Trading Under Fire: Houthi Attacks and the Disruption of Global Shipping

Nathalie El-Bazzal*

August 30, 2026

Abstract

This paper examines how the 2023 Houthi attacks in the Red Sea affected global trade and maritime shipping. Exploiting the Red Sea crisis as an exogenous shock to a major maritime chokepoint, I combine AIS-Kpler vessel-level data covering more than 1.2 million voyages between 2019 and 2025 with international trade data from UN Comtrade. Using a difference-in-differences approach that compares Asia–Europe routes structurally exposed to the Red Sea/Suez corridor with other maritime routes, I find that trade volumes on exposed routes decline by approximately 5.8% after the shock. I further show that vessels become about 10 percentage points more likely to reroute or adopt substantially longer maritime trajectories. At the micro level, the usual cargo-consolidation response to longer delays is weaker on exposed routes and becomes significantly more constrained following the shock. Finally, the effects vary across products, with medium-value goods experiencing the largest trade contractions. Overall, the results highlight how disruptions at critical maritime chokepoints propagate through shipping networks and global value chains, reshaping both trade volumes and shipping behavior.

Keywords: Maritime trade; Geopolitical shocks; Shipping networks; Trade disruptions
JEL classification: F14; F15; F51; L92

1 Introduction

This paper studies how geopolitical disruptions affecting strategic maritime chokepoints reshape global trade volumes and shipping behavior. In October 2023, the landscape of global maritime security shifted abruptly. Following the outbreak of the Israel-Hamas conflict, the Ansar Allah movement (commonly known as the Houthis) initiated a campaign of targeted attacks on commercial vessels transiting the Red Sea. Starting with missile launches on October 19, 2023,

*Corresponding author: Nathalie El-Bazzal, Paris School of Economics and University of Paris 8
Email address: nathalie.elbazzal@psemail.eu

and escalating to the hijacking of the *Galaxy Leader* on November 19, this localized disruption in the Red Sea severely affected the Suez Canal corridor, which connects Asia and Europe and accounts for approximately 12% of global trade and 30% of container traffic.

Unlike traditional maritime threats such as piracy, this episode reflects a form of state-linked asymmetric warfare. The Houthis, supported by Iran and embedded within a broader geopolitical alliance often referred to as the "Axis of Resistance," deploy anti-ship ballistic missiles and unmanned aerial vehicles to target commercial shipping. This technological asymmetry allows a non-state actor to disrupt one of the most critical arteries of global trade without possessing conventional naval power. By exploiting the narrow geography of the Bab el-Mandeb Strait, the disruption effectively imposes a sudden and localized constraint on the Europe-Asia corridor, forcing vessels to adapt their routing decisions in real time.

This episode constitutes a large and plausibly exogenous shock to maritime transport conditions. It provides a unique opportunity to study how geopolitical risk propagates through the transport layer of the global economy. More broadly, it highlights a shift in the nature of maritime vulnerability: rather than reflecting stable geographical constraints, trade routes can become abruptly unstable due to geopolitical events, thereby affecting the organization of global value chains.

A growing literature shows that geopolitical tensions affect trade, global value chains, and firms' sourcing decisions (Handley and Limão, 2017; Alfaro and Chor, 2023). At the same time, a large body of work emphasizes the role of transport frictions-including distance, infrastructure, and time-in shaping trade volumes (Anderson and van Wincoop, 2003; Limao and Venables, 2001; Feyrer, 2021; Hummels and Schaur, 2013). More recent contributions highlight the importance of transport networks and their endogenous response to shocks. Allen and Arkolakis (2022) show that changes in transport infrastructure can generate large welfare effects through network adjustments, while Heiland et al. (2025) document how local disruptions in shipping networks can propagate globally through reorganization of trade routes.

Despite these advances, most existing studies rely on aggregate trade data or customs records and therefore cannot directly observe how transport networks adjust to sudden disruptions. As a result, the micro-level mechanisms through which geopolitical shocks affect shipping behavior and propagate to trade outcomes remain largely unexplored.

Understanding these mechanisms is crucial. When a major maritime corridor becomes risky, shipping firms face a set of operational decisions: reroute vessels through longer but safer paths, delay shipments, or adjust shipment sizes to cope with increased uncertainty. These adjustments affect transit times and the reliability of delivery, which in turn influence trade volumes. In particular, the response to disruptions may involve non-linear adjustments, including rerouting, congestion, and breakdowns in standard logistical relationships.

This paper addresses the following research question: how does a localized geopolitical disruption affecting a strategic maritime chokepoint propagate through shipping networks and translate into trade outcomes?

I test three main hypotheses. First, the disruption increases transit times and delays on exposed routes, leading to a contraction in trade volumes. Second, this effect is heterogeneous across goods, with differential responses across product value categories. Existing evidence suggests that transport costs interact with product value density, with higher-value goods generally better able to absorb per-unit shipping costs (Hummels and Skiba, 2004). I therefore allow for potentially non-linear responses across value segments. Third, the shock disrupts standard operational adjustment margins in shipping, weakening the relationship between transit time and shipment size at the vessel level. This mechanism is consistent with models of inventory management and lumpy trade, where firms adjust shipment sizes in response to time frictions (Alessandria et al., 2010).

To address these questions, I construct a novel dataset based on high-frequency AIS data covering approximately 2.7 million raw observations and more than 1.2 million reconstructed voyages between 2019 and 2025. AIS data provide a granular, vessel-centric view of maritime transport, allowing me to track individual ships continuously across routes, corridors, and maritime zones. Unlike traditional trade datasets, AIS data allow me to directly observe the micro-foundations of trade adjustments in real time.

Empirically, I exploit variation in pre-shock exposure to the Red Sea corridor within a difference-in-differences framework. Asia–Europe and Europe–Asia routes, which are structurally dependent on the Suez Canal, are classified as treated, while other maritime routes serve as the control group. This identification strategy allows me to estimate the differential impact of the disruption on both shipping behavior and trade volumes.

I find that transit times increase sharply on exposed routes following the shock, consistent with increased rerouting and the use of longer alternative maritime trajectories. Trade volumes on these routes decline by approximately 5.8%. The effect emerges progressively over time, consistent with an initial phase of adjustment followed by persistent disruption. At the micro level, vessels typically respond to longer transit times by increasing shipment sizes, consistent with standard logistical adjustments. However, this relationship breaks down on exposed routes, indicating that geopolitical risk disrupts the ability of shipping firms to optimize capacity. Finally, I document substantial heterogeneity across goods, with medium-value products exhibiting the largest contractions, while high-value goods appear more resilient and low-value commodities display more moderate but significant declines.

This paper contributes to three strands of the literature. First, it contributes to the literature on transport frictions and trade, which emphasizes the role of time, geography, and infrastructure in shaping trade volumes (Anderson and van Wincoop, 2003; Limao and Venables, 2001; Feyrer, 2021; Hummels and Schaur, 2013; Allen and Arkolakis, 2022). Second, it contributes to the literature on geopolitics and economic outcomes, which shows that conflict and political risk have substantial economic effects (Collier, 1999; Abadie and Gardeazabal, 2003; Besley and Mueller, 2012), and more recently on trade fragmentation (Handley and Limão, 2017; Caldara and Iacoviello, 2022; Alfaro and Chor, 2023). Third, it contributes to

the literature on transport networks and endogenous adjustment, highlighting how shipping systems reorganize in response to shocks (Brancaccio et al., 2020; Heiland et al., 2025; Schøyen and Bråthen, 2011). By combining vessel-level data with a causal identification strategy, this paper provides a direct empirical link between geopolitical shocks, shipping behavior, and global trade outcomes.

The remainder of the paper is organized as follows. Section 2 presents the data and empirical strategy. Section 3 reports the main results. Section 4 presents robustness checks. Section 5 discusses the findings and their economic implications. Section 6 concludes.

