation. Alternatively, one can let the monopolist maximize host-country welfare, depending on the institutional application.

Equilibrium. An equilibrium consists of:

1. prices, quantities, and trade shares for the usual trade model;
2. route shares $\pi_{ij,sr}$ and route-level shipments $X_{ij,sr}$;
3. route congestion and capacity terms Ξ_r and Ψ;
4. a vector of monopoly tolls $\{\phi_q\}_{q\in\mathcal{R}^M}$ solving the fee-setting problem (12) under the chosen pricing convention;

such that firm optimization, household demand, market-clearing, the route-choice equations (10)–(11), the endogenous freight-cost equations (6)–(9), and monopoly best responses (12) with revenue given by (13) all hold simultaneously.

Computation in changes. Following the quantitative trade literature in the wake of Dekle et al. (2008), the model can be solved in changes using exact hat algebra. Relative to a standard trade model implementation, the additional fixed-point block is the endogenous transport wedge system:

$$\hat{d}_{ij,sr} = \hat{\bar{d}}_{ij,sr} \cdot \hat{\Xi}_r^{\chi_s \lambda_s} \cdot \hat{\Psi}^{\chi_s \gamma_s} \cdot \hat{\phi}_r^{\chi_s},$$

where $\hat{\bar{d}}_{ij,sr} = \hat{\bar{f}}_{ij,sr}^{\chi_s}$ if baseline freight wedges are shocked exogenously.

Together with the route-share equations (11) and the aggregation (10), where $\hat{\Xi}_r$ and $\hat{\Psi}$ are implied by (6)–(7). When monopoly-controlled routes are present, this transport fixed point is nested inside an outer fixed point over monopoly wedges $\{\hat{\phi}_q\}_{q \in \mathcal{R}^M}$ under Nash pricing; under pricing in isolation, a single round of best responses at reference wedges replaces the outer fixed point. This preserves the standard structure of the model while making freight costs and route rents endogenous to disruptions, rerouting, and bottleneck pricing.

Two implementation details matter for counterfactuals around a baseline that already carries tolls. First, toll wedges are levels while hat algebra works in changes: when the calibrated baseline embeds fee levels ϕ_q^{base}, a counterfactual fee ϕ_q enters route costs as the change $\hat{\phi}_q = \phi_q / \phi_q^{\text{base}}$, while toll revenue (13) is evaluated at the level ϕ_q — so that re-solving the baseline at its own fees reproduces it exactly, with unchanged route costs and unchanged rents. Second, the equilibrium in changes holds nominal trade imbalances fixed at their baseline values (as in Dekle et al., 2008) and uses world value added as the numeraire.

4.5 Mode choice above route choice

In reality, shipping companies often first choose a transportation *mode*—sea, air, rail, or road—and then, conditional on the mode, a specific route. This matters when substitution across modes reflects a different economic margin from substitution across routes within a mode. For instance, switching from sea to air is conceptually distinct from rerouting maritime trade from the Suez Canal to the Cape route.

To capture this distinction, we extend the transport block by introducing a nested choice structure, as in the multimodal framework of Fuchs and Wong (2026). Following Tolva (2026), let $m \in \mathcal{N}_{ij}$ denote a feasible transport mode between i and j, and let $r \in \mathcal{M}_{ijm}$ denote a feasible route within mode m. The lower nest governs route choice within a mode, while the upper nest governs mode choice across the resulting mode-specific transport composites.

Mode-specific route composite. For each sector s, bilateral pair (i, j), and mode m, define the within-mode composite iceberg cost as

$$d_{ij,sm} = \left[\sum_{r \in \mathcal{M}_{ijm}} d_{ij,smr}^{-\eta_s^R} \right]^{-1/\eta_s^R}, \tag{14}$$

$$\pi_{ij,smr} = \frac{d_{ij,smr}^{-\eta_s^R}}{\sum_{\rho \in \mathcal{M}_{ijm}} d_{ij,sm\rho}^{-\eta_s^R}}, \tag{15}$$

where $\eta_s^R > 0$ is the sector-specific elasticity of substitution across routes within a mode, $d_{ij,smr}$ is the route-level iceberg cost, and $\pi_{ij,smr}$ is the share of mode-m shipments allocated to route r.

Mode choice. Given the within-mode composite costs $d_{ij,sm}$, importers choose across available modes according to

$$d_{ij,s} = \left[\sum_{m \in \mathcal{N}_{ij}} d_{ij,sm}^{-\eta_s^M} \right]^{-1/\eta_s^M}, \tag{16}$$

$$\mu_{ij,sm} = \frac{d_{ij,sm}^{-\eta_s^M}}{\sum_{\nu \in \mathcal{N}_{ij}} d_{ij,s\nu}^{-\eta_s^M}}, \tag{17}$$

where $\eta_s^M > 0$ is the sector-specific elasticity of substitution across modes and $\mu_{ij,sm}$ is the share of sector-s shipments from i to j that use mode m. The bilateral transport wedge entering the trade model is therefore still $d_{ij,s}$, but it is now a two-level composite: first across routes within modes, then across modes.

Route-level freight costs. At the lowest level of the transport nest, route-specific iceberg costs may depend on route congestion, mode-specific capacity pressure, global capacity pressure, and route-specific tolls:

$$d_{ij,smr} \;=\; \bar{d}_{ij,smr} \cdot \left(\Xi_{mr}/\bar{\Xi}_{mr}\right)^{\chi_s \lambda_s^R} \cdot \left(\Psi_m/\bar{\Psi}_m\right)^{\chi_s \gamma_s^M} \cdot \left(\Psi/\bar{\Psi}\right)^{\chi_s \gamma_s^G} \cdot \phi_{mr}^{\chi_s}, \tag{18}$$

where $\bar{d}_{ij,smr} \geq 1$ is the baseline route-level cost, Ξ_{mr} is congestion on route r within mode m, Ψ_m is mode-specific capacity tightness, Ψ is global transport-capacity tightness, and ϕ_{mr} is a route-specific toll or rent wedge. For monopoly-controlled mode-route pairs, ϕ_{mr} is determined endogenously under either pricing convention of Section 4.4; otherwise $\phi_{mr} = 1$. The parameters $\lambda_s^R \geq 0$, $\gamma_s^M \geq 0$, and $\gamma_s^G \geq 0$ govern the elasticities of costs with respect to route congestion, mode-level capacity pressure, and global capacity pressure, respectively. The freight cost share $\chi_s \in [0, 1]$ governs the pass-through of the route-mode freight wedge to the cargo-value iceberg, allowing low-shipping-share sectors (e.g.

electronics) to absorb route shocks differently from bulk.

The corresponding flow aggregates are

$$\Xi_{mr} \equiv \sum_{i,j,s} \delta^{C}_{ij,smr}\, X_{ij,smr}, \tag{19}$$

$$\Psi_m \equiv \sum_{i,j,s,r} \delta^{M}_{ij,smr}\, X_{ij,smr},$$

$$\Psi \equiv \sum_{i,j,s,m,r} \delta_{ij,smr}\, X_{ij,smr}, \tag{20}$$

where $X_{ij,smr}$ denotes shipments from i to j in sector s using mode m and route r, while $\delta^{C}_{ij,smr}$, $\delta^{M}_{ij,smr}$, and $\delta_{ij,smr}$ are the route-congestion, mode-specific, and global capacity weights, respectively.

Embedding in the trade model. The transport block continues to enter the standard trade model only through the bilateral trade wedge. The difference is that the bilateral iceberg cost is now obtained from the nested transport structure in (14)–(16). For example, in our environment, bilateral sourcing shares become

$$\pi_{ij,s} = T_{i,s}\big(w_i d_{ij,s}\big)^{-\theta_s} P_{j,s}^{\theta_s},$$

with route- and mode-level shipments given by

$$X_{ij,sm} = \mu_{ij,sm}\, X_{ij,s},$$

$$X_{ij,smr} = \pi_{ij,smr}\, X_{ij,sm}.$$

These flows feed back into (19)–(20), thereby endogenizing transport costs through congestion and capacity effects.

Special cases. The route-only model from Section 4.3 is nested as a special case. If there is only one mode for each bilateral pair, then $|\mathcal{N}_{ij}| = 1$, $\mu_{ij,sm} = 1$, and $d_{ij,s} = d_{ij,sm}$. Likewise, if a mode has only one feasible route, then $|\mathcal{M}_{ijm}| = 1$, $\pi_{ij,smr} = 1$, and $d_{ij,sm} = d_{ij,smr}$. Setting each pair to a single transport mode ($|\mathcal{N}_{ij}| = 1$, so that $\mu_{ij,sm} = 1$) therefore switches off mode choice and returns the model to the route-only benchmark of Section 4.3.

Computation in changes. In exact hat algebra, the additional transport fixed-point block becomes

$$\hat{d}_{ij,smr} = \hat{\bar{d}}_{ij,smr} \cdot \hat{\Xi}_{mr}^{\chi_s \lambda_s^R} \cdot \hat{\Psi}_m^{\chi_s \gamma_s^M} \cdot \hat{\Psi}^{\chi_s \gamma_s^G} \cdot \hat{\phi}_{mr}^{\chi_s},$$

together with the within-mode route-share equations (14)–(15) and the across-mode aggregation (16)–(17). When some mode-route pairs are monopoly controlled, the transport block is nested inside an outer fixed point over the monopoly wedges $\{\hat{\phi}_{mr}\}$. Relative to the route-only specification, the additional objects are the mode-specific composite wedges $d_{ij,sm}$, the mode shares $\mu_{ij,sm}$, the mode-specific capacity terms Ψ_m, and potentially a vector of endogenous monopoly wedges ϕ_{mr}. The rest of the general-equilibrium structure remains unchanged.

