Trade Chokepoints

Blaize Giangiulio*

November 2026

Abstract

Trade chokepoints have become targets of armed disruption, although existing work treats such episodes as exogenous hazards. This paper models chokepoint disruption as a strategic choice. In a multi-country model, a disruptor country may pay to raise trade costs between other countries, diverting demand toward itself. Disruption therefore improves the disruptor's terms of trade even absent any geopolitical motive, and the gain grows with the corridor's importance to world trade. Retaliation that raises the disruptor's own trade costs extinguishes this economic motive while leaving a geopolitical motive untouched. I take the model to the 2023–24 Red Sea crisis using a new, network-based measure of logistical distance that captures each country pair's exposure to disruption of the Bab el-Mandeb Strait. Trade fell on exposed pairs, most sharply among Western-aligned pairs, and welfare losses rise with exposure while Iran's own gain is small. Treating the observed campaign as Iran's optimal choice reveals extraordinary leverage: slight intensification would gain Iran $48 to $78 million per year while costing the average Western economy $1.6 to $8.1 billion. Pure self-interest can rationalize the campaign only if it cost Iran under $78 million per year, below even the lowest estimate of Iran's annual support to the Houthis. Rationalizing it therefore requires a strictly positive weight on harming adversaries: in the benchmark calibration, Iran is willing to give up one dollar of its own welfare to destroy roughly $12 of Western welfare.

Keywords: chokepoints, trade costs, geoeconomics, structural gravity, general equilibrium
JEL Codes: F13, F14, F51, F52

*Drexel University. I thank Yoto Yotov, Constantinos Syropoulos, Ohyun Kwon, Jack Liang, Paul Ko, Peter Herman, and Minuk Kim for guidance throughout this project, and participants at the Philadelphia Trade Group for helpful comments. All errors are my own. The most recent version of this paper is available at https://blaizegiangiulio.com.

1 Introduction

A few narrow transportation passages—the Straits of Hormuz and Malacca, the Suez and Panama Canals, and the Bab el-Mandeb—funnel a disproportionate share of global trade (UNCTAD, 2024). Any actor capable of disrupting them thus possesses unique leverage over distant economies. The recent Red Sea Crisis greatly publicized this leverage. Following the outbreak of the Israel–Hamas war in October 2023, Iran-backed Houthi militants began attacking commercial vessels in the Bab el-Mandeb Strait, the narrow southern entrance to the Red Sea through which nearly all Suez-bound traffic must pass (Figure 1).[1] The corridor handles roughly 10 percent of world maritime trade and 22 percent of global container trade (UNCTAD, 2024).[2] Within months, a state-backed militia had redirected a large share of that traffic: by mid-2024, tonnage transiting the Suez Canal had fallen 70 percent below its 2023 average, while passages around the Cape of Good Hope rose 89 percent (UNCTAD, 2024), a detour that lengthens the voyage from Mumbai to London by 71 percent. The episode demonstrates how vulnerable global trade is to a relatively low-cost chokepoint disruption.

The Houthis declared that their campaign was in solidarity with Hamas, and the militants who carried out one of its most prominent early operations, the hijacking of the Bahamian-flagged Galaxy Leader, released footage of a helicopter flying a Palestinian flag beside a Yemeni one. The campaign was thus a *deliberate* choice by a state-backed actor against an important economic corridor. While existing work quantifies the effects of chokepoint disruptions on trade, treating them as exogenous shocks, this paper asks why an actor would choose to disrupt a chokepoint: under what conditions does an actor benefit from raising the cost of trade between others, and what does an observed disruption reveal about whether the disruptor's motives are economic or geopolitical—greed or grievance?

To answer these questions, I develop a multi-country endowment economy with quasilinear preferences over a freely traded numeraire and a CES bundle of differentiated national goods.[3] One country, the disruptor, has the unique ability to pay to raise bilateral trade costs uniformly

[1] The strait is 18 miles wide at its narrowest, limiting large vessel traffic to two two-mile-wide lanes (U.S. Energy Information Administration, 2019).

[2] Approximately 7.5 million barrels per day of oil and 36 billion cubic meters of liquefied natural gas—around 10 percent and 8 percent of global seaborne oil and gas trade, respectively—transited the canal between January and October 2023 (International Energy Agency, 2023). Weekly Suez container transits fell from approximately 110 around the onset of the attacks to 30 within three months (AXSMarine, 2024).