2 Data and Empirical Strategy

2.1 Data

2.1.1 Maritime shipping data (AIS/Kpler)

My primary dataset is derived from AIS-based vessel tracking data enriched with proprietary information from Kpler, covering global dry bulk shipping volumes over the period 2019-2025. AIS (Automatic Identification System) data consist of high-frequency signals transmitted by vessels, reporting information on ship identity, geographic position, speed, and port calls. These signals are collected via satellite and terrestrial receivers and processed to reconstruct vessel trajectories at the global level. Kpler complements these data using proprietary algorithms and additional sources, including port-level records, shipping documents, and commercial intelligence, allowing the identification of cargo types, estimated cargo volumes, and detailed origin–destination shipping flows. [1]

The raw dataset contains 2,761,737 port-level observations. After restricting the sample to observations with non-missing cargo information, the estimation sample includes 2,459,615 observations.

I then reconstruct shipment-level voyages by linking consecutive port calls for each vessel. This procedure allows me to identify origin-destination voyages, compute transit times, and measure transported quantities, yielding more than 1.2 million reconstructed voyages over the period 2019-2025.

These micro-level observations are subsequently aggregated into a monthly panel at the route-product level, resulting in 58,428 route-product-month observations.

To incorporate product-level characteristics, AIS/Kpler product categories are matched to the Harmonized System at the 6-digit level (HS6). The final matched sample used for heterogeneity analysis contains 35,537 observations.

Trade volumes are proxied by total cargo volumes transported along each route-product-

[1] Both AIS data and Kpler datasets are proprietary and not publicly available.

month cell. For each route r, product p, and month t, I define:

$$TradeTons_{rpt} = \sum_{v \in (r,p,t)} Cargo_v, \tag{1}$$

where v indexes voyages departing from the origin port in month t. This measure captures shipments at the time of departure, avoiding mechanical reductions in measured trade that could arise from transit delays alone. The baseline trade variable is therefore built directly from AIS-based shipping quantities rather than from customs records.

I also construct the logarithm of trade volumes, $\ln(TradeTons_{rpt})$, and its first difference, $\Delta \ln(TradeTons_{rpt})$, which is the main dependent variable in the trade regressions.

I further construct measures of shipping disruptions by comparing realized transit times to pre-shock baselines. Travel time is defined as:

$$TravelDays_v = End_v - Start_v. \tag{2}$$

Abnormal delays are defined as deviations from average travel times observed in 2022 for each route and calendar month:

$$AbnormalDelay_{rt} = TravelDays_{rt} - Baseline_{r,m}, \tag{3}$$

where $Baseline_{r,m}$ is the average travel time for route r in month m during 2022. This construction accounts for route-specific seasonality and captures deviations from pre-shock transport conditions.

The analysis focuses on dry bulk commodities such as coal, iron ore, and grains, which are predominantly transported by sea because of their low value-to-weight ratio and large shipment scale. These markets are characterized by significant shipping frictions and limited short-run adjustment capacity, implying that maritime disruptions can generate substantial real effects on trade volumes (Kalouptsidi, 2014).

2.1.2 Trade composition and unit values (UN Comtrade)

To capture product heterogeneity, I match AIS/Kpler product categories to the Harmonized System at the 6-digit level (HS6). This mapping allows me to combine shipping data with external information on product-level values from the UN Comtrade database.

Using UN Comtrade data for the period 2019-2022, I compute product-level unit values as the ratio of trade value in U.S. dollars to physical quantities in kilograms, aggregated at the HS6 level across all reporting countries. Formally, for each product p:

$$UV_p = \frac{\text{Total trade value}_p}{\text{Total traded weight}_p}. \tag{4}$$

These unit values provide a proxy for product value density and are merged with the AIS-based dataset using HS6 product codes. I then construct value terciles-low, medium, and high value-to study heterogeneity in the response of trade volumes to transport disruptions.

Importantly, I do not use UN Comtrade to measure bilateral trade volumes in the regressions. Trade quantities are measured exclusively from AIS/Kpler data. UN Comtrade is used only to characterize the economic value of the goods transported in the AIS-based dataset.

2.2 Empirical strategy

The empirical strategy exploits the Red Sea disruption starting in late 2023 as an exogenous shock to maritime transport conditions. This shock, driven by geopolitical events affecting commercial navigation in the Red Sea corridor, led to a sharp increase in transit times and rerouting behavior, particularly for routes connecting Asia and Europe.

2.2.1 Treatment and control groups

The identification strategy relies on a Difference-in-Differences framework comparing routes that are structurally exposed to the Red Sea corridor with routes that are not.

The treatment group consists of routes connecting Asia and Europe. These routes are structurally dependent on the Suez Canal and the Red Sea corridor and are therefore directly affected by the disruption. In the absence of the Suez route, vessels operating on these corridors must reroute around the Cape of Good Hope or adopt longer alternative maritime routes, leading to substantial increases in transit times and transport costs.

The control group includes all other maritime routes, such as Asia-Americas, Europe-Americas, or intra-regional routes, that do not rely on the Red Sea passage. These routes provide a natural counterfactual, as they are not directly exposed to the disruption and therefore capture broader trends in global trade and shipping conditions.

Formally, I define a treatment indicator $Treated_r$ equal to one for Asia-Europe routes and zero otherwise. I define a post-shock indicator $Post_t$ equal to one from November 2023 onward, corresponding to the first full month of the disruption:

$$Treated_r = 1[\text{route } r \text{ connects Asia and Europe}], \tag{5}$$

$$Post_t = 1[t \geq \text{November 2023}]. \tag{6}$$

This design ensures that treatment status is determined by pre-existing route characteristics rather than by post-shock routing decisions.

2.2.2 Baseline trade specification

I estimate the impact of the disruption on trade volumes using the following specification:[2]

$$\Delta \ln(TradeTons_{rpt}) = \alpha + \beta(Treated_r \times Post_t) + \gamma_{rp} + \delta_t + \epsilon_{rpt}. \tag{7}$$

The dependent variable is the monthly change in log trade volumes at the route-product level. Estimating the model in first differences removes persistent differences in trade levels across route-product pairs. Following the inference considerations emphasized by Bertrand et al. (2004), standard errors are clustered at the route level to account for within-route correlation over time.

The coefficient of interest, β, captures the differential change in trade volumes on treated routes relative to control routes after the shock. The specification includes route-product fixed effects γ_{rp}, which absorb time-invariant heterogeneity across routes and products, and time fixed effects δ_t, which capture global shocks affecting all routes simultaneously. Standard errors are clustered at the route level.

2.2.3 Dynamic effects

To analyze adjustment dynamics, I estimate a dynamic specification:

$$\Delta \ln(TradeTons_{rpt}) = \alpha + \sum_k \beta_k(Phase_{k,t} \times Treated_r) + \gamma_{rp} + \delta_t + \epsilon_{rpt}. \tag{8}$$

Here, $Phase_{k,t}$ denotes different phases following the shock. This specification allows me to trace how the treatment effect evolves over time as the disruption persists.

2.2.4 Micro-level shipping adjustments

I complement the route-level analysis with voyage-level regressions that examine how shipping behavior adjusts to delays and whether this adjustment changes following the Red Sea shock. The micro-level specification is:

$$\begin{aligned} \ln(Cargo_{it}) = \alpha_i + \mu_r + \delta_t + \theta AbnormalDelay_{it} \\ + \gamma\,(AbnormalDelay_{it} \times Treated_i) \\ + \rho\,(Treated_i \times Post_t) \\ + \lambda\,(AbnormalDelay_{it} \times Post_t) \\ + \beta\,(AbnormalDelay_{it} \times Treated_i \times Post_t) + \epsilon_{it}. \end{aligned} \tag{9}$$

The dependent variable is the logarithm of cargo transported in voyage i at time t. The

[2]Because trade is measured at departure, the estimated decline captures a reduction in shipments initiated, not merely delays in arrival. The coefficient β thus reflects a real contraction in trade activity rather than a mechanical consequence of longer transit times.

coefficient θ captures the relationship between abnormal delays and cargo size on control routes before the shock. The coefficient γ captures how this relationship differs on treated routes before the shock, while λ captures how the delay–cargo relationship changes after the shock on control routes. The coefficient ρ captures the differential post-shock change in cargo levels on treated routes.