4.6 Predictions

For a shock at a chokepoint, the model implies a set of comparative statics that can be confronted with the evidence of Section 3. On impact, trade falls for affected country pairs, the more so the larger the maritime share of their trade and the freight cost share of the goods traded. As shipments reallocate towards alternative passages, congestion on those passages and tighter global shipping capacity raise freight costs — also for country pairs whose routes are not directly disrupted. Bilateral trade between affected pairs is hence partly restored through rerouting, even as freight rates stay elevated. At the disrupted chokepoint itself, throughput and toll revenue decline: the shock erodes the relative advantage of the bottleneck route and with it the rents its operator can extract. These are the on-impact comparative statics the static model delivers; the dynamic recovery of trade over the months after the shock, and any accompanying expansion of global shipping capacity, lie outside the static framework and are not solved for here. The stylized facts of Section 3 — the initial decline and subsequent recovery of trade between affected pairs, the rerouting visible in vessel positions, and the rise in freight rates — are consistent with these mechanics.

5 Calibration

5.1 Data and simulation inputs

Baseline input–output data. The baseline equilibrium is the GTAP 11 database, covering 160 countries and 65 sectors with full input–output linkages. Its country detail is what makes the chokepoint scenarios meaningful: Egypt, Panama, Iran, Oman, and each of the Gulf economies enter as separate regions, so toll revenues accrue to their actual collectors and incidence can be traced country by country. The reference year of the accounts is 2017; trade shares and route exposure are, however, taken from more recent data as described below.

Route shares. Baseline route shares $\pi_{ij,sr}$ are constructed from ship-level AIS trajectories and port-call sequences (Appendix A). For each origin–destination pair we classify observed voyages into four passages—Suez, Panama, Cape of Good Hope, and a direct route—and compute route shares from trip counts. Country pairs without observed AIS voyages are assigned to the direct route. Because the AIS data do not identify cargo at the product level, route shares are common across sectors within a country pair; sectoral variation in exposure comes from mode shares and from the freight cost shares χ_s.

Mode shares. Trade is split across sea, air, and other transport at the country-pair–sector level using predicted transport modes for HS6 product flows, aggregated to the sector classification. Pairs and sectors without mode information—predominantly services—are assigned to the other-transport mode, which is not exposed to maritime chokepoints. The mode-substitution elasticities are set at the sector level following Section 5.2, with a pooled value of 2.44 for sectors—predominantly services—without a sectoral estimate.

Chokepoint exposure. Exposure to the Strait of Hormuz is assigned from port geography. Each country receives a Gulf-side share g_i of its seaborne trade ($g_i = 1$ for Kuwait, Iraq, Qatar, and Bahrain; 0.95 for Iran; 0.80 for Saudi Arabia; 0.75 for the United Arab Emirates; 0 otherwise), and the share of trade between i and j transiting the strait is $e_{ij} = g_i(1 - g_j) + g_j(1 - g_i)$, so that intra-Gulf trade is unexposed. Exposed trade is carried on composite routes (Hormuz–Suez, Hormuz–Cape, and so on) that inherit the AIS route shares scaled by e_{ij}: a Hormuz transit precedes the choice of onward passage, tolls on the strait and on the canals stack multiplicatively, and a closure of an onward passage also closes its Hormuz composite.

Canal fees and the route-choice elasticity. Because toll revenues are not recorded in the input–output accounts, the Suez and Panama wedges are calibrated in a preliminary equilibrium in which each authority sets its revenue-maximizing fee on the baseline data, pricing in isolation (the non-strategic convention of Section 4.4). The route-choice elasticity on the cargo iceberg, μ^d, is then pinned by a revenue moment: we choose μ^d such that the resulting revenue-maximizing Suez fee generates the observed Suez Canal revenue of \$10.25 billion (2023).[6] This yields $\mu^d = 350$, $\phi_{\text{Suez}} = 1.59$, and—as an untargeted validation check—Panama toll revenue of \$3.62 billion against an observed \$3.18 billion (FY2024),[7] with $\phi_{\text{Panama}} = 1.16$. The counterfactuals hold these wedges fixed; the Hormuz fee is set exogenously at the levels described in Section 6.1.

[6] Calendar-year 2023 Suez Canal Authority toll revenue, at its pre-crisis peak; the calibrated fee yields \$10.5 billion, within about 2.5 percent of the target. Suez revenue fell sharply in 2024 as Red Sea traffic collapsed.

[7] FY2024 (October 2023–September 2024) Panama Canal toll revenue was depressed by drought-related transit restrictions, so the unconstrained revenue-maximizing level implied by the model is expected to lie above it.

5.2 Elasticities

Trade elasticites θ_s For the Fréchet elasticities underlying the productivity differences, we use elasticities from Fontagné et al. (2022) directly aggregated at the GTAP sectoral level. For every sector s, we use $\theta_s = \sigma_s - 1$ where σ_s is the long-run elasticity of substitution based on the impact of tariffs in Fontagné et al. (2022).

Transport modality elasticity η_s^M Ko et al. (2025) use a long-run instrumental variable based on air and sea distances to estimate sectoral elasticities of substitution between modes that vary from 0.4 to 11.4. These sectoral values are applied in the simulations, mapped to GTAP sectors; sectors without an estimate — in particular services — use the pooled value of 2.44.[8] Table 4 reports the values.

Trade route congestion externalities λ_s Congestion raises the cost of shipping by raising the time it takes to ship a product through a specific route; in the model this is the route-congestion elasticity λ_s in (9). (The separate mode-capacity elasticity γ_s^M of the mode-route extension is set to zero in the quantification.) Jung et al. (2025) gauge how much a prolonged drought in the Panama Canal raised the time it took ships to travel through the canal. Our approach is similar: we estimate how much congestion episodes driven by high traffic delayed passages, excluding episodes such as the Houthi attacks or the *Ever Given* blockage of the Suez Canal. Let $\Xi_{Suez,t}$ be the total number of ships going through Suez at time t, and let $\tau_{Suez,t}$ be their average measured travel time. We estimate the following regression:

$$\log \tau_{Suez,t} = \alpha_\tau \log \Xi_{Suez,t} + \varepsilon_t.$$

The parameter α_τ gauges how much longer a passage takes when traffic is higher. We estimate it to be close to 0.38. However, α_τ does not directly correspond to our route-congestion elasticity λ_s: while labor costs and ship utilization time can reasonably be assumed proportional to time, other costs such as fuel, loading, or unloading should not be affected by congestion.[9] To convert the estimate α_τ into a congestion externality λ, we hence scale it by the relative importance of time-dependent costs. We estimate this share from the long-run relationship between the logarithm of travel costs and the logarithm of distance, using average freight prices by route over a long period. Let β_τ summarize the impact of distance on costs in the long run. Our estimate of β_τ is 0.35, based on the

[8] This aggregate elasticity of substitution is close to the estimate by Tolva (2026), that is equal to 2.6. That estimate is based on a shock similar to the ones we model: it is based on a difference-in-difference exploiting the closure of Russian airspace in the wake of the war in Ukraine in 2022.

[9] Fuel consumption, for example, increases only moderately with congestion and can even decline, as lower speeds reduce consumption.

Table 4: Transformed IV transport-mode elasticities: $-\beta^{IV}$

Sector	Description	Sectoral Elasticity
TOT	All sectors pooled	2.44
		(0.53)
AFP	Agriculture and food processing	4.87
		(2.18)
BPH	Basic pharmaceutical products	11.36
		(2.47)
CHM	Chemical products	3.13
		(0.60)
EEQ	Electrical equipment	2.88
		(1.11)
ELE	Computer, electronic and optical	1.67
		(0.95)
ENG	Energy	11.08
		(1.17)
FMP	Metal products	4.08
		(0.98)
I_S	Ferrous metals	5.21
		(1.13)
LEA	Leather products	0.78
		(0.72)
MVH	Motor vehicles and parts	0.60
		(0.27)
NFM	Metals n.e.c.	5.81
		(1.71)
NMM	Mineral products n.e.c.	2.68
		(0.44)
OFD	Food products n.e.c.	0.39
		(0.39)
OME	Machinery and equipment n.e.c.	3.17
		(0.89)
OMF	Manufactures n.e.c.	4.34
		(1.28)
OTH	Goods and services n.e.c.	1.60
		(1.61)
OTN	Transport equipment n.e.c.	4.47
		(1.52)
PPP	Paper products, publishing	1.01
		(0.59)
RPP	Rubber and plastic products	2.46
		(0.79)
TEX	Textiles	4.05
		(0.87)
WAP	Wearing apparel	3.21
		(1.13)

Notes: Sectoral rows report the transformed parameter $-\beta^{IV}$, where β^{IV} is the IV coefficient reported in Table 2 of Ko et al. (2025), with the standard errors of the original. The "TOT" (all-sectors-pooled) row reports the pooled estimate of Tolva (2026), used as the fallback for sectors without a sectoral estimate.

correlation between the logarithm of Freightos price data for seven containerized routes, averaged from 2022 to 2025, and the logarithm of distance between the corresponding ports.[10] We then compute $\lambda = \alpha_\tau \times \beta_\tau \approx 0.135$.