[3] Quasilinear preferences shut down income effects, delivering closed-form solutions. Section 3.8 numerically demonstrates that the disruptor's gain survives richer preferences, including a standard Armington model.

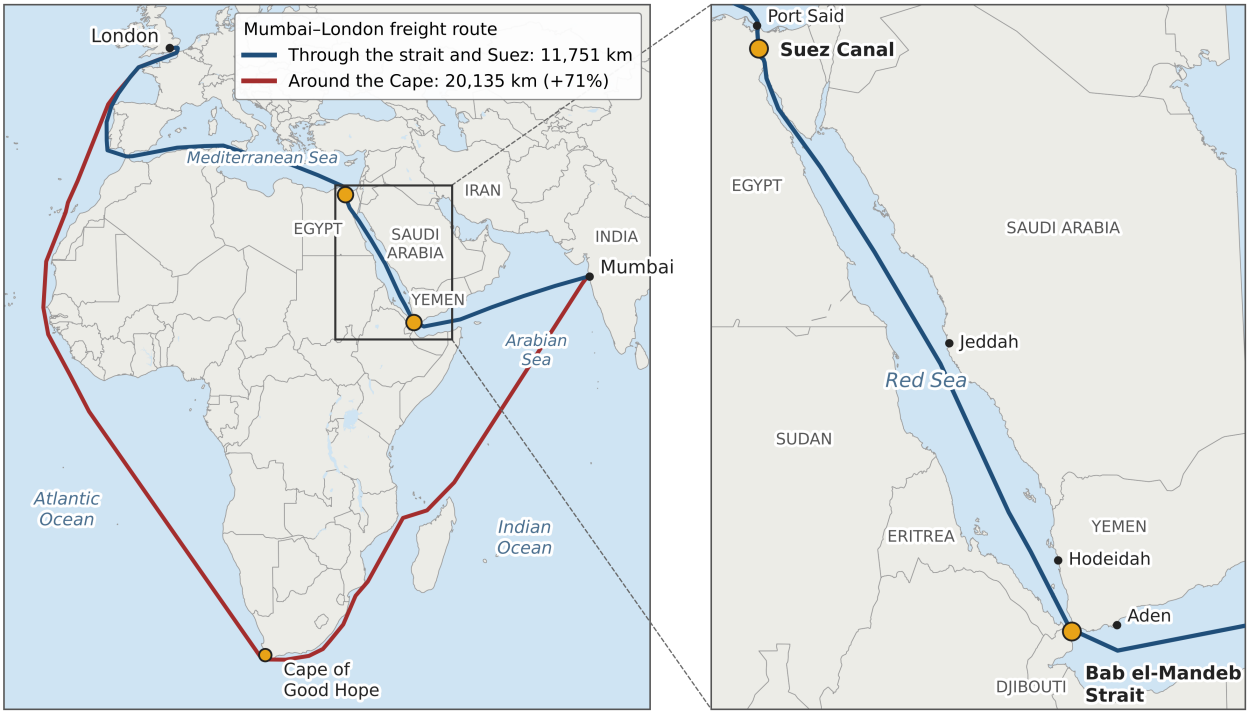


Figure 1: The Red Sea corridor and the Mumbai–London freight route. Left: the least-cost freight route between Mumbai and London on the global road and maritime network used to construct logistical distance (Section 4.1), through the Bab el-Mandeb Strait and the Suez Canal (blue) and, with the strait removed from the network, around the Cape of Good Hope (red). Right: the Red Sea corridor between the two chokepoints. Route lengths are network distances; this route enters the logistical distance between India and the United Kingdom with a large population weight.

among the members of a bloc through disruption ($\mathcal{M} = \{2, ..., M\}$), up to the point at which the disrupted route becomes as expensive as its bypass—the analogue of vessels abandoning the Bab el-Mandeb for the Cape of Good Hope.