The coefficient of primary interest, β, is the triple interaction between abnormal delay, treatment status, and the post-shock period. It captures whether the relationship between abnormal delays and cargo size changes differentially on treated routes following the Red Sea disruption. A negative coefficient indicates that the cargo-consolidation response to longer delays becomes relatively weaker on exposed routes after the shock.

The specification includes vessel, route, and month fixed effects. Vessel fixed effects absorb time-invariant vessel characteristics, route fixed effects account for persistent differences across maritime routes, and month fixed effects capture common time shocks. Because treatment status is time invariant and the post-shock indicator is common across routes within a given month, their standalone effects are absorbed by the fixed effects. Standard errors are clustered at the vessel level.

This framework allows me to test whether the Red Sea disruption altered the ability of exposed routes to adjust to longer transit times through cargo consolidation, consistent with the literature on time as a trade cost (Hummels and Schaur, 2013) and on supply-chain transmission (Amiti and Weinstein, 2011).

2.2.5 Rerouting behavior

To study network-level adjustments, I estimate the following specification:

$$Reroute_{it} = \alpha_i + \delta_t + \beta(Treated_i \times Post_t) + \epsilon_{it}, \tag{10}$$

where $Reroute_{it}$ is a binary indicator capturing whether a voyage deviates from its usual maritime trajectory. Specifically, the variable equals one if a vessel either deviates from the Suez route, for instance by transiting around the Cape of Good Hope, or exhibits a substantial increase in transit time relative to its pre-shock baseline (computed at the route level), and zero otherwise.

This definition allows the rerouting measure to capture both direct geographical deviations (i.e., bypassing the Suez Canal via the Cape of Good Hope) and indirect adjustments through longer alternative trajectories.

The coefficient β captures the change in the likelihood of rerouting (expressed in percentage points) for exposed routes after the shock.

2.2.6 Identification

Identification relies on the assumption that, absent the Red Sea disruption, treated and control routes would have followed parallel trends. I assess this assumption through pre-trend analyses in the dynamic specification and through placebo tests.

The empirical strategy follows the identification logic of Feyrer (2021), who exploits exogenous variation in maritime routes, and Donaldson (2018), who studies the effect of transport infrastructure on trade. In my setting, the exogeneity of the shock stems from its geopolitical origin, while treatment status is defined using pre-shock route characteristics. Overall, this framework allows me to trace how the maritime disruption translates into shipping delays, rerouting adjustments, and lower trade volumes on exposed routes.

3 Results

3.1 Shipping disruptions: micro evidence from AIS data

This section begins by documenting the micro-level shipping disruptions caused by the Red Sea shock before turning to their aggregate trade implications. The analysis highlights two complementary adjustment margins: internal operational adjustments and network-level rerouting.

My identification exploits cross-route variation in exposure to the Red Sea corridor. Treated routes correspond to routes that structurally rely on the Suez Canal in the pre-shock period, while control routes do not. This distinction allows me to isolate the effect of the geopolitical disruption from general shipping fluctuations.

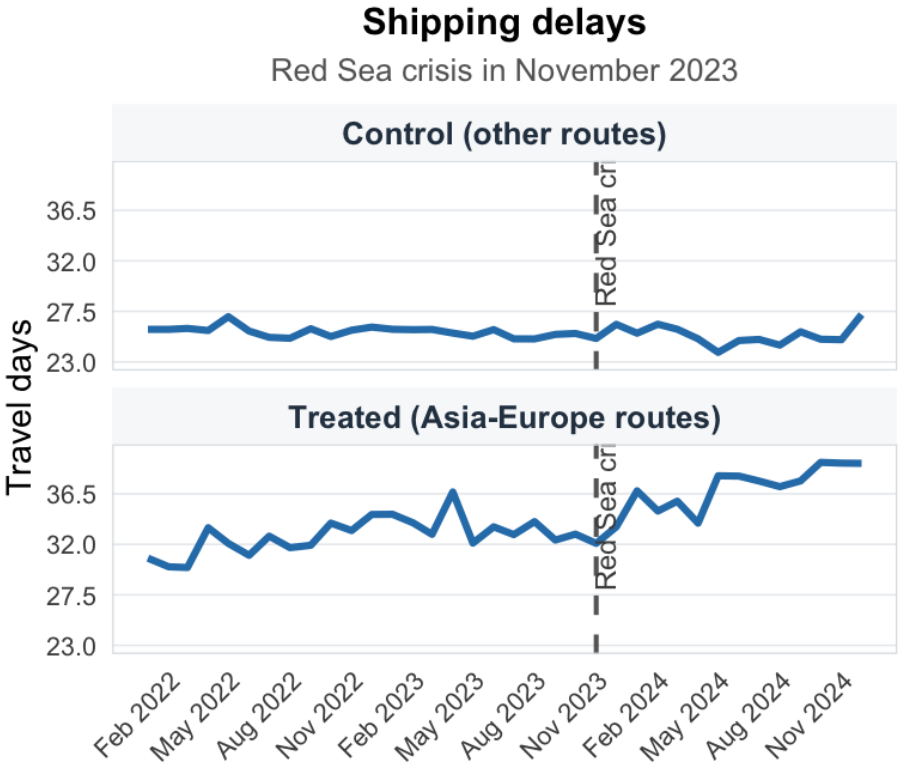


Figure 1: Shipping delays around the Red Sea crisis. The figure reports average travel times for treated routes (i.e., routes structurally dependent on the Suez/Red Sea corridor) and control routes. The vertical dashed line indicates the onset of the Red Sea disruption in November 2023. Travel times increase sharply for treated routes following the shock, while control routes remain relatively stable.

Figure 1 shows a clear increase in travel times for treated routes following the onset of the disruption, while control routes remain broadly stable.

Table 1 exploits more than 1.2 million vessel-level voyages reconstructed from AIS data. Column (1) shows that longer delays are associated with larger cargo loads per voyage. The estimated coefficient of 0.0028 implies that a one-day increase in abnormal delay is associated with a 0.28 percent increase in cargo tons, reflecting shipment consolidation. This mechanism reflects standard operational adjustments to longer transit times and is consistent with the literature on time as a trade cost (Hummels and Schaur, 2013).

Column (2) confirms this mechanism using an extensive-margin proxy. The coefficient of 0.0013 indicates that longer delays increase the probability that a voyage belongs to the upper half of the cargo distribution, indicating a reallocation toward larger shipments.

Column (3) examines whether this consolidation response differs across exposed and non-exposed routes and whether this difference changes after the Red Sea shock. Before the shock, the relationship between abnormal delays and cargo size is already significantly weaker on treated routes. The baseline coefficient on abnormal delay is positive (0.0028), while its interaction with treatment is negative (-0.0035), implying a slightly negative relationship on treated

routes before the shock.

Importantly, the triple interaction between abnormal delay, treatment status, and the post-shock period is negative and statistically significant (-0.0013). This indicates that the difference in consolidation responses between treated and control routes widens further after the disruption. While control routes continue to respond to longer delays by increasing cargo per voyage, this adjustment margin becomes even weaker on exposed Asia–Europe routes after the shock. The evidence therefore suggests that the Red Sea disruption further constrained an operational adjustment margin that was already weaker on exposed routes.

While the existing literature shows that delays and supply chain frictions affect shipment behavior (Hummels and Schaur, 2013; Amiti and Weinstein, 2011), these results highlight an additional dimension of adjustment to geopolitical disruptions. Standard consolidation responses are more limited on routes exposed to a critical maritime corridor, and this difference becomes significantly larger following the Red Sea shock.

Taken together, these results show that standard adjustment mechanisms to longer transit times are more limited on exposed routes and that this gap widens significantly following a geographically concentrated geopolitical disruption affecting a key maritime chokepoint. Figure 2 illustrates this divergence: while control routes maintain a positive consolidation response both before and after the shock, the response on exposed Asia–Europe routes becomes more negative following the Red Sea disruption.