5.2.1 Inverse supply elasticity γ_s^G

Parameter γ_s^G determines how much an increase in global use of maritime capacity increases the cost of shipping globally in the medium-run and reflects rigid capacity-constraints when there is an unexpected increase in demand for freight, in the same spirit as Brancaccio et al. (2025). Freight prices and capacity use are simultaneously determined by the demand system

$$\begin{aligned} \log \Xi_t &= -\,\alpha_\Xi^D \,\log P_t^G + \varepsilon_t^D, \\ \log P_t^G &= \quad \gamma_s^G \,\log \Xi_t + \varepsilon_t^S, \end{aligned} \tag{21}$$

where $\alpha_\Xi^D > 0$ is the demand elasticity and γ_s^G is the inverse supply elasticity. P_t^G is the freight price, that we approximate with the containerized freight index. To measure Ξ_t, we use combined data from Census (US imports), Comext (EU imports and EU exports), with modality of transport, complemented by our geographic data, to estimate total maritime imports in ton-kilometers each month.[11] Our estimation uses data from January 2016 to April 2026. Let $\mathbf{y}_t = (\log \Xi_t,\, \log P_t^G)'$ follow the reduced-form VAR

$$\mathbf{y}_t = \mathbf{c} + \sum_{\ell=1}^{L} \mathbf{A}_\ell\, \mathbf{y}_{t-\ell} + \mathbf{u}_t, \qquad \mathbb{E}\big[\mathbf{u}_t \mathbf{u}_t'\big] = \Sigma, \tag{22}$$

with the lag order L chosen by Bayesian information criterion. We recover the structural shocks $\boldsymbol{\varepsilon}_t = (\varepsilon_t^D, \varepsilon_t^S)'$ through $\mathbf{u}_t = \mathbf{B}\,\boldsymbol{\varepsilon}_t$, where $\mathbf{B}$ is the lower-triangular Cholesky factor of Σ under the ordering $(\log \Xi_t, \log P_t^G)$. Lower triangularity imposes the single exclusion restriction $\partial \log \Xi_t / \partial \varepsilon_t^S = 0$: within the month, capacity use is predetermined with respect to the current freight price, so the innovation to $\log \Xi_t$ is a pure demand shock. A demand shock then shifts quantity along the supply schedule (21), and the impact response of price identifies the supply elasticity,

$$\gamma_0 \;\equiv\; \frac{\mathrm{IRF}_P(0)}{\mathrm{IRF}_\Xi(0)} \;=\; \frac{b_{21}}{b_{11}} \;=\; \gamma_s^G. \tag{23}$$

Letting $\boldsymbol{\Phi}_h$ denote the reduced-form moving-average matrices of (22) and $\mathrm{IRF}_k(h)$ the response of variable k to a demand shock at horizon h (from $\boldsymbol{\Theta}_h = \boldsymbol{\Phi}_h \mathbf{B}$), the *cumulative*,

[10] To include round-trip data, we take the average of the round-trip travel for every pair of destinations.

[11] We avoid double counting by making sure that we have mirror flows for an observation, we use as a priority import data. Using imports implies that flows are potentially recorded with a larger lag relative to freight contracts.

Table 5: Dynamic supply elasticity $\gamma_s^G(h)$ from the recursive VAR

Horizon h (months)	$\gamma_s^G(h)$	90% CI	95% CI
Ξ_t: *ton-km index* ($L = 3$, $N = 124$)			
0 (impact)	0.437	$[-0.029,\ 0.920]$	$[-0.105,\ 1.025]$
1	0.930	$[-0.017,\ 1.938]$	$[-0.201,\ 2.125]$
3	0.715	$[-0.652,\ 1.935]$	$[-0.901,\ 2.221]$
6	0.168	$[-1.845,\ 1.736]$	$[-2.135,\ 2.037]$

Notes: P^G is the Drewry World Container Index; Ξ is our own maritime transport-work index based on trade modality data from Census and Comext, both in logs. $\gamma_s^G(h)$ is defined in (24); $\gamma_s^G(0) = \gamma_0$ is the impact elasticity (23). Intervals are percentile bootstrap intervals ($2{,}000$ replications) under the recursive ordering $(\log \Xi_t, \log P_t^G)$.

or dynamic, supply elasticity over h months is

$$\gamma_s^G(h) \ \equiv \ \frac{\sum_{j=0}^{h} \mathrm{IRF}_P(j)}{\sum_{j=0}^{h} \mathrm{IRF}_\Xi(j)}, \qquad \gamma_s^G(0) = \gamma_0. \tag{24}$$

Table 5 reports $\gamma_s^G(h)$, with percentile confidence intervals from $2{,}000$ recursive-residual bootstrap replications. The point estimates are positive and hump-shaped in the horizon, consistent with mild congestion pricing—a one-percent demand-driven rise in capacity use raises freight prices by roughly 0.44 percent on impact and by up to 0.7–0.9 percent cumulatively over one to three months. They are, however, imprecise: every interval contains zero, so we cannot reject a perfectly elastic global supply of shipping ($\gamma_s^G = 0$).

For the simulations we adopt the impact elasticity, $\gamma_s^G = 0.44$, uniformly across sectors; the cumulative responses (0.7–0.9) suggest that this is, if anything, a conservative choice for disruptions on maritime freight in the medium run.

5.3 Freight cost shares

The freight cost share χ_s governs the share of the delivered price accounted for by route-level shipping cost. We assign it at the sector level from the transport-cost evidence in Hummels (2007) and the ad-valorem maritime transport cost estimates underlying the UNCTAD and OECD transport-cost databases: bulk commodities (grain, ore, fertilizer) carry χ in the range 0.20–0.40, basic chemicals and steel 0.05–0.10, manufactures 0.02–0.04, and high-value categories such as electronics and pharmaceuticals 0.01, with a trade-weighted mean of about 0.09. This cross-sector heterogeneity is the key driver of differential exposure to maritime route shocks in our counterfactuals.

6 Counterfactual simulations

We solve the model in changes for 160 countries and 65 sectors, using the data, calibrated fees, and elasticities of Section 5. The calibrated Suez and Panama wedges are held fixed throughout; welfare is measured as the change in real income relative to the baseline equilibrium.

6.1 Scenarios

Red Sea closure. The Suez route becomes unavailable for all country pairs, and trade reallocates across the remaining routes—primarily the Cape of Good Hope. The scenario captures a persistent version of the 2023–24 Houthi disruption, in which effectively no cargo transits the Red Sea.

Strait of Hormuz transit fee. A toll is levied on all seaborne trade transiting the Strait of Hormuz, with revenue accruing in equal parts to Iran and Oman, the two littoral states. We consider three intensities of the fee ϕ_{Hormuz} on the freight bill of exposed cargo: 1.05, the scale of actual canal tolls such as Suez; 1.35, the scale of the transit fees reportedly under discussion in Tehran since March 2026 (around $1.5–2 million per crossing); and 1.50 as an upper intensity, still below crisis-level war-risk premia.

6.2 Results

6.2.1 A Red Sea closure

Closing the Suez route reduces global real income only mildly: cargo reroutes around the Cape of Good Hope and through Panama, whose flows rise by 18 and 20 percent, at a freight penalty that translates into small delivered-price increases once weighted by the freight cost shares χ_s. The distributional effects are concentrated. Egypt is the largest loser at -3.02 percent of real income, almost entirely through the collapse of canal toll revenue rather than through its own trade costs: of the -3.02 percent, -2.76 percentage points are the direct loss of Suez rents and only -0.27 the residual general-equilibrium and trade-cost channel (Table 10). This incidence is conditional on the Canal Authority pricing at its revenue-maximizing fee in the baseline: precisely because the Authority already captures much of the shorter route's surplus, its disruption falls on the collector rather than on shippers. Transshipment- and Suez-dependent economies follow at an order of magnitude smaller losses (Qatar -0.20 percent, Malta -0.05 percent), while the large European and Asian trading economies lose less than 0.02 percent each—rerouting flexibility, the modest freight share in delivered prices, and the density of alternative routes absorb the shock. Figure 8 maps the incidence, Figure 9 the associated rerouting, and Figure 10 the most and least affected economies; Appendix C reports the corresponding tables.

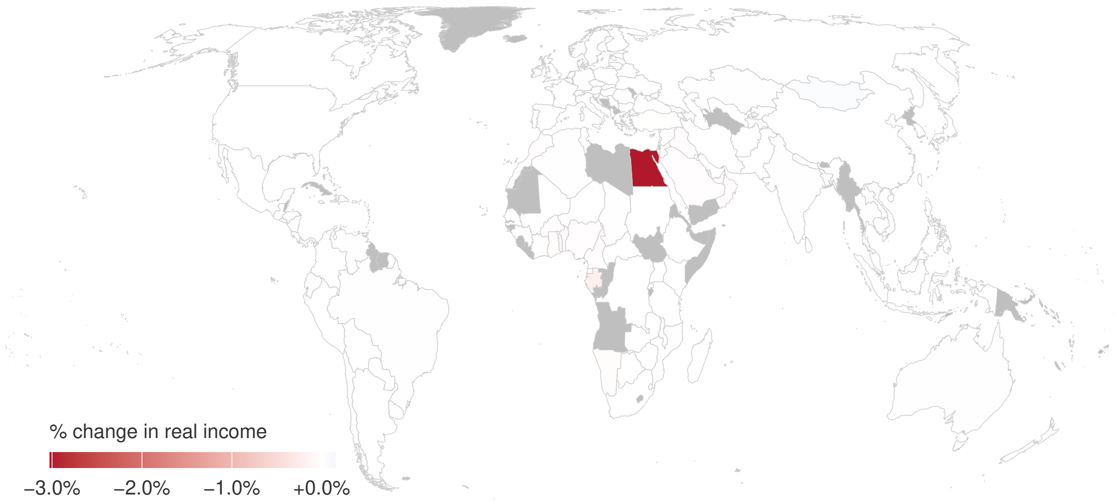

Figure 8: Welfare change under a Red Sea closure. Percent change in real income relative to the baseline equilibrium; grey indicates no data.

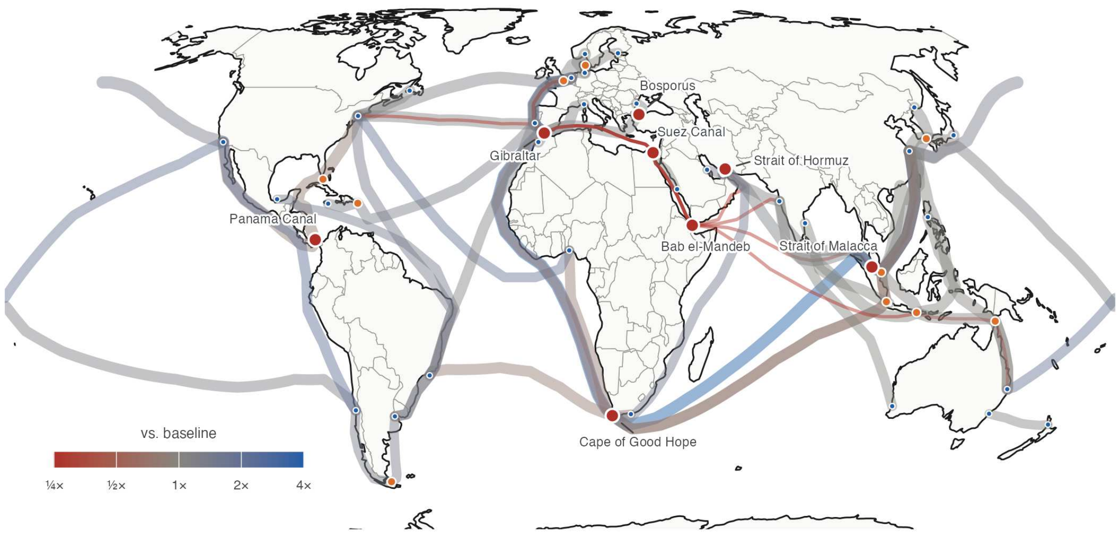

Figure 9: Seaborne trade flows under a Red Sea closure. Line width is proportional to the log of post-shock shipment value; colours indicate the change relative to the baseline (red: decline, blue: increase).