The mechanism is most easily understood with three countries: a disruptor, its target, and a third country. The disruptor raises the cost of trade between the target and the third country. This induces both to substitute some of their purchases toward the relatively cheaper disruptor. This added demand then drives up the price of the disruptor's goods, improving its terms of trade, as those of the target and third country deteriorate. In the many-country setting, the disruptor instead raises trade costs among all members of a bloc such that countries outside of the bloc, including the disruptor, unambiguously benefit through the same channel.

The model yields four results. First, even absent any political motives, disruption is profitable: raising trade costs among bloc members diverts their demand toward the disruptor and other by-standers, improving their terms of trade. Second, a geopolitical benefit from harming a target

unambiguously raises the chosen intensity of disruption. Third, the disruptor's gain scales quadratically in the size of the disrupted bloc: severing a single trade relationship is negligible in a large world, while disrupting a corridor that carries a fixed share of world trade delivers a first-order gain, rationalizing why observed disruptions exclusively target major arteries. A fourth result concerns the possibility of retaliation. If the bloc countries commit to increasing the disruptor's trade costs by a fraction of the disruption they suffer, the disruptor's economic gain linearly decreases in that fraction and vanishes at a threshold strictly less than one. When retaliation is also costly to the bloc, the target included, the resulting decrease in the target's welfare increases the grievance-related payout, such that retaliation effectively deters greed but not grievance. The results are not unique to the Armington structure: they hold verbatim on Eaton–Kortum foundations, where trade adjusts on the extensive margin. Under a Ricardian framework, the gain requires only that the disruptor trade with the bloc.

The model's equilibrium admits a structural gravity representation (Anderson and van Wincoop, 2003), with one departure from the canon: trade costs on the disrupted links are endogenous, determined by the disruptor's optimal choice. The model is thus an augmented gravity system: a standard, estimable gravity block closed by a chokepoint first-order condition. This structure dictates the empirical division of labor. The gravity block identifies the trade response from non-strategic variation in trade costs. However, treating the observed disruption as the disruptor's optimal choice reveals how much weight it must place on harming its rivals, given incurred disruption costs. Welfare changes, in turn, admit a sufficient-statistic decomposition in the spirit of Arkolakis et al. (2012) and Costinot and Rodríguez-Clare (2014): the disruptor's terms-of-trade gain carries a predicted sign, a predicted incidence across bystanders and bloc members, and a predicted scaling in corridor size.

I take the framework to the Red Sea crisis. Measuring which country pairs were exposed requires a measure of distance that responds to infrastructure, which the standard great-circle measure does not: the straight line from Mumbai to London is the same whether or not the Suez Canal is open. I instead use logistical distance (Giangiulio et al., 2025), the length of the cheapest freight route between cities over a global network of highways and shipping lanes, aggregated to country pairs (Mayer and Zignago, 2011). Recomputing every route with the Bab el-Mandeb closed gives each country pair a measure of its exposure that is determined by geography alone: zero for pairs whose

trade never relied on the strait, such as Mexico and the United States, and large for pairs forced around the Cape of Good Hope, such as India and the United Kingdom, whose Mumbai–London route lengthens by 71 percent (Figure 1). Because the canal also carries tolls and delays that no distance measure records, pairs facing only small detours likely avoided it even before the crisis, so I only treat pairs whose routes substantially lengthen as exposed, with the cutoff chosen so that these pairs carry the share of world trade that independent estimates attribute to the corridor. Comparing trade on exposed pairs with trade on all others before and after the attacks,[4] I find that their trade fell by roughly five percent, with the decline concentrated among Western-aligned pairs, the campaign's declared targets. The decline began immediately and was largest in the first four months, then weakened by the spring of 2024 even though ships continued to reroute. That pattern points to a transient adjustment cost on top of the permanent cost of a longer route. The design builds on Feyrer's study of the 1967–75 Suez closure (Feyrer, 2021), using a continuous, theory-founded measure of exposure in place of a binary one.

The partial-equilibrium estimates recover the direct cost of the disruption but not its reallocation across the world economy, which is precisely the channel through which some countries may gain. Feeding the network-implied trade-cost shocks through the model's general equilibrium via exact hat algebra (Dekle et al., 2008), I find that welfare losses are increasing in exposure. Roughly two thirds of countries lose, and the corridor-dependent economies of the Gulf, the Horn of Africa, and South Asia lose the most. A minority of countries whose trade rarely transits the strait gain through the terms-of-trade channel that the theory predicts. Iran itself gains, but only about five thousandths of a percent of expenditure.