Table 1: Micro Evidence from AIS Vessel-Level Data

	(1) Load Size ln(Cargo Tons)	(2) Large Shipment Large Shipment	(3) Heterogeneity ln(Cargo Tons)
Abnormal Delay	0.0028***	0.0013***	0.0028***
	(0.00005)	(0.00003)	(0.00006)
Abnormal Delay × Treated			-0.0035***
			(0.0003)
Abnormal Delay × Post			0.0006***
			(0.0001)
Abnormal Delay × Treated × Post			-0.0013***
			(0.0004)
Vessel FE	Yes	Yes	Yes
Route FE	Yes	Yes	Yes
Month FE	Yes	Yes	Yes
Observations	1,249,651	1,249,651	1,249,528

Notes: Abnormal delay is measured in days. Treated routes correspond to Asia–Europe and Europe–Asia routes exposed to the Red Sea–Suez corridor. Post equals one from November 2023 onward. Column (3) allows the relationship between abnormal delay and cargo size to vary across treated and control routes and before and after the shock. Vessel, route, and month fixed effects are included as indicated. Standard errors clustered at the vessel level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

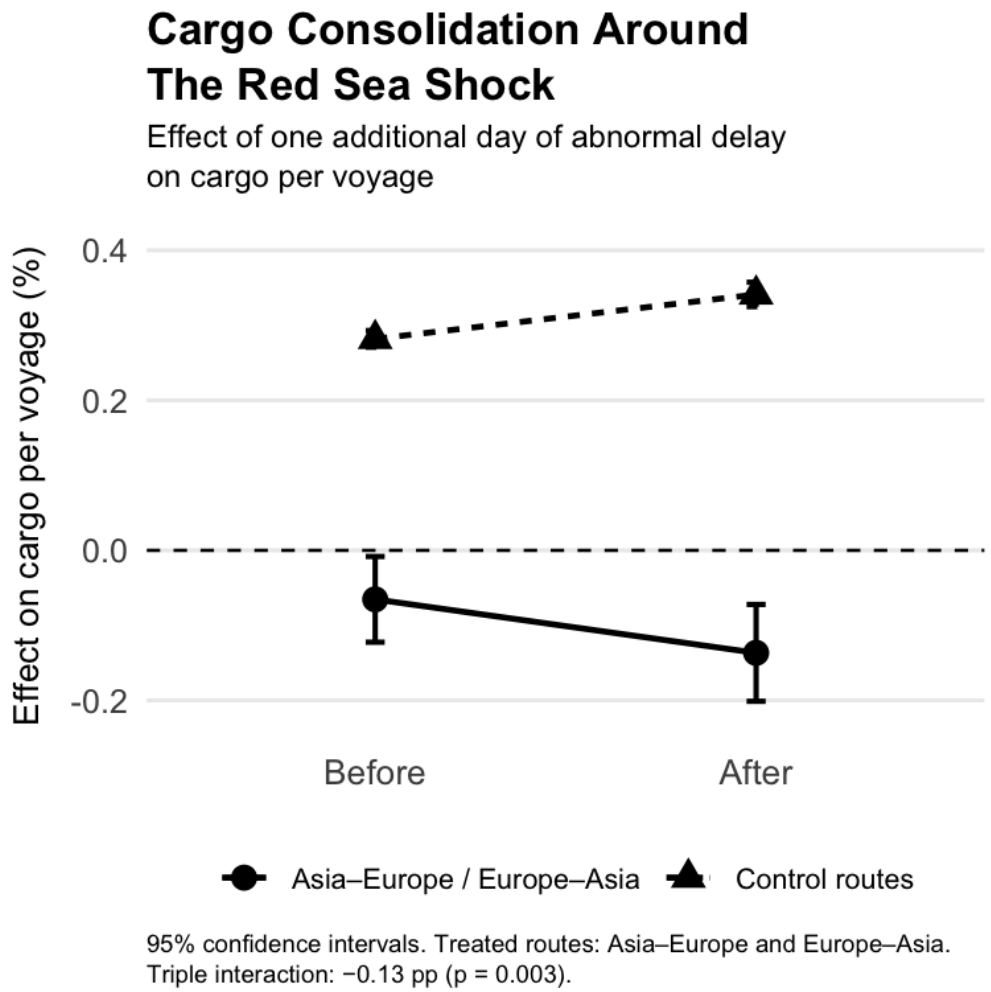


Figure 2: Cargo Consolidation Before and After the Red Sea Shock

3.2 Rerouting dynamics and network adjustments

In addition to internal operational adjustments, vessels respond to the Red Sea disruption by altering their maritime trajectories. In particular, exposed vessels substitute the Suez Canal route and reroute around the Cape of Good Hope. However, such rerouting is inherently costly in terms of transit time, fuel consumption, and operational constraints, suggesting that this adjustment margin remains limited in the short run.

To quantify this mechanism, Table 2 estimates the effect of the shock on rerouting behavior using a difference-in-differences framework.[3]

The coefficient on Treated × Post is positive and statistically significant (0.108), indicating that vessels operating on exposed routes are about 10.8 percentage points more likely to reroute

[3]Rerouting is defined as a binary indicator equal to one if a voyage either deviates from the Suez route, for instance by transiting around the Cape of Good Hope or exhibits a substantial increase in transit time relative to its pre-shock baseline (defined at the route level). This definition captures both direct route substitution and indirect rerouting through longer alternative trajectories.

and adopt longer alternative maritime routes following the onset of the Red Sea crisis.

While economically meaningful, this magnitude remains moderate relative to the scale of the disruption. This suggests that rerouting constitutes a real but constrained adjustment margin, consistent with the presence of substantial frictions in maritime transport networks. In particular, capacity constraints, contractual rigidities, and the high operational costs associated with longer routes may limit the ability of shipping firms to fully substitute away from the Suez corridor.

Table 2: Rerouting Behavior

	Rerouted
Treated × Post	0.108***
	(0.0157)
Route FE	Yes
Month FE	Yes
Observations	588,413

Notes: The dependent variable is a binary indicator equal to one if a voyage is rerouted. Standard errors clustered at the route level in parentheses.

Figure 3 provides a dynamic view of these adjustments using an event-study specification. The estimates show a clear increase in rerouting immediately following the onset of the crisis, with no evidence of differential pre-trends between treated and control routes. This supports the parallel trends assumption and strengthens the causal interpretation of the estimates.