6.2.2 A transit fee in the Strait of Hormuz

The Strait of Hormuz differs from Suez and Panama in that, for the Gulf economies, there is no alternative route: a toll cannot be evaded by rerouting, only attenuated by switching transport modes or redirecting trade towards partners whose shipments do not transit the

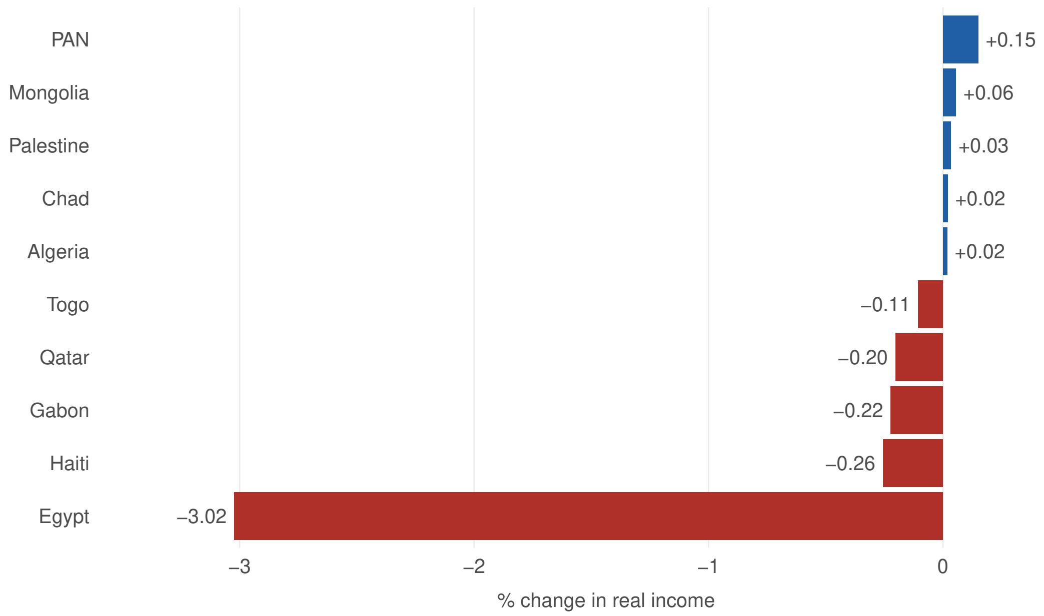

Figure 10: Red Sea closure: the five largest welfare losses and gains (percent change in real income).

strait. Incidence therefore falls almost entirely on the Gulf itself, in proportion to each country's exposure. At the intermediate fee ($\phi_{\text{Hormuz}} = 1.35$), Qatar loses 3.9 percent of real income, Kuwait 2.7 percent, and Iraq 2.5 percent; Bahrain, the United Arab Emirates, and Saudi Arabia—the latter two partly served by Red Sea and Fujairah ports outside the strait—lose 0.6, 0.2, and 0.1 percent, respectively. Figure 11 maps the incidence, Figure 12 shows the associated reallocation of seaborne trade across routes, and Figure 13 reports the five largest welfare losses and gains.

Three features of the incidence stand out. First, effects on third countries are negligible: Germany, France, and the Netherlands each lose about 0.01 percent of real income, and even the most exposed advanced economies—Korea, through LNG imports, and India—lose less than 0.06 percent. The strait is critical for the Gulf, not, in welfare terms, for its customers, who substitute towards other suppliers. Second, the beneficiaries of that substitution are rival hydrocarbon exporters outside the Gulf: Equatorial Guinea and Brunei gain almost 1 percent of real income at the intermediate fee. Third, the two collectors fare very differently despite the equal revenue split. Oman, whose own trade largely bypasses the strait, receives a nearly pure transfer worth $+12.2$ percent of real income; Iran, with an exposure of 0.95, levies a large part of the fee on its own trade and nets $+1.5$ percent. Toll revenue rises from $4.8 billion at the low fee to $14.6 billion at the intermediate and $17.0 billion at the high fee, while Hormuz-transiting trade falls by 19, 60, and 67 percent, respectively—at the fee level under discussion, the toll would already be close to the revenue-maximizing region of its Laffer curve.

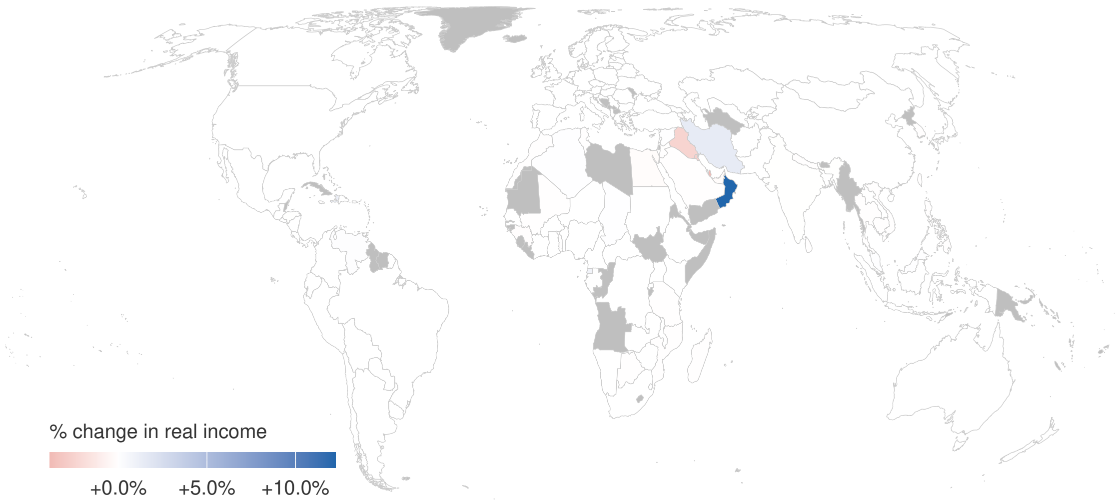

Figure 11: Welfare change under a Strait of Hormuz transit fee ($\phi_{\text{Hormuz}} = 1.35$, the scale of the fees under discussion since March 2026). Percent change in real income relative to the baseline equilibrium.

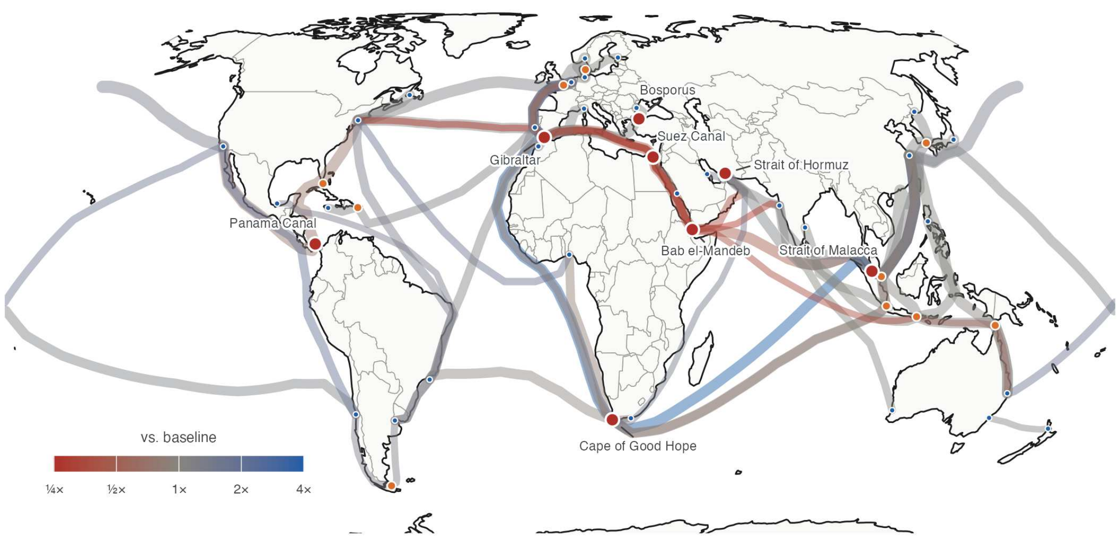

Figure 12: Seaborne trade flows under a Strait of Hormuz transit fee ($\phi_{\text{Hormuz}} = 1.35$). Line width is proportional to the log of post-shock shipment value; colours indicate the change relative to the baseline (red: decline, blue: increase).