Finally, I invert the disruptor's first-order condition to recover the strategic primitives consistent with the observed event, treating the nineteen Western-aligned countries as the target and weighting their harm by GDP, yielding two results. The ability to disrupt is extraordinarily leveraged: at the margin, slightly intensifying disruption gains Iran between $48 and $78 million per year while costing the GDP-weighted average Western economy between $1.6 and $8.1 billion—roughly 34 to 103 dollars of harm inflicted per dollar of own gain. Additionally, because the first-order condition binds with equality at the chosen intensity, pure economic self-interest rationalizes the observed

[4] I estimate a structural gravity model by Poisson pseudo-maximum likelihood (Santos Silva and Tenreyro, 2006) with exporter-time, importer-time, and directional-pair fixed effects, following Head and Mayer (2014) and Yotov et al. (2016); the time-varying country effects absorb multilateral resistance (Baldwin and Taglioni, 2006).

campaign at exactly one marginal cost. With convex costs and no harmful retaliation from bloc countries, a purely greedy Iran could have sustained the campaign only if it cost under $64 million per year in the benchmark calibration, and under $78 million at any trade elasticity in the calibrated range. Independent estimates place Iran's annual support for the Houthis at $100 to $300 million—between 1.3 and 6.3 times that bound—and that support, which spans cash, fuel, weapons, and the training of Houthi fighters by Iran's Revolutionary Guard, is itself the latest installment of an investment in the Houthis dating to at least 2009. Greed alone therefore cannot rationalize the disruption even in the case most favorable to it: pairing the lowest trade elasticity in the calibrated range, which yields the highest greed-only bound, with the lowest cost estimate still leaves Iran's observed support 28 percent above what pure self-interest can justify, before counting any of its in-kind contributions. Mapping the external cost estimates to geopolitical weights, the implied λ ranges from 0.003 to 0.155 under GDP weights. At the benchmark elasticity and upper-bound cost, Iran is willing to sacrifice one dollar of its own welfare to destroy roughly $12 of Western welfare. These weights are floors: any retaliation Iran anticipated lowers the greed-only bound further.

This paper makes three contributions. It is the first to model a chokepoint disruption as an endogenous choice and to characterize when that choice is beneficial, delivering closed-form incidence results, a corridor-scaling result that explains which chokepoints are worth disrupting, and a deterrence result that separates the economic motive, which retaliation can stifle, from the geopolitical one. It introduces logistical distance as a tool for measuring exposure to infrastructure shocks and uses it to provide estimates of the Red Sea crisis's effects on trade and welfare that exploit continuous, geography-determined variation in exposure. And it shows how inverting a disruptor's optimality condition turns a quantitative trade model into a device for reading motives off observed behavior, yielding a quantitative bound on the role of economic self-interest in a live geopolitical episode.

The remainder of the paper proceeds as follows. Section 2 discusses how this work relates to the literature. Section 3 introduces the model, derives its analytical results, including the retaliation game, and establishes their numerical robustness; Appendix B re-derives them on Ricardian foundations. Section 4 constructs the treatment, presents the partial-equilibrium estimates, and computes the general-equilibrium counterfactuals. Section 5 inverts the disruptor's optimality condition to recover its implied geopolitical weight. Section 6 concludes.

2 Related Literature

This paper relates to three strands of economic literature: (1) geoeconomics and the economics of statecraft and power, (2) trade costs and the effects of distance shocks, and (3) structural gravity estimation and quantitative trade.