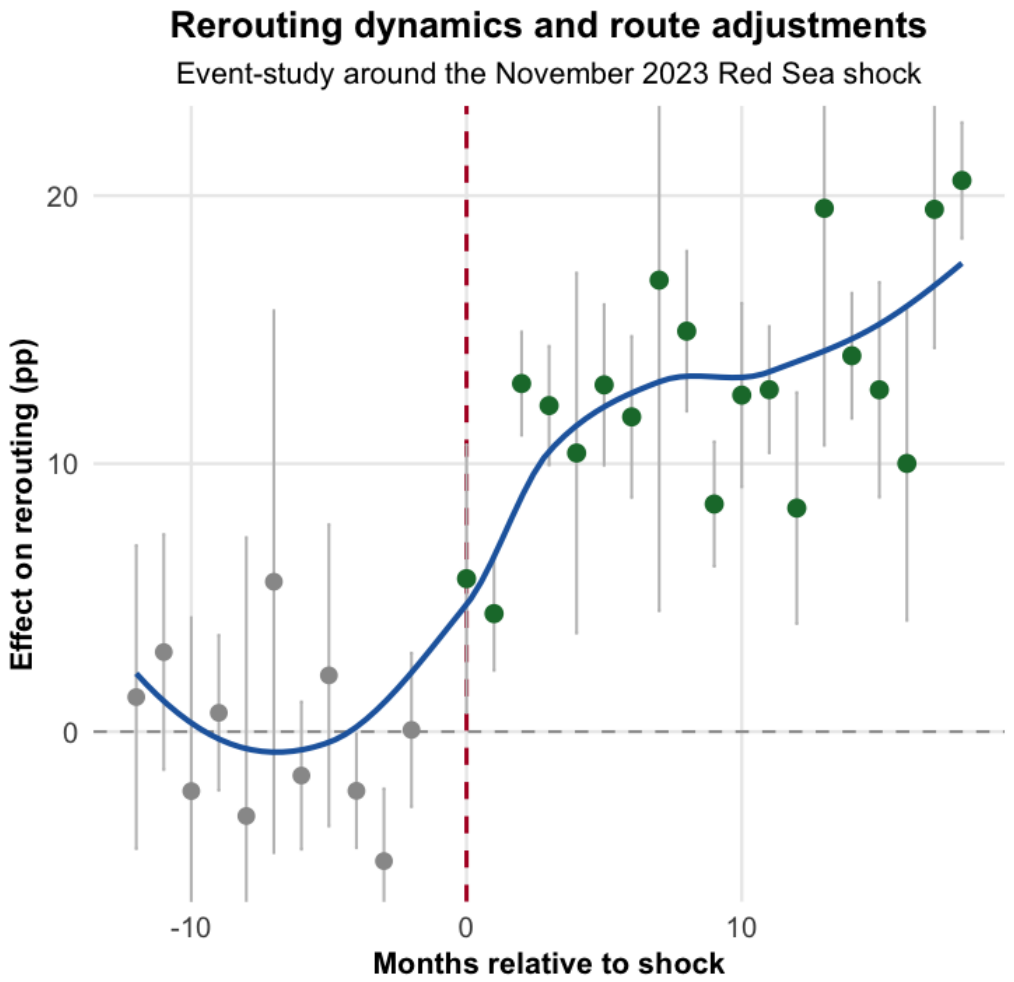


Figure 3: **Rerouting dynamics and route adjustments** This figure reports event-study estimates of the effect of the November 2023 Red Sea shock on rerouting behavior, expressed in percentage points. The vertical dashed line marks the onset of the crisis. Rerouting increases sharply after the shock, with no evidence of pre-trends. However, the magnitude remains moderate, suggesting that switching maritime routes is costly and operationally constrained. Error bars represent 95% confidence intervals.

Overall, rerouting increases significantly after the shock but remains quantitatively limited relative to the scale of the disruption. This pattern highlights that, while maritime networks exhibit some degree of flexibility, their ability to reconfigure in response to localized geopolitical disruptions is constrained. As a result, shocks affecting critical chokepoints such as the Red Sea do not simply reallocate trade volumes across routes, but instead generate additional frictions that contribute to the observed decline in trade volumes.

3.3 Trade effects of the Red Sea shock

I now turn to the implications for international trade volumes. I estimate the model in first differences, which removes persistent differences in trade levels across route-product pairs. Following the inference considerations emphasized by Bertrand et al. (2004), standard errors are clustered at the route level to account for within-route correlation over time.

Table 3 presents the main results. Column (1) shows that the coefficient on *Treated* × *Post* is negative and statistically significant. The estimate of −0.0588 implies a statistically and economically meaningful decline of approximately 5.8 percent in trade volumes on exposed routes following the shock. Since the analysis is conducted at the route-product level, this effect reflects an average contraction across all traded goods. Because trade is measured at the time of departure, this decline captures a reduction in shipments initiated, rather than delays in arrival. This interpretation indicates that the estimated decline reflects fewer shipments initiated on exposed routes rather than a mechanical consequence of longer arrival times.

This result is consistent with the broader literature showing that transport time is a key determinant of trade volumes (Hummels and Schaur, 2013; Feyrer, 2021). More generally, it aligns with evidence that disruptions to transport infrastructure and shipping conditions reduce trade volumes by increasing effective trade frictions (Limao and Venables, 2001).

Column (2) presents a dynamic specification. It replaces the aggregate post-treatment indicator with phase-specific interactions, allowing the treatment effect to vary over time. The phases are defined at the monthly level, with November 2023 corresponding to the first fully treated month, December 2023 capturing the immediate adjustment period, and January 2024 onward reflecting persistent disruption. Since these phase indicators fully span the post-treatment period, the baseline Treated × Post term is omitted to avoid perfect collinearity.

The coefficient for Phase 3 is negative and statistically significant (−0.0659), while earlier phases are imprecisely estimated. This pattern suggests that trade declines materialize progressively as disruptions persist. In the short run, shipping adjustments may be limited due to pre-existing contracts, scheduling rigidities, and inventory smoothing. As the disruption continues, however, firms adjust shipment decisions, leading to a more pronounced contraction in trade volumes.

Figure 4 in Appendix A displays the dynamic effects by phase.

Table 3: Trade Effects of the Red Sea Shock (First-Difference)

	(1) Main DiD $\Delta \ln$(Trade Tons)	(2) Dynamic DiD $\Delta \ln$(Trade Tons)
Treated × Post	-0.0588**	
	(0.0207)	
Treated × Phase 1		-0.0653
		(0.1147)
Treated × Phase 2		0.1067
		(0.1767)
Treated × Phase 3		-0.0659**
		(0.0187)
Route-product FE	Yes	Yes
Month FE	Yes	Yes
Observations	20,753	20,753

3.4 Heterogeneity by product value

Finally, Table 4 explores heterogeneity across product types using unit values constructed from UN Comtrade data.

Low-value goods experience a statistically significant contraction (-0.0367), while medium-value goods exhibit a stronger negative effect (-0.0817). In contrast, the effect for high-value goods is smaller in magnitude and not statistically significant (-0.0361).

The results reveal a non-monotonic relationship between product value density and vulnerability to the disruption. Medium-value goods experience the largest contraction, followed by low-value goods, while high-value goods appear comparatively resilient. The resilience of high-value goods is consistent with the Alchian-Allen mechanism, whereby per-unit transport costs weigh less heavily on higher-value shipments (Hummels and Skiba, 2004).

More generally, the literature shows that sensitivity to shipping frictions depends not only on value density but also on time sensitivity and firms' inventory and shipment decisions (Hummels and Schaur, 2013; Alessandria et al., 2010). The particularly strong response of medium-value goods therefore represents an additional empirical finding that is not captured by a simple monotonic relationship between product value and vulnerability to transport disruptions.

Understanding the mechanisms underlying this non-monotonic pattern, including potential differences in time sensitivity, inventory management, and sourcing flexibility across products, remains an important avenue for future research.

Figure 5 in Appendix A illustrates this heterogeneous response. Table 13 reports the main products included in each value bucket. Products are classified into low-, medium-, and high-value groups based on terciles of the distribution of HS6-level unit values computed from UN Comtrade data.

Table 4: Heterogeneity by Product Value (First-Difference)

	Low Value $\Delta \ln$(Trade)	Mid Value $\Delta \ln$(Trade)	High Value $\Delta \ln$(Trade)
Treated × Post	-0.0367**	-0.0817**	-0.0361
	(0.0130)	(0.0232)	(0.0309)
Route-product FE	Yes	Yes	Yes
Month FE	Yes	Yes	Yes
Observations	7,539	6,994	6,220

Notes: Estimates from first-difference specifications. Standard errors clustered at the route level in parentheses. Products are classified into low-, medium-, and high-value groups using terciles of the distribution of HS6-level unit values computed from UN Comtrade data. *** p<0.01, ** p<0.05, * p<0.1.

4 Robustness

4.1 Placebo test

To assess whether the main results are driven by spurious time variation rather than the actual Red Sea disruption, I conduct a placebo test by assigning a fake shock to November 2022. If the identification strategy is valid, the placebo coefficient should be close to zero and statistically insignificant.

Table 5 shows that the placebo coefficient is 0.0013 and statistically insignificant, providing no evidence of a spurious treatment effect prior to the actual disruption.

Following recent recommendations in the difference-in-differences literature (Hagemann, 2019), which emphasize the importance of placebo inference to validate identifying assumptions, I also implement additional placebo tests using alternative fake shock dates. The results, reported in Table 7 in the Appendix, show coefficients that are small and statistically insignificant across all placebo specifications, reinforcing the absence of spurious pre-trends and supporting the credibility of the identification strategy.