General-equilibrium spillovers across chokepoints. The corridor-level reallocation in Figure 12 reveals that a toll on one chokepoint propagates to the others even though total flows on the remaining routes barely move. Gulf–Mediterranean trade transits Hormuz *and then* Suez, so the Hormuz fee stacks on top of the Suez toll: Gulf-linked traffic

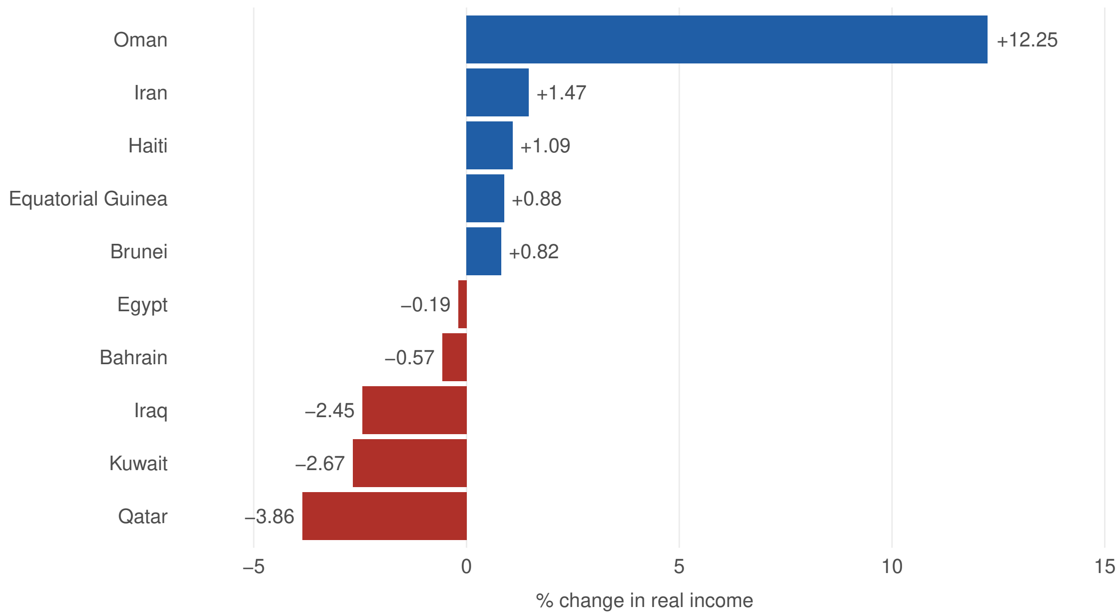

Figure 13: Strait of Hormuz transit fee ($\phi_{\text{Hormuz}} = 1.35$): the five largest welfare losses and gains (percent change in real income).

through the Suez Canal (the Hormuz–Suez composite) falls by 54 percent ($33.9 billion, about 8 percent of total Suez traffic) at the intermediate fee, and with it the Suez Canal Authority's toll income—an Iranian transit fee taxes the canal at one remove. Egypt's net welfare effect is a loss of about 0.2 percent, as the drop in Gulf-linked Suez toll traffic reduces its canal income. Conversely, the traffic gains around the Cape of Good Hope do not reflect Gulf cargo evading the strait—no such route exists—but the geographic footprint of supplier substitution: flows linked to West African hydrocarbon exporters rise as importers redirect demand towards them, while Gulf-linked flows on the same corridor decline. Both spillovers are only visible because chokepoints are linked through the route structure rather than modelled in isolation. Figure 14 summarizes the route-level reallocation across all scenarios.

Caveats. Four modelling choices qualify these numbers. First, exposure to the strait is assigned from port geography, not estimated (Section 5.1); the Gulf-side shares of Saudi Arabia, the United Arab Emirates, and Iran in particular are judgement calls. Second, the bypass margin of these partially exposed countries—rerouting through Red Sea, Fujairah, or Jask ports—substitutes at the route-choice elasticity μ, which likely overstates short-run bypass capacity and thus understates their losses. Third, the Hormuz fee is exogenous throughout: unlike the calibrated Suez and Panama wedges, it is not chosen optimally by the collectors, and canal fees do not re-optimize in response to any scenario. The welfare numbers should therefore be read as the incidence of given fee levels, not as the outcome of a fee-setting game. Fourth, the route shares and the route-choice elasticity are estimated

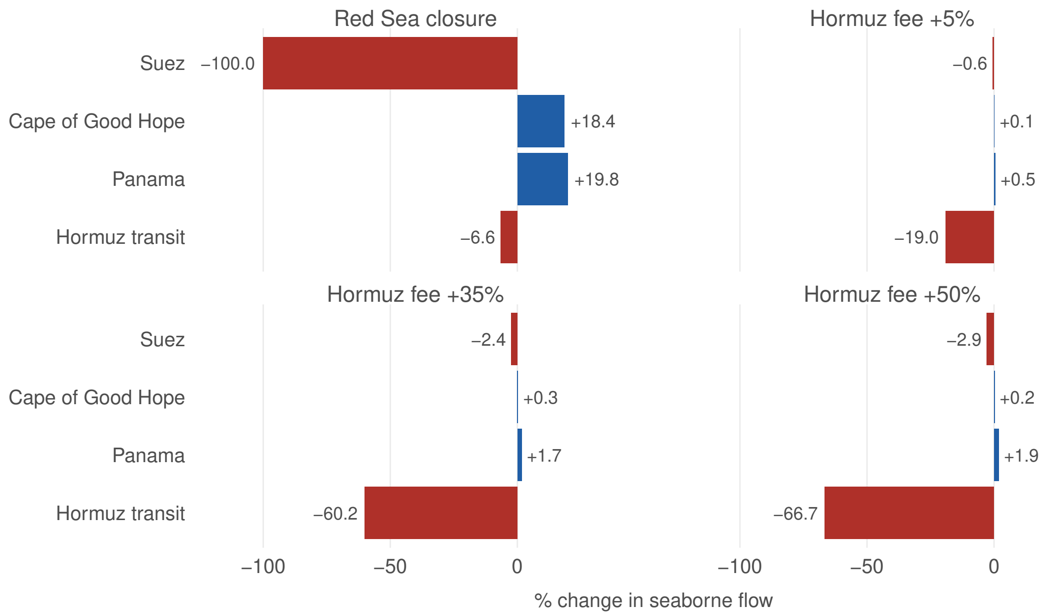

Figure 14: Seaborne trade by route: percent change in flow relative to the baseline, by scenario. Composite Hormuz routes are attributed to their onward passage; the last bar reports total trade transiting the Strait of Hormuz.

from container vessels, whereas Hormuz trade is predominantly crude, refined products, and LNG; because tanker and LNG routing is less elastic and more contract-bound, the concentration of losses on the Gulf is, if anything, a lower bound. Relatedly, the model's high long-run elasticities of substitution allow importers to switch toward unaffected suppliers at relatively low cost, shifting much of the burden onto Gulf producers and likely understating short-run importer losses.

6.3 Who bears the closure: rents concentrate the loss on the collector

The incidence of a Red Sea closure in Section 6.2.1 rests on one assumption: that the Canal Authority prices at its revenue-maximizing fee in the baseline, so that the disruption destroys a rent rather than raising shippers' costs. To isolate this channel we re-solve the closure in a counterfactual world in which no chokepoint carries a monopoly toll: the Suez and Panama wedges are pinned at $\phi = 1$, and the baseline is the raw assembled equilibrium with no embedded transit rents. The closure shock itself is unchanged—the Suez route becomes unavailable and cargo reroutes around the Cape of Good Hope and through Panama—but there is now no toll revenue to collapse, so the loss is the pure real-resource cost of the longer haul.

The comparison overturns the incidence. Egypt's loss falls from -3.02 to -0.18 percent of real income once its canal-toll revenue is stripped out of the baseline, leaving only the second-order effect of the closure on its own trade. The losses that remain, and the new ones that appear, fall on the economies whose trade actually lengthens: Gabon deepens

Table 6: Red Sea closure with and without chokepoint rents

	Red Sea closure	
	with fees	without fees
Egypt	−3.02	−0.18
Qatar	−0.20	−1.09
Malta	−0.05	−0.22
Singapore	−0.01	−0.07
United Arab Emirates	−0.04	−0.21
Haiti	−0.26	−1.09
Gabon	−0.22	−0.64
Togo	−0.11	−0.33

Notes: Percent change in real income relative to the baseline equilibrium. The "with fees" column reproduces the Red Sea closure of Section 6.2.1, in which the Suez and Panama wedges are held at their calibrated revenue-maximizing values. The "no fees" column re-solves the same closure with both wedges pinned at $\phi = 1$ and no monopoly rents embedded in the baseline, so that the reported loss is the real-resource cost of rerouting alone.

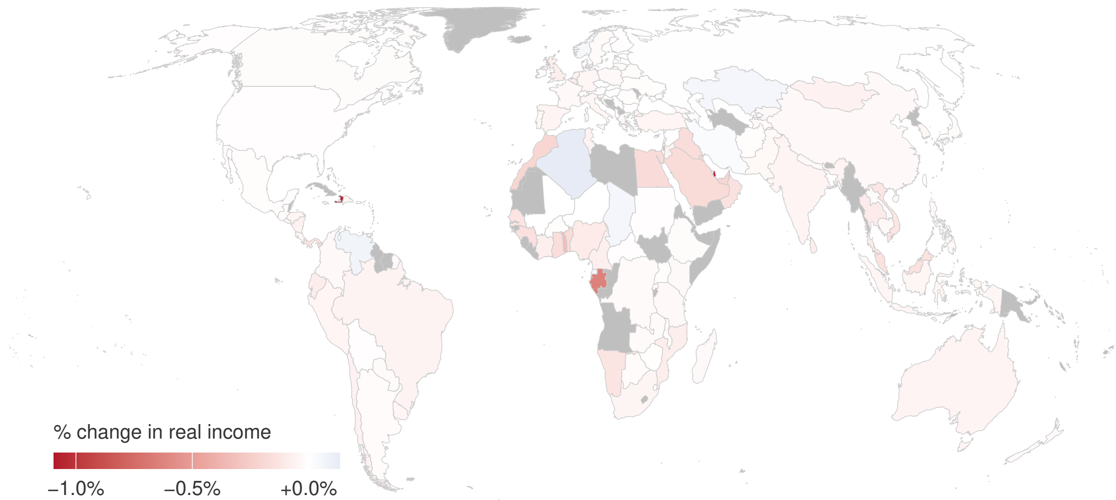

Figure 15: Welfare change under a Red Sea closure without chokepoint rents. Percent change in real income relative to a baseline with the Suez and Panama wedges pinned at $\phi = 1$; grey indicates no data.

from -0.22 to -0.64 percent and Malta from -0.05 to -0.22 percent. Qatar shifts the same way, from -0.20 to -1.09 percent, as does the United Arab Emirates, from -0.04 to -0.21 percent. Table 6 sets the closure with and without fees side by side; Figure 15 maps the no-fee incidence and Figure 16 shows the most and least affected economies. In the world without rents the burden of the closure spreads outward from the collector to the countries whose cargo must travel farther.