2.1 Geoeconomics

In his seminal work, Hirschman (1945) provides an early understanding of how states can exploit international trade to further their international standing and increase power. More recent works in the geoeconomics literature have focused on understanding the relationship between international trade and armed conflict. Martin et al. (2008) provide a theoretical and empirical assessment of the "liberal" political science perspective that international trade promotes peace. They ultimately find that, for nearby countries, trade tends to increase the probability of an armed conflict, a result that proves robust across numerous specifications. Glick and Taylor (2010) examine the impact that conflicts dating back to 1870 have had on international trade for both warring and neutral countries. To do this, the authors run a gravity specification that accounts for lagged and contemporaneous presence of bilateral war. They find that war has a significant, destructive, and lasting effect on trade between country pairs in which at least one country is a combatant. In a similar vein, Qureshi (2013) uses a gravity model to assess the impact that armed conflict in a neighboring country can bear on trade for both importers and exporters, ultimately finding that conflict bears significant deleterious and persistent spillover effects on trade in neighboring countries. More recently, Clayton et al. (2026) develop a geoeconomic model for understanding how hegemons can exploit their economic power over other countries to achieve policy goals and better their global standing when contracts are off limits. Similarly, Becko et al. (2025) show how in a bipolar world, two hegemons' optimal tariff choices reflect strategic incentives to maximize geopolitical alignment. In a recent *Annual Review of Economics* survey of geoeconomics, Mohr and Trebesch (2025) list "sabotage," like the destruction of critical infrastructure, as a geoeconomic tool, but they note that the economic and political consequences of its use are underdeveloped. Overall, the geoeconomic literature has tended to focus on how large, hegemonic countries exploit their economic prowess to further their policy objectives, while my paper focuses on how a smaller country with access

to a chokepoint can exploit that access to better itself economically while harming adversaries. The closest paper in spirit is Becko (2024), who draws a connection between trade restrictions as sanctions and as terms-of-trade manipulation. In my work, however, the disruptor manipulates its terms of trade by raising other countries' trade costs through geography rather than by taxing its own trade, so the gain accrues to every bystander and creates a free-rider margin that does not exist in the sanctions setting.

2.2 Trade Costs and Distance Shocks

As early as Ravenstein (1885), economists have known that the extent of bilateral interaction and flows between countries tends to decrease with higher distance due to increased friction. Eaton and Kortum (2002) provide an early gravity-based demonstration of the significant role that geography plays in trade patterns and welfare. Anderson and van Wincoop (2004) outline the inherent difficulty of modeling trade costs and the central role of distance in doing so accurately. A smaller strand of literature has sought to understand the effects of shocks to trade costs arising from distance changes. Feyrer (2021) exploits the unexpected eight-year closure of the Suez Canal to understand how pseudo-random changes in shipping distance between countries impacted both trade and welfare in the gravity setting. Similarly, Besedeš et al. (2024) analyze how conflict-related flight-path changes that increased air distance between certain country pairs impacted costly air trade. Some papers take a theoretical approach to understanding how infrastructure impacts trade. Donaldson (2018) exploits the construction of India's colonial railroad network to gauge how large transportation infrastructure projects decrease the cost of trade and impact welfare. Allen and Arkolakis (2014) develop a framework to understand the geospatial distribution of economic activity within the continental United States. They use this model, along with detailed geographic data, to estimate bilateral trade costs and perform counterfactual analysis that gauges the extent to which geography impacts regional price heterogeneity and how a theoretical removal of the interstate highway system would impact welfare. In a follow-up, Allen and Arkolakis (2022) develop a quantitative spatial model that incorporates traffic-related congestion to evaluate the welfare gains associated with network improvements. Brancaccio et al. (2025) study the post-COVID volatility of shipping markets and argue that rigidities in vessel and port capacity make transportation supply inelastic, so that prices respond sharply to relatively small demand shocks and disruptions

propagate into trade and inflation. Only relatively recently have publications sought to seriously integrate infrastructure-related realities like chokepoint disruptions and congestion into developing an understanding of bilateral friction and trade costs.