Table 5: Placebo test (fake shock in November 2022)

	$\Delta \ln(\text{Trade})$
Treated × Post (Placebo)	0.0013
	(0.0168)
Route-product FE	Yes
Month FE	Yes
Observations	9,556

Notes: The placebo test assigns a fake shock to November 2022. The coefficient is not statistically significant, supporting the validity of the identification strategy. Standard errors clustered at the route level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

4.2 Additional robustness checks

To further assess the validity of the main results, I conduct several additional robustness checks. First, I restrict the sample to a balanced panel of routes that are consistently observed throughout the sample period. This ensures that the results are not driven by routes entering or exiting the sample over time. Column (2) of Table 6 shows that the coefficient remains stable and significant at -0.0656.

Second, I exclude trans-Pacific routes (Americas-Asia and Asia-Americas), which could potentially be affected by other shocks such as disruptions to the Panama Canal. This test

reduces concerns that the estimated effect is driven by contemporaneous disruptions affecting other major maritime corridors. Column (3) reports a coefficient of −0.0666, which is statistically significant and close to the baseline estimate.

Third, I use an alternative treatment definition based on pre-shock exposure to the Red Sea / Suez corridor. Exposure is measured as the share of voyages on a given route transiting through the Suez Canal prior to November 2023. I then classify routes with pre-shock exposure above the median as treated. This addresses concerns that the results depend on the baseline treatment definition. Column (4) shows a coefficient of −0.0252, which remains statistically significant at the 5% level.

Fourth, I cluster standard errors at the route-product level instead of the route level to account for potential correlation at a finer level of aggregation. Column (5) confirms that inference is robust to the clustering choice, with a coefficient of −0.0644 significant at the 1% level.

Fifth, I exclude the smallest 10% of trade volumes to ensure that the results are not driven by noisy observations with low trade volumes.[4] Column (6) yields a coefficient of −0.0750, which is larger than the baseline estimate and significant at the 1% level.

Table 6: Additional robustness checks

	(1) Main $\Delta \ln$(Trade)	(2) Balanced $\Delta \ln$(Trade)	(3) Excl. Panama $\Delta \ln$(Trade)	(4) Exposure-based $\Delta \ln$(Trade)	(5) Alt. cluster $\Delta \ln$(Trade)	(6) Large volumes $\Delta \ln$(Trade)
Treated × Post	-0.0588**	-0.0656***	-0.0666***		-0.0644***	-0.0750***
	(0.0207)	(0.0218)	(0.0210)		(0.0240)	(0.0107)
High Exposure × Post				-0.0252**		
				(0.0121)		
Route-product FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	20,753	10,782	10,259	11,495	11,495	10,379

Notes: Column (1) reports the baseline first-difference specification. Column (2) uses a balanced panel of routes. Column (3) excludes trans-Pacific routes. Column (4) uses an alternative treatment definition based on pre-shock exposure to the Red Sea / Suez corridor, measured as the share of voyages transiting through the Suez Canal prior to November 2023. Routes with exposure above the median are classified as treated. Column (5) clusters standard errors at the route-product level. Column (6) excludes the smallest 10% of trade volumes. Standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The results in Table 6 show that the estimated decline in trade volumes on exposed routes is robust across alternative specifications. All specifications yield negative and statistically significant coefficients. The balanced panel specification (−0.0656) and the specification excluding trans-Pacific routes (−0.0666) produce estimates close to the baseline effect of −0.0588. The exposure-based treatment specification yields a smaller but still statistically significant coefficient (−0.0252), showing that the negative effect persists under an alternative definition of exposure. The alternative clustering specification confirms that inference is not sensitive to the

[4]Very small trade flows may exhibit greater relative variability, making the estimates more sensitive to measurement noise.

clustering choice (-0.0644, $p < 0.01$). Finally, excluding the smallest trade volumes yields a larger coefficient (-0.0750, $p < 0.01$), suggesting that the estimated decline is stronger when focusing on economically meaningful trade volumes.

5 Discussion

The results provide a coherent picture of how geopolitical disruptions propagate through maritime transport systems and affect trade outcomes. Three main insights emerge from the empirical analysis.

First, the Red Sea disruption operates primarily through the transport layer of the global economy. The AIS evidence shows a clear increase in transit times on exposed routes, together with important differences in vessels' ability to adjust to longer delays. While longer delays are generally associated with larger cargo loads per voyage, consistent with shipment consolidation, this adjustment margin is already significantly weaker on exposed Asia-Europe routes before the shock. Importantly, the difference widens significantly after the Red Sea disruption, as shown by the negative and significant interaction between abnormal delay, treatment status, and the post-shock period. This suggests that the disruption further constrained an operational adjustment margin that was already more limited on exposed routes.

This finding complements the literature on time as a trade cost (Hummels and Schaur, 2013) by showing that the consequences of transport disruptions depend not only on the increase in transit time itself, but also on the ability of shipping networks to adjust operationally to longer voyages. In the presence of severe geopolitical disruptions, such adjustment margins may become increasingly constrained.

Second, network-level adjustments play an important role in response to the shock. Rerouting increases significantly following the disruption, with vessels on exposed routes becoming about 10 percentage points more likely to reroute or adopt substantially longer maritime trajectories. This indicates that maritime networks exhibit some degree of flexibility in response to localized disruptions. However, this adjustment coexists with a substantial increase in transit times and a decline in trade volumes, suggesting that rerouting does not eliminate the increase in effective transport frictions. This result is consistent with recent work emphasizing the role of transport networks in propagating shocks (Heiland et al., 2025). In this context, the Red Sea disruption leads both to network reconfiguration and to an increase in overall trade frictions.

Third, these transport disruptions translate into economically meaningful trade effects. The estimated decline of approximately 5.8% in trade volumes on exposed routes is consistent with the idea that delays and uncertainty act as implicit trade barriers. Trade effects are heterogeneous across goods: medium-value products exhibit the strongest contraction, while low-value goods are also significantly affected and high-value goods appear more resilient. This pattern points to a non-linear relationship between product value and sensitivity to transport disruptions. While the relative resilience of high-value goods is consistent with existing

evidence on the role of value density in transport costs, the particularly strong contraction among medium-value goods constitutes an additional empirical finding. Understanding the mechanisms behind this pattern, including potential differences in time sensitivity, inventory management, and sourcing flexibility across products, remains an important avenue for future research.

Taken together, these findings suggest that geopolitical shocks affecting maritime chokepoints operate through a combination of increased transport frictions, constrained operational adjustment, and network reconfiguration. Vessels respond to the disruption by rerouting, but exposed routes simultaneously exhibit a further weakening of their cargo-consolidation response. These adjustments are not sufficient to offset the broader impact on trade volumes.

6 Conclusion

This paper studies how a localized geopolitical disruption affecting a strategic maritime chokepoint propagates through shipping networks and translates into global trade outcomes. Exploiting the Red Sea crisis as an exogenous shock, I combine high-frequency AIS data with a difference-in-differences framework to provide new evidence on the micro-foundations of trade disruptions.

Three main findings emerge. First, the disruption leads to a significant increase in transit times on exposed routes, reflecting a deterioration in shipping conditions. At the vessel level, longer delays are generally associated with cargo consolidation, but this adjustment margin is substantially weaker on exposed routes and becomes significantly more constrained following the Red Sea shock. Second, vessels respond through increased rerouting, indicating a degree of network flexibility. Third, despite this network adjustment, trade volumes on exposed routes decline by approximately 5.8%, suggesting that rerouting alone is not sufficient to neutralize the effects of the disruption.

These results highlight a key mechanism through which geopolitical shocks affect trade: they not only increase transport frictions, but also alter the margins through which shipping networks adjust to those frictions. The coexistence of longer transit times, constrained cargo consolidation, increased rerouting, and lower trade volumes illustrates how a localized disruption can propagate from vessel-level operations to aggregate trade outcomes. In this context, maritime chokepoints emerge as critical nodes of vulnerability in global value chains.