The contrast has a clean interpretation, but not a reassuring one for the collector: pricing

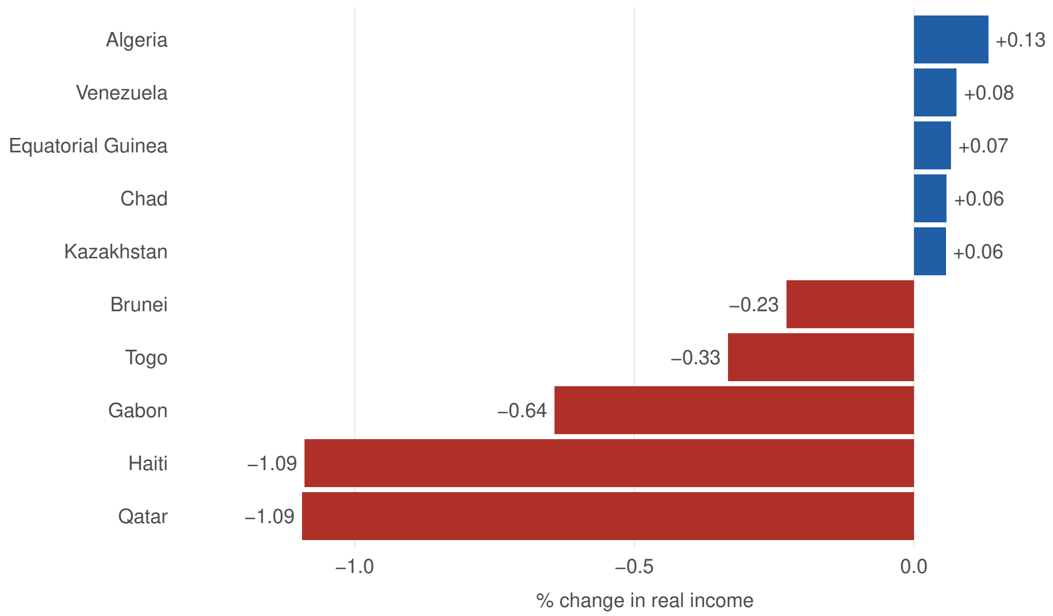

Figure 16: Red Sea closure without chokepoint rents: the five largest welfare losses and gains (percent change in real income).

the chokepoint concentrates the disruption loss on the party that prices it. Because the calibrated Suez fee is disciplined by the surplus of the shorter route, levying that surplus as a rent hands the Canal Authority a claim worth $10.5 billion a year that runs to essentially zero the moment the route closes—so Egypt's loss reaches -3.02 percent of real income, of which -2.76 percentage points are the collapse of that rent rather than any shock to its own trade (Table 10). For the shippers the fee does act as a cushion: in the no-fee world the same closure lands harder on the economies whose cargo must travel farther—Qatar deepens from -0.20 to -1.09 percent, the United Arab Emirates from -0.04 to -0.21, Gabon from -0.22 to -0.64—losses that in the fees world are muted because part of the route surplus is skimmed upstream as a toll. The rent therefore does not insure the collector; it makes the collector the residual risk-bearer, absorbing as a concentrated payout what would otherwise be spread thinly across many trading partners. Stripping the rent away does not make the closure cheaper—global real income falls by a similar amount either way—but moves its incidence off the single collector and back onto the parties whose trade the closure actually disrupts.

7 Conclusion

This paper gauges the local and global consequences of disruptions at maritime chokepoints. We document the trade effects of the Red Sea crisis — an initial decline in seaborne trade between affected country pairs, rerouting around the Cape of Good Hope, and a recovery within months — and build a quantitative trade model in which freight costs respond to congestion, global shipping capacity, and monopoly toll-setting at priced chokepoints. Calibrating the route-choice elasticity to observed canal revenues disciplines the model with the margin that matters most for chokepoint economics: how easily cargo substitutes away when transit becomes more expensive.

The counterfactuals deliver one central lesson: chokepoints are heterogeneous, and their welfare relevance cannot be read off transit volumes. A permanent closure of the Red Sea passage — one of the largest disruptions to maritime trade in decades — lowers real income of the large trading economies by less than 0.02 percent, because cargo reroutes around the Cape of Good Hope at a freight penalty that is small relative to the value of the goods carried. The bulk of the loss falls instead on Egypt, whose canal-toll revenue collapses: where an infrastructure owner pricing to maximize revenue already extracts much of the surplus of a shorter route, a disruption is, in welfare terms, a shock to the owner rather than to the shippers — in our calibration, a loss of 3.0 percent of Egyptian real income. In the case of the Strait of Hormuz, a chokepoint without maritime substitutes, the logic reverses. A transit fee is borne almost entirely by the Gulf economies themselves, third countries substitute towards other suppliers, and the two would-be collectors fare very differently depending on how much of their own trade transits the strait. The firm-level evidence for Turkey reinforces this reading: the response to the closure was short-lived, but the accompanying shift toward air transport persisted, pointing to a lasting change in relative transport margins of the kind the model makes endogenous.

These results come with caveats. Exposure to the Strait of Hormuz is assigned from port geography rather than estimated, canal authorities do not re-optimize their fees across scenarios, and the calibrated substitution margins reflect a crisis that lasted months rather than years. Within these limits, the framework offers a transparent way to gauge which chokepoints matter, for whom, and through which channel — rerouting, freight costs, or rents. Treating all chokepoints as identical overstates the cost of disruptions to priced canals with feasible detours and understates the exposure of regions that sit behind passages without alternatives.

References

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

Asturias, Jose, "Endogenous Transportation Costs," *European Economic Review*, 2020, *123*, 103366.

Attinasi, Maria Grazia, Lukas Boeckelmann, Lorenz Emter, Massimo Ferrari Minesso, Rinalds Gerinovics, Vanessa Gunnella, Baptiste Meunier, and Roberta Serafini, "Sailing through Storms: The Fallout of Red Sea Disruptions for Global Trade and Inflation," VoxEU Column, CEPR April 2024.

Baqaee, David Rezza and Emmanuel Farhi, "Networks, Barriers, and Trade," *Econometrica*, 2024, *92* (2), 505–541.

Bernhofen, Daniel M., Zouheir El-Sahli, and Richard Kneller, "Estimating the Effects of the Container Revolution on World Trade," *Journal of International Economics*, 2016, *98*, 36–50.

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

_ , _ , and _ , "Rigidities in Transportation and Supply Chain Disruptions," *AEA Papers and Proceedings*, 2025, *115*, 543–550.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *Review of Economic Studies*, 2015, *82* (1), 1–44.

Chowdhry, Sonali, Inga Heiland, and Hendrik Mahlkow, "Quantitative Trade with Ships," Kiel Working Paper 2314, Kiel Institute for the World Economy 2026.

Conte, Maddalena, Pierre Cotterlaz, and Thierry Mayer, "The CEPII Gravity Database," CEPII Working Paper 2022-05, CEPII 2022.

Cristoforoni, Enrico, Marco Errico, Federico Rodari, and Edoardo Tolva, "Oligopolies in Trade and Transportation: Implications for the Gains from Trade," 2025. Working paper, revise and resubmit at the Journal of International Economics.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Global Rebalancing with Gravity: Measuring the Burden of Adjustment," *IMF Staff Papers*, 2008, *55* (3), 511–540.

Djankov, Simeon, Caroline Freund, and Cong S. Pham, "Trading on Time," *The Review of Economics and Statistics*, 2010, *92* (1), 166–173.

Dunn, Jason and Fernando Leibovici, "Navigating the Waves of Global Shipping: Drivers and Aggregate Implications," Working Paper 2023-002, Federal Reserve Bank of St. Louis 2025. Revised February 2025.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

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

Feyrer, James, "Trade and income—exploiting time series in geography," *American Economic Journal: Applied Economics*, 2019, *11* (4), 1–35.

_ , "Distance, trade, and income—The 1967 to 1975 closing of the Suez Canal as a natural experiment," *Journal of Development Economics*, 2021, *153*, 102708.

Fontagné, Lionel, Houssein Guimbard, and Gianluca Orefice, "Tariff-based product-level trade elasticities," *Journal of International Economics*, 2022, *137*, 103593.

Fouquin, Michel and Jules Hugot, "Two Centuries of Bilateral Trade and Gravity Data: 1827–2014," CEPII Working Paper 2016-14, CEPII 2016.

Fuchs, Simon and Woan Foong Wong, "Multimodal Transport Networks," NBER Working Paper 35065, National Bureau of Economic Research 2026.

_ **, Fernando Leibovici, and Woan Foong Wong**, "Geopolitical Risk and Foreign Dependence in Global Shipping," 2026. Presented at the NBER Summer Institute (International Trade & Investment), July 2026.

Gaffuri, Julien and Eurostat, "SeaRoute: Compute shortest maritime routes between ports," `https://github.com/eurostat/searoute` 2022. European Union Public Licence v1.2.

Ganapati, Sharat and Woan Foong Wong, "How Far Goods Travel: Global Transport and Supply Chains from 1965–2020," *Journal of Economic Perspectives*, 2023, *37* (3), 3–30.

_ , _ **, and Oren Ziv**, "Entrepôt: Hubs, Scale, and Trade Costs," *American Economic Journal: Macroeconomics*, 2024, *16* (4), 239–278.

Halili, Gent, "searoute-py: A Python package for generating the shortest sea route between two points on Earth," `https://github.com/genthalili/searoute-py` 2026. Version 1.6.0.

Head, Keith and Thierry Mayer, "Gravity Equations: Workhorse, Toolkit, and Cookbook," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics*, Vol. 4, Elsevier, 2014, chapter 3, pp. 131–195.

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

Hugot, Jules and Camilo Umana Dajud, "Trade Costs and the Suez and Panama Canals," CEPII Working Paper 2016-29, CEPII 2016.

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

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

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

Jacks, David S and Krishna Pendakur, "Global trade and the maritime transport revolution," *The Review of Economics and Statistics*, 2010, *92* (4), 745–755.

Jacks, David S., Christopher M. Meissner, and Nikolaus Wolf, "Suez," NBER Working Paper 33250, National Bureau of Economic Research 2024.