2.3 Structural Gravity and Quantitative Trade

With the advent of gravity modeling in the international trade literature, economists have vastly changed how they model trade costs. Early gravity papers accounted for bilateral trade costs by including trade-cost variables, like distance, colonial connection, and contiguity, directly in their regressions. Anderson and van Wincoop (2003) establish the theoretical foundation for the multilateral resistance terms that any consistent gravity specification must control for. Baier and Bergstrand (2007) discuss possible omitted-variable bias from traditional specifications that include distance when considering the effect of trade agreements between countries—closer countries with fewer frictions may be more likely to both sign agreements and trade more—and thus opt for the inclusion of directional fixed effects for each importer-exporter pair, accounting for all time-invariant features of bilateral trade. As outlined by Larch et al. (2025), using directional pair fixed effects has become standard practice for properly estimating gravity models, with one issue being that decades-long trade panels may have large time-varying bilateral trade costs, due to shifts in specialization and international relations, that are omitted with a time-invariant fixed effect. On the estimation side, Santos Silva and Tenreyro (2006) establish that Poisson pseudo-maximum likelihood is the appropriate estimator for gravity models given the heteroskedasticity endemic to trade data and the prevalence of zero flows. This paper also builds on the quantitative trade literature that uses estimated elasticities to conduct general-equilibrium counterfactuals. Dekle et al. (2008) develop the exact hat algebra approach that expresses the counterfactual equilibrium in changes, requiring only observed expenditure shares and a trade elasticity rather than the model's unobservable primitives. Arkolakis et al. (2012) show that in a large class of quantitative trade models, welfare changes are summarized by the change in the own-trade share and the trade elasticity; Costinot and Rodríguez-Clare (2014) extend this framework, including to settings with unbalanced trade. The welfare decomposition in this paper inherits the algebra of that extension. The difference here is that the terms-of-trade component is the disruptor's motive of which the model provides insight.

3 The Model

3.1 Environment

Notation. There are $N+1$ countries indexed $n, i, j, k \in \{1, \ldots, N+1\}$. Country 1 is the disruptor, who has the ability to disrupt transportation costs between countries in the disrupted bloc $\mathcal{M} \equiv \{2, \ldots, M\}$, with $3 \leq M \leq N+1$. The number of disrupted countries is $|\mathcal{M}| \equiv m = M - 1$, and $b \equiv N + 1 - m$ is the number of bystanders, the disruptor included. The set of disrupted links is all of the international linkages between bloc countries, $\mathcal{D} \equiv \{(n, i) \in \mathcal{M}^2 : n \neq i\}$, so that $|\mathcal{D}| = m(m-1)$.

Goods and trade costs. Each country is endowed with $y_n > 0$ units of a unique, differentiated national good and $T_n > 0$ units of a freely traded numeraire whose price is normalized to one. The mill price in country i is p_i. There are standard iceberg trade costs $\tau_{ni} \geq 1$ with $\tau_{nn} = 1$ for all n. Delivered prices of a country-i good shipped to country n are thus $p_{ni} = \tau_{ni} p_i$. In order to deliver c_{ni} units of the good, country i must ship $m_{ni} \equiv \tau_{ni} c_{ni}$ units.

Preferences. The representative country-n consumer has quasilinear preferences over consumption of the numeraire x_0^n and a CES aggregate of differentiated goods C_n,

$$U_n = x_0^n + v(C_n), \qquad C_n = \left[\sum_{i=1}^{N+1} \beta_i c_{ni}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \qquad v' > 0,\ v'' < 0,\ \sigma > 1. \tag{1}$$

For much of this paper, I assume the log tier, where $v(C_n) = E \ln C_n$ and $E > 0$. Doing so simplifies much of the analysis without contradicting the nature of my findings. Section 3.8 verifies this numerically. These results are not unique to the Armington structure: Appendix B shows that every result below holds verbatim in an Eaton–Kortum economy, with the trade elasticity θ replacing $\sigma - 1$. Under EK framework the disruptor's gain simply requires it to have ex-ante market access to the disrupted bloc.

3.2 The Consumer's Problem

Country n chooses a feasible quantity of the numeraire and a vector of differentiated good quantities that maximize its utility:

$$\max_{x_0^n \geq 0,\ \{c_{ni}\}} x_0^n + v(C_n) \quad \text{s.t.} \quad x_0^n + \sum_i p_{ni} c_{ni} = I_n \equiv T_n + p_n y_n. \tag{2}$$

Define country n's price index as

$$P_n \equiv \left[\sum_i \beta_i^\sigma p_{ni}^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \tag{3}$$

Conditional demand and expenditure shares for country-i goods in country n are then

$$x_{ni} = p_{ni} c_{ni} = \beta_i^\sigma p_{ni}^{1-\sigma} P_n^\sigma C_n \tag{4}$$

$$s_{ni} \equiv \frac{x_{ni}}{P_n C_n} = \frac{\beta_i^\sigma (\tau_{ni} p_i)^{1-\sigma}}{\sum_k \beta_k^\sigma (\tau_{nk} p_k)^{1-\sigma}}, \qquad \sum_i s_{ni} = 1. \tag{5}$$