More broadly, the findings point to a growing role of geopolitical risk in shaping the physical structure of global trade. Maritime corridors are no longer neutral infrastructures, but may become strategically contested and conditionally accessible.

An important avenue for future research is to examine whether the decline in trade volumes observed on exposed routes is accompanied by a reallocation of trade toward alternative partners or routes. Understanding such adjustments would provide further insight into how firms reorganize global value chains in response to geopolitical risk.

References

Abadie, A. and Gardeazabal, J. (2003). The economic costs of conflict: A case study of the basque country. *American economic review*, 93(1):113–132.

Alessandria, G., Kaboski, J. P., and Midrigan, V. (2010). Inventories, lumpy trade, and large devaluations. *American Economic Review*, 100(5):2304–2339.

Alfaro, L. and Chor, D. (2023). Global supply chains: The looming "great reallocation". *NBER Working Paper*.

Allen, T. and Arkolakis, C. (2022). The welfare effects of transportation infrastructure improvements. *The Review of Economic Studies*, 89(6):2911–2957.

Amiti, M. and Weinstein, D. E. (2011). Exports and financial shocks. *Quarterly Journal of Economics*, 126(4):1841–1877.

Anderson, J. E. and van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93(1):170–192.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*, 119(1):249–275.

Besley, T. and Mueller, H. (2012). Estimating the peace dividend: The impact of violence on house prices in northern ireland. *American Economic Review*, 102(2):810–833.

Brancaccio, G., Kalouptsidi, M., and Papageorgiou, T. (2020). Geography, transportation, and endogenous trade costs. *Econometrica*, 88(2):657–691.

Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4):1194–1225.

Collier, P. (1999). On the economic consequences of civil war. *Oxford economic papers*, 51(1):168–183.

Donaldson, D. (2018). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review*, 108(4-5):899–934.

Feyrer, J. (2021). Distance, trade, and income—the 1967 to 1975 closing of the suez canal as a natural experiment. *Journal of Development Economics*, 153:102708.

Hagemann, A. (2019). Placebo inference on treatment effects when the number of clusters is small. *Journal of Econometrics*, 213(1):190–209.

Handley, K. and Limão, N. (2017). Policy uncertainty, trade, and welfare: Theory and evidence for china and the united states. *American Economic Review*, 107(9):2731–2783.

Heiland, I., Moxnes, A., Ulltveit-Moe, K. H., and Zi, Y. (2025). Trade from space: Shipping networks and the global implications of local shocks. *Review of Economics and Statistics*, pages 1–45.

Hummels, D. and Skiba, A. (2004). Shipping the good apples out? an empirical confirmation of the alchian-allen conjecture. *Journal of political Economy*, 112(6):1384–1402.

Hummels, D. L. and Schaur, G. (2013). Time as a trade barrier. *American Economic Review*, 103(7):2935–2959.

Kalouptsidi, M. (2014). Time to build and fluctuations in bulk shipping. *American Economic Review*, 104(2):564–608.

Limao, N. and Venables, A. J. (2001). Infrastructure, geographical disadvantage, transport costs, and trade. *World Bank Economic Review*, 15(3):451–479.

Schøyen, H. and Bråthen, S. (2011). The northern sea route versus the suez canal: cases from bulk shipping. *Journal of Transport Geography*, 19(4):977–983.