Jung, Paul, Kijin Kim, Ah-Hyun Jo, and Joseph Seong-Hyun Cho, "Panama Canal Drought and Supply Chain Disruptions in Asia–United States Trade: Evidence from Micro-Level Trade Shipments and Vessel Trajectory Data," ADB Economics Working Paper Series 822, Asian Development Bank 2025.

Kalouptsidi, Myrto, "Time to Build and Fluctuations in Bulk Shipping," *American Economic Review*, 2014, *104* (2), 564–608.

Massoni, Hubert, "Climate Trade Costs: Extreme Weather, Transportation, and Supply Chains," 2025. Job Market Paper, University of Bologna, October 2025.

Maurer, Stephan and Ferdinand Rauch, "Economic geography aspects of the Panama Canal," *Oxford Economic Papers*, 2023, *75* (1), 142–162.

Mayer, Thierry and Soledad Zignago, "Notes on CEPII's Distances Measures: The GeoDist Database," CEPII Working Paper 2011-25, CEPII 2011.

McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in Paul Zarembka, ed., *Frontiers in Econometrics*, New York: Academic Press, 1974, pp. 105–142.

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

Pascali, Luigi, "The wind of change: Maritime technology, trade, and economic development," *American Economic Review*, 2017, *107* (9), 2821–2854.

Pilgrim, Graham, Emmanuelle Guidetti, and Annabelle Mourougane, "An Ocean of Data: The Potential of Data on Vessel Traffic," OECD Statistics Working Papers 2024/02, OECD Publishing March 2024.

Tolva, Edoardo, "One Way or Another: Modes of Transport and International Trade," 2026. Working paper, University of Warwick / ISEG, Universidade de Lisboa.

Verhoef, Erik, Peter Nijkamp, and Piet Rietveld, "Second-Best Congestion Pricing: The Case of an Untolled Alternative," *Journal of Urban Economics*, 1996, *40* (3), 279–302.

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

won Ko, Jeong, Sergey Nigai, and Alec Truax, "Sectoral Transport Mode Elasticities and International Trade," 2025. Working paper, revise and resubmit at the Journal of International Economics.

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

A Route Share Construction

Raw AIS position data contain — among other things — a persistent ship identifier (the IMO number), the vessel's position in latitude and longitude, its heading and speed, and a timestamp. To obtain port calls, we construct port boundaries via density-based clustering of AIS positions around a known port centroid, conceptually similar to Pilgrim et al. (2024). A vessel dwelling inside a port boundary for a minimum time span is considered to have called that port; arrival and departure times are pinned down by the first and last observation of the vessel inside the boundary. Having obtained vessel-specific port calls, a ship's voyage is its sequence of port calls over time.

We use the sequence of port calls to build bilateral trips: for each vessel, we walk through its history of port calls and extract all valid trips from origin country o to destination country d, including any intermediate ports. A trip from o to d is valid if (i) a port call in o is eventually followed by a port call in d, and (ii) no premature port call is made in o or d. This truncates trips where vessels serve domestic ports following an international trip. Trips starting after the event are tagged as post, all others as pre; trips straddling the event are hence assigned wholly to the pre period.

We use the bilateral trips to generate route shares. Using the *searoute-py* package (Halili, 2026), a Python implementation of the SeaRoute routing engine (Gaffuri and Eurostat, 2022), we compute the shortest route between o and d. If the shortest route leads through at least one maritime bottleneck such as the Suez Canal, we also compute the restricted counterfactual route. If a shortest route leads through multiple bottlenecks, we compute counterfactual routes for all possible restriction combinations. Our assumption is that vessel operators will always opt for the shortest possible route in order to minimise time at sea, fuel expenditure, etc.[12]

B Additive freight-share decomposition

The multiplicative iceberg $d_{ij,sr} = f_{ij,sr}^{\chi_s}$ in (5) embeds the freight cost share into the trade-cost wedge through a Cobb–Douglas-type aggregator of cargo and freight. An economically tighter primitive treats cargo and freight as Leontief complements at the level of moving one shipment: total delivered cost equals goods cost plus freight cost, in fixed proportions. The two are equivalent at the calibration baseline ($f = 1 \Rightarrow d = 1$ under both) and to first order at small wedges, but diverge under large freight wedges.

[12]Note that this assumption may be violated if certain routes are riskier, for instance due to weather conditions.

Iceberg decomposition. Under the additive (Leontief) decomposition, the cargo-value iceberg satisfies

$$d_{ij,sr} \;=\; 1 \;+\; \chi_s\bigl(f_{ij,sr} - 1\bigr), \tag{3'}$$

where $f_{ij,sr}$ is the multiplicative freight wedge from (8). At $\chi_s = 1$ the iceberg coincides with the freight wedge; at $\chi_s \to 0$ the iceberg collapses to 1.

Monopoly revenue. The corresponding monopoly revenue equation is

$$\Pi_q(\phi_q; \phi_{-q}) \;=\; \sum_{i,j,s} \frac{\chi_s(\phi_q - 1)}{1 + \chi_s(\phi_q - 1)}\, X_{ij,sq}(\phi_q, \phi_{-q}). \tag{14'}$$

The toy-model fee structure from Section 4.1 is recovered exactly when $\chi_s = 1$ on a single bottleneck route.[13]

Hat algebra. The exact-hat transport-block equation becomes nonlinear in the freight wedge:

$$\hat{d}_{ij,sr} \;=\; \frac{1 + \chi_s\bigl(f_{ij,sr}\,\hat{\hat{f}}_{ij,sr} \cdot \hat{\Xi}_r^{\lambda_s} \cdot \hat{\Psi}^{\gamma_s} \cdot \hat{\phi}_r - 1\bigr)}{d_{ij,sr}}. \tag{15'}$$

The transport-block fixed point under the additive decomposition iterates on the freight wedge $f_{ij,sr}$ in multiplicative form, with the cargo-value iceberg recomposed via (3') at each outer iteration.

Mode-route extension. Under the additive decomposition the mode-route level analogue of (18) is

$$d_{ij,smr} \;=\; 1 \;+\; \chi_s\bigl(f_{ij,smr} - 1\bigr), \tag{24'}$$

with $f_{ij,smr}$ defined analogously to (8) at the mode-route level.

Equivalence at small wedges. A first-order expansion of (5) around $f = 1$ gives $d \approx 1 + \chi(f-1) + O((f-1)^2)$, identical to (3') up to second order. The two decompositions therefore deliver the same elasticities at the calibration baseline and the same comparative statics for marginal shocks. Divergence is second-order in $f - 1$ and proportional to $\chi(1-\chi)$.

Numerical comparison. Table 7 reports the cargo-value iceberg under both decompositions, evaluated at the model's calibrated baseline Suez markup $\phi_r = 1.59$ and a range of sector-typical freight cost shares.

[13] On routes tolled by several authorities, the markup in (14') is evaluated at the combined wedge $\Phi_q = \prod_{a \in A(q)} \phi_a$ and divided across authorities in proportion to $\phi_a - 1$ — the additive analogue of the log-share division in (13).

Table 7: Cargo-value iceberg d under multiplicative vs. additive decomposition.

Representative sector	χ	$d = \phi_r^\chi$	$d = 1 + \chi(\phi_r - 1)$
Electronics, pharmaceuticals	0.01	1.0046	1.0059
General manufactures	0.03	1.0139	1.0176
Steel, basic chemicals	0.08	1.0376	1.0468
Bulk (grain, ore, fertilizer)	0.30	1.1483	1.1756
Shipping-only limit	1.00	1.5855	1.5855

Notes: The cargo-value iceberg is evaluated at the model's calibrated baseline Suez markup $\phi_r = 1.59$ and a range of sector-typical freight cost shares χ. Columns report the multiplicative decomposition $d = \phi_r^\chi$ and the additive decomposition $d = 1 + \chi(\phi_r - 1)$.

The two formulations agree to sub-percentage-point precision for containerised sectors and diverge by under two percentage points for bulk. The relative ranking of sectors by exposure is identical under both. In the hat-algebra implementation the additive update is applied as $\hat{d} = 1 + \chi_s(\hat{f} - 1)$, which coincides with (3') exactly when the baseline freight wedge is one; for shocked baselines the multiplicative decomposition is the exact specification and is the one used in all reported results.

C Additional counterfactual results

Table 8 reports welfare changes for the most and least affected economies across the scenarios of Section 6; Table 9 reports the associated route-level reallocation of seaborne trade. Figures 17 and 18 show the welfare and route-flow maps for the remaining Hormuz fee intensities.

Two further numbers are worth recording. Under the Red Sea closure, flows through the Panama Canal rise by 20 percent, slightly more than around the Cape—part of the Asia–Europe trade re-optimizes westwards. Under the Hormuz fee, the throughput response is strongly convex in the fee: the low fee shaves 19 percent off Hormuz-transiting trade, while the intermediate fee cuts it by 60 percent. Revenue accordingly flattens between the intermediate and high fee ($14.6 vs. $17.0 billion, for a fee rising from 1.35 to 1.50).

Table 8: Welfare effects of chokepoint scenarios

	Red Sea	Strait of Hormuz transit fee		
	closure	$\phi = 1.05$	$\phi = 1.35$	$\phi = 1.50$
Gulf economies				
Qatar	−0.20	−0.44	−3.86	−5.42
Kuwait	−0.04	−0.26	−2.67	−3.84
Iraq	−0.02	−0.25	−2.45	−3.52
Bahrain	−0.02	−0.07	−0.57	−0.80
United Arab Emirates	−0.04	−0.06	−0.19	−0.20
Saudi Arabia	−0.02	−0.05	−0.06	−0.04
Toll collectors				
Iran	+0.01	+0.48	+1.47	+1.71
Oman	−0.02	+4.01	+12.25	+14.24
Egypt	−3.02	−0.06	−0.19	−0.23
Rival hydrocarbon exporters				
Equatorial Guinea	+0.00	+0.10	+0.88	+1.38
Brunei	−0.02	+0.11	+0.82	+1.14
Large trading economies				
Germany	−0.01	−0.00	−0.01	−0.01
France	−0.00	−0.00	−0.01	−0.01
Netherlands	−0.01	−0.00	−0.01	−0.01
United States	+0.00	−0.00	−0.00	−0.00
China	−0.00	−0.00	−0.01	−0.02
Japan	−0.00	−0.01	−0.02	−0.02
Korea	−0.00	−0.01	−0.06	−0.08
India	−0.01	−0.01	−0.03	−0.03

Notes: Percent change in real income relative to the baseline equilibrium. The Hormuz fee $\phi_{\text{Hormuz}} \in \{1.05, 1.35, 1.50\}$ applies to the freight bill of cargo transiting the strait, with revenue split equally between Iran and Oman.

Table 9: Seaborne trade by route

	Red Sea	Strait of Hormuz transit fee		
Route	closure	$\phi = 1.05$	$\phi = 1.35$	$\phi = 1.50$
Suez	−100.0	−0.6	−2.4	−2.9
Cape of Good Hope	+18.4	+0.1	+0.3	+0.2
Panama	+19.8	+0.5	+1.7	+1.9
Hormuz transit	−6.6	−19.0	−60.2	−66.7

Notes: Percent change in flow relative to the baseline. Composite Hormuz routes are attributed to their onward passage. The last row reports total trade transiting the Strait of Hormuz.

Table 10: Welfare decomposition for the toll collectors

	Red Sea		$\phi = 1.05$		$\phi = 1.35$		$\phi = 1.50$	
	ΔW	$\Delta\tau$	ΔW	$\Delta\tau$	ΔW	$\Delta\tau$	ΔW	$\Delta\tau$
Canal toll collectors								
Egypt	-3.02	-2.76	-0.06	-0.04	-0.19	-0.15	-0.23	-0.18
Panama	$+0.15$	$+0.19$	$+0.01$	$+0.02$	$+0.02$	$+0.04$	$+0.01$	$+0.04$
Iran	$+0.01$	$+0.00$	$+0.48$	$+0.43$	$+1.47$	$+1.31$	$+1.71$	$+1.52$
Oman	-0.02	$+0.00$	$+4.01$	$+3.73$	$+12.25$	$+11.35$	$+14.24$	$+13.17$
World	-0.02	-0.01	$+0.00$	$+0.01$	-0.00	$+0.02$	-0.01	$+0.02$

Notes: Decomposition of the welfare change (percent of real income, relative to the baseline equilibrium) into the direct toll-rent channel and the residual general-equilibrium and trade-cost channel. ΔW is the total welfare change; $\Delta\tau$ is the direct rent component (the change in the country's canal-toll revenue as a share of income); the residual is $\Delta W - \Delta\tau$.

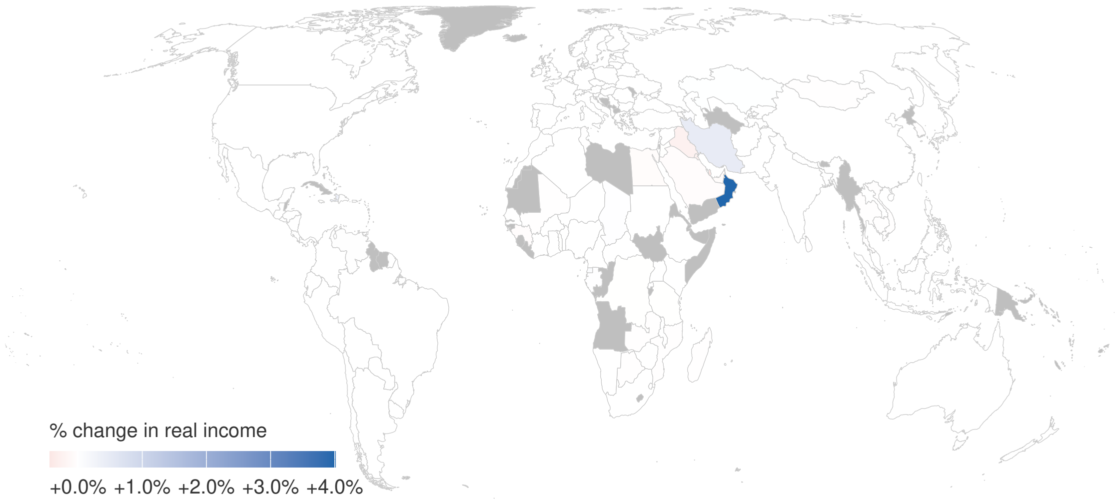

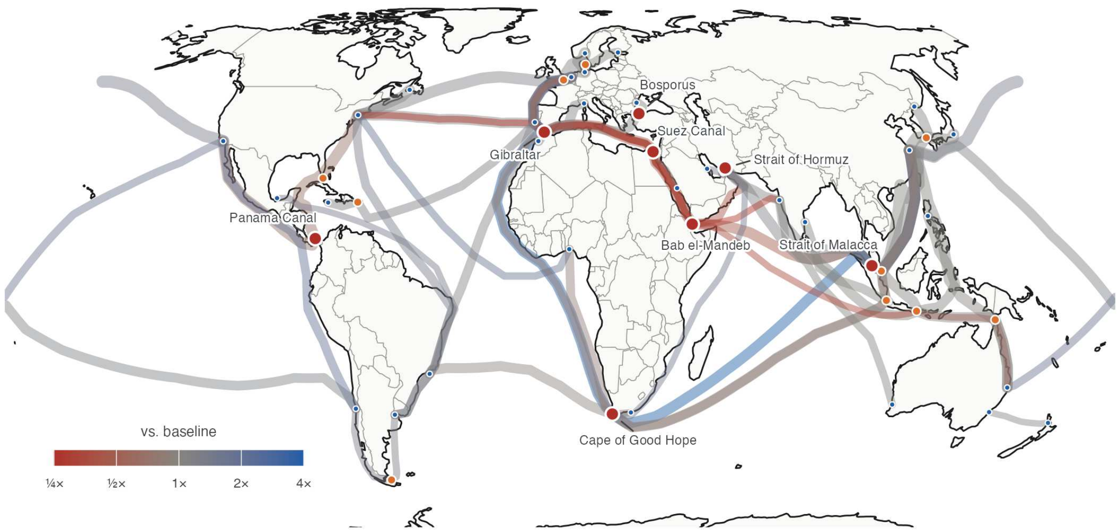

Figure 17: Strait of Hormuz transit fee, $\phi_{\text{Hormuz}} = 1.05$: welfare change (top) and seaborne trade flows (bottom).

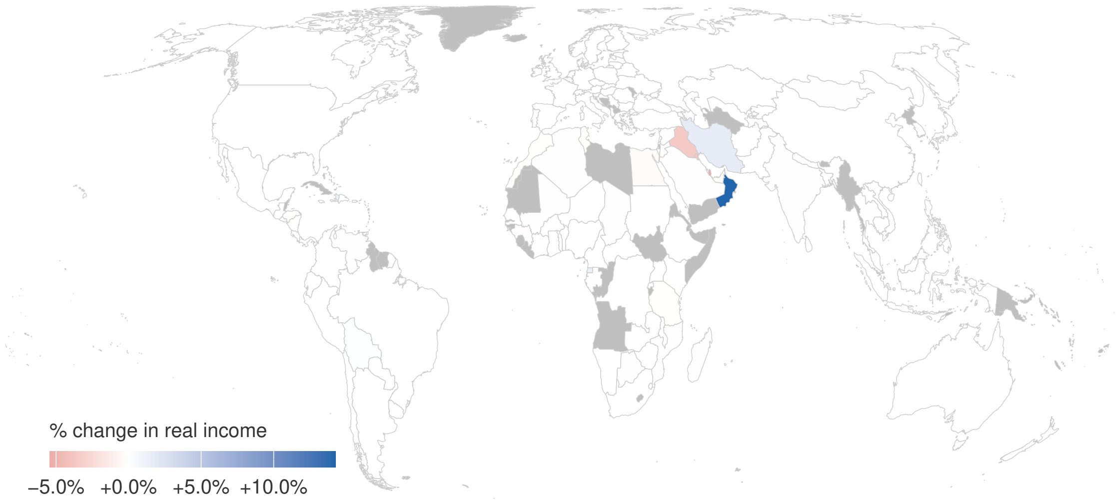

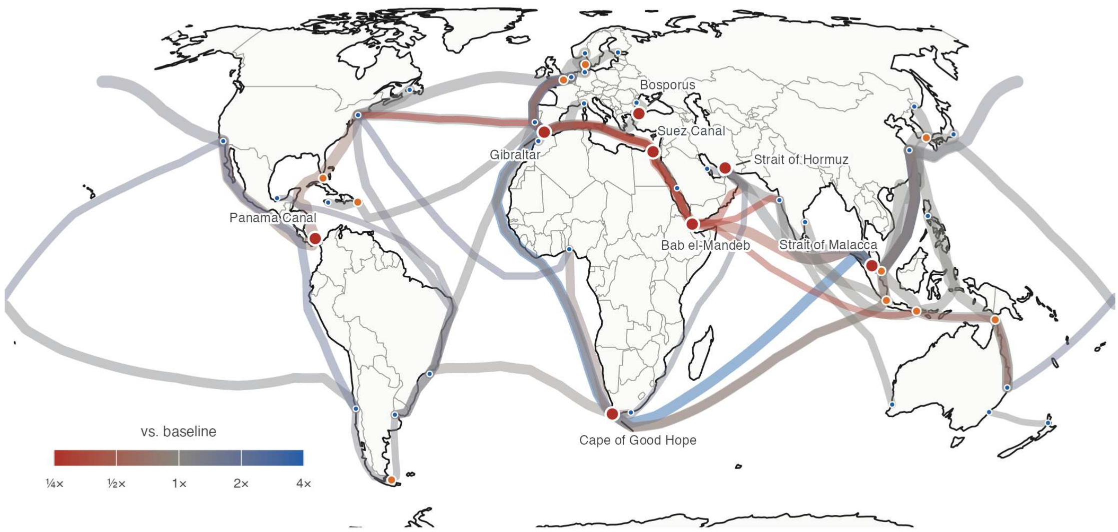

Figure 18: Strait of Hormuz transit fee, $\phi_{\text{Hormuz}} = 1.50$: welfare change (top) and seaborne trade flows (bottom).