A Appendix A: Additional Figures

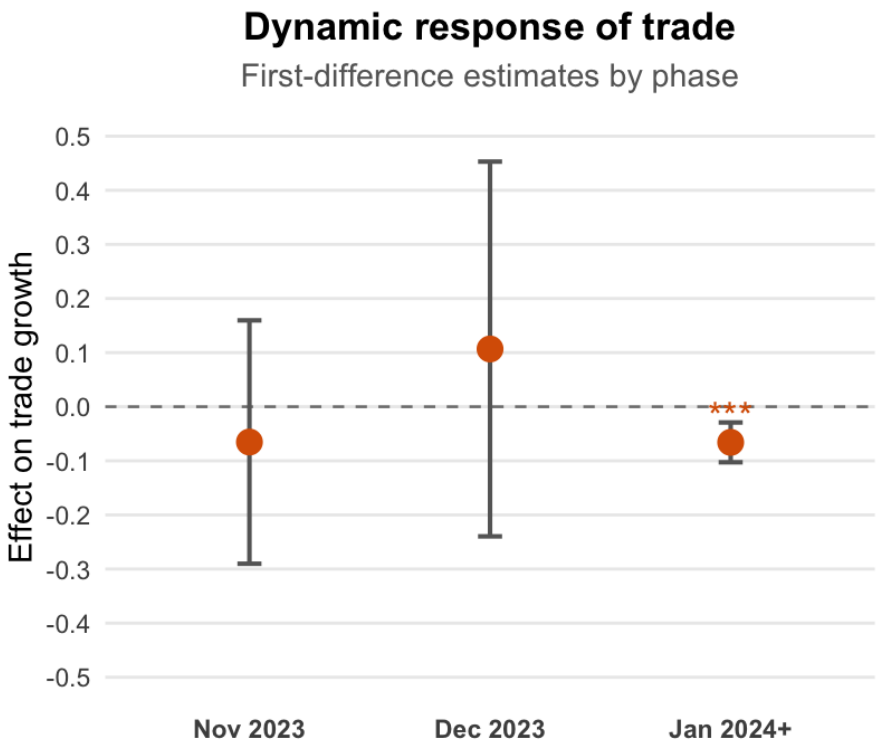


Figure 4: Dynamic response of trade to the Red Sea disruption

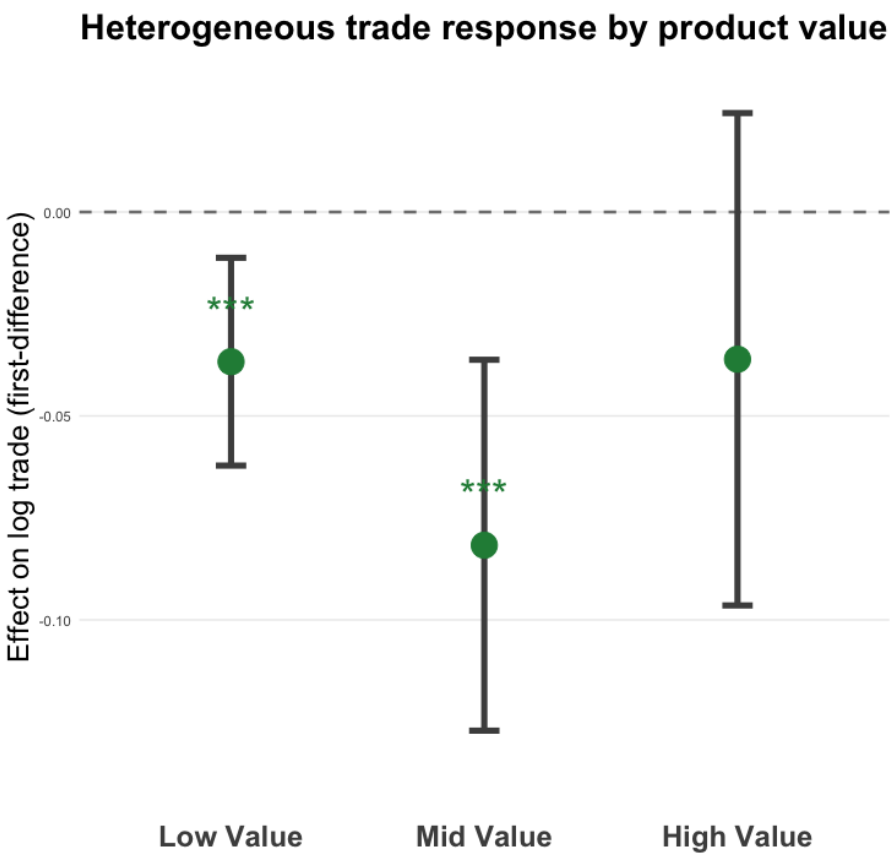


Figure 5: Heterogeneity by product value

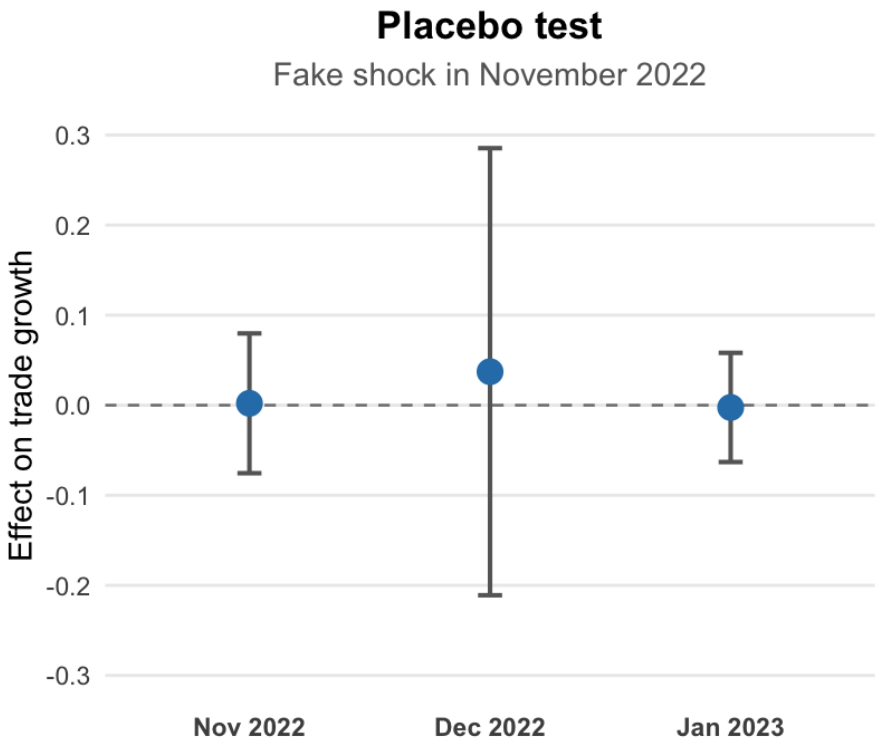


Figure 6: Placebo test: fake shock in November 2022

Table 7: Multiple Placebo Tests

	Nov 2021	May 2022	Nov 2022
Treated × Post	0.0407	-0.0029	0.0013
	(0.1287)	(0.0237)	(0.0168)
Route-product FE	Yes	Yes	Yes
Month FE	Yes	Yes	Yes
Observations	7,216	9,900	9,556

B Data Construction

This appendix details the construction of the main variables used in the empirical analysis.

Voyages are reconstructed from AIS port calls by linking consecutive port observations at the vessel level. A voyage is defined as the movement between two consecutive port calls, with departure and arrival times corresponding to the end and start of port stays.

Travel time is computed as the difference between arrival and departure timestamps, measured in days. I restrict the sample to plausible voyages with duration between 1 and 100 days.

Abnormal delays are defined as the deviation from a pre-shock baseline computed in 2022:

$$\text{Delay}_{rt} = \text{TravelTime}_{rt} - \text{Baseline}_{r,m}$$

where the baseline is the average travel time for route r and calendar month m in 2022. This definition controls for route-specific seasonality and allows the analysis to isolate changes in transport performance associated with the Red Sea disruption.

Routes connecting Asia and Europe are defined as treated, as they structurally rely on the Suez Canal prior to the shock. All other routes serve as the control group.

The post period is defined as starting in November 2023, corresponding to the onset of persistent attacks on commercial vessels in the Red Sea and the subsequent disruption of maritime traffic.

B.1 Descriptive statistics: Rerouting

Table 8: Cost of rerouting: travel days on exposed routes

Period	Mean	Median	Observations
Pre-shock (before Nov 2023)	45.42	44.11	10,827
Post-shock (Nov 2023 onward)	51.08	49.25	5,069
Change		+5.66 days (+12.5%)	

Notes: The table reports average and median travel times (in days) for voyages on Asia–Europe routes before and after the onset of the Red Sea disruption in November 2023. Transit times increase substantially in the post-shock period. This pattern is consistent with rerouting behavior and the use of longer or more congested alternative maritime routes, even when vessels do not fully deviate via the Cape of Good Hope.

B.2 Sample Construction

Table 9: Sample construction

Step	Observations
Raw AIS/Kpler observations	2,761,737
Estimation sample (non-missing cargo)	2,459,615
Route-product-month panel	58,428
HS6-matched sample	35,537

Notes: This table reports the successive steps of sample construction, from raw AIS/Kpler observations to the final estimation samples. The HS6-matched sample is used for the heterogeneity analysis by product value.

B.3 Representativeness of Estimation Samples

Table 10: Representativeness of estimation samples

Variable	Control	Full sample	High value	Low value	Matched sample	Treated
Mean log trade	11.809	11.764	12.097	12.237	12.327	11.207
Mean delay	-0.113	-0.015	-0.011	0.012	0.023	1.220
Treated share	0.000	0.074	0.078	0.075	0.083	1.000
Observations	54,124	58,428	10,230	12,418	35,537	4,304

Notes: This table compares the full route-product-month sample, the HS6-matched estimation sample, treated and control routes, and product groups by value. The matched sample represents approximately 61% of the total observations and is concentrated on the economically most relevant commodities. The main trade results are estimated on the full sample and therefore do not rely on the product matching.

B.4 Voyage Duration on Exposed Routes

Table 11: Voyage duration on exposed routes

Route	Period	Mean	Median	Observations
Asia-Europe	Pre	31.82	26.95	8,337
Asia-Europe	Post	38.93	35.61	3,889
Europe-Asia	Pre	30.97	31.05	9,538
Europe-Asia	Post	35.33	34.13	4,320

Notes: The pre-shock period includes voyages departing before November 2023, while the post-shock period includes voyages departing thereafter. Voyage duration increases substantially after the shock on exposed routes, consistent with rerouting and congestion effects.

B.5 Cargo Composition

Table 12: Cargo composition by group

Group	Full sample Obs. share	Estimation sample Obs. share	Full sample Cargo share	Estimation sample Cargo share
Coal	22.1%	24.8%	32.8%	32.8%
Iron Ore	10.8%	12.1%	30.2%	30.2%
Grains	8.1%	9.0%	5.7%	5.7%
Cement/Clinker	7.5%	8.4%	2.8%	2.8%
Steel	7.4%	8.3%	3.4%	3.4%
Metals	5.7%	6.4%	6.9%	6.9%
Fertilizer	4.8%	5.3%	3.3%	3.3%
Minerals	4.2%	4.7%	3.4%	3.4%
Oilseeds	3.9%	4.4%	3.9%	3.9%
Grains/Oilseeds	2.9%	3.3%	1.2%	1.2%

Notes: The table reports the composition of the AIS/Kpler sample by cargo group. Shares are calculated over observations with non-missing cargo information. The sample is dominated by standard dry bulk commodities, especially coal and iron ore, which account for the largest shares of transported tonnage.

B.5.1 Product Classification

Table 13: Product Classification by Value Buckets

Value Bucket	Product	Observations
Low	Thermal coal	1,869
Low	Metallurgical coal	1,785
Low	Iron ore fines	1,551
Low	Cement	1,438
Low	Clinker	1,243
Mid	Wheat	1,803
Mid	Nitrogen fertilizers	1,532
Mid	Corn	1,360
Mid	Iron ore pellets	1,354
Mid	Barley	1,349
High	Aluminium	1,853
High	Phosphorus fertilizers	1,511
High	Sugar	1,405
High	Copper	1,303
High	Soybean meal	1,226

Notes: Products are classified into value buckets based on terciles of the global distribution of unit values computed from UN Comtrade data at the HS6 level. The table reports the most frequent products in each category.