Substituting the fact that numeraire consumption is residual income, $x_0^n = I_n - P_n C_n$, into Equation (1) implies that the optimally chosen C_n must solve $\max_{C_n} I_n - P_n C_n + v(C_n)$, which gives the first-order condition

$$v'(C_n) = P_n \quad \Longrightarrow \quad C_n = (v')^{-1}(P_n). \tag{6}$$

I then define the bundle-size elasticity

$$\varepsilon \equiv -\frac{d \ln C_n}{d \ln P_n} = -\frac{P_n}{C_n} \frac{1}{v''(C_n)} \geq 0. \tag{7}$$

Under the log tier, this simplifies to $\varepsilon = 1$, and differentiated-goods spending is the constant $P_n C_n = E$. Combining Equation (4) with $m_{ni} = c_{ni} \tau_{ni}$ gives

$$m_{ni} = \beta_i^\sigma \tau_{ni}^{1-\sigma} p_i^{-\sigma} P_n^\sigma C_n, \qquad \ln m_{ni} = \sigma \ln \beta_i + (1-\sigma) \ln \tau_{ni} - \sigma \ln p_i + \sigma \ln P_n + \ln C_n. \tag{8}$$

Differentiating at fixed trade costs and applying $d \ln C_n = -\varepsilon \, d \ln P_n$ gives

$$d \ln m_{ni} = -\sigma \, d \ln p_i + \gamma \, d \ln P_n, \qquad \gamma \equiv \sigma - \varepsilon. \tag{9}$$

Further, log-differentiating country n's price index with respect to mill price p_j yields

$$\frac{d\ln P_n}{d\ln p_j} = \frac{1}{1-\sigma}\cdot\frac{\beta_j^{\sigma}(1-\sigma)(\tau_{nj}p_j)^{1-\sigma}}{\sum_k \beta_k^{\sigma}(\tau_{nk}p_k)^{1-\sigma}} = s_{nj}. \tag{10}$$

Price changes in goods that hold a higher share in country n's expenditure yield a larger change in its price level. Equations (9)–(10) are the only inputs the comparative statics below require.

3.3 Disruption and Equilibrium

The disruption technology. I now introduce the disruptor's decision. The disruptor chooses an intensity $d \geq 0$ that raises the bilateral trade cost of every link in the disrupted set $\mathcal{D}$ symmetrically, so that the matrix of bilateral trade costs changes according to

$$\tau_{ni}(d) = \min\{\tau_{ni}^0 e^{g(d)}, \bar{\tau}_{ni}\} \quad \forall\ (n,i) \in \mathcal{D}, \tag{11}$$

where $g(0) = 0$, $g' > 0$, and $\bar{\tau}_{ni}$ is the cost of the next-best available route between n and i. Define $\delta' \equiv g'(d) = \partial \ln \tau_{ni}(d)/\partial d > 0$, the common percentage increase in cost on every disrupted link below its cap, and let $\bar{d}$ denote the smallest intensity at which every disrupted link is capped. Three features deserve emphasis. First, the shock is log-symmetric: every disrupted link's cost rises by the same percentage, so a link's absolute cost change depends on its baseline, with higher baseline costs receiving larger absolute increases. Second, none of the disruptor's own links are disrupted; the structure assumes that the disruptor can degrade the corridor without harming its own market access. Third, the bypass caps $\bar{\tau}_{ni}$ bound the achievable damage. In the Red Sea application they correspond to Cape of Good Hope routing. Disruption is also costly: the disruptor pays $\kappa(d)$ in numeraire, with $\kappa(0) = 0$, $\kappa' > 0$, and $\kappa'' > 0$.

Equilibrium under disruption. Given d, goods markets clear when each country's endowment equals gross shipments demanded of its good,

$$y_i = X_i(p;d) \equiv \sum_n m_{ni}, \qquad i = 1,\ldots,N+1, \tag{12}$$

and the numeraire market clears by Walras's law. A *competitive equilibrium given* d is a price

vector $\{p_i(d)\}$ and an allocation $\{c_{ni}, x_0^n\}$ such that every consumer maximizes utility and (12) holds. The numeraire pins absolute prices, so the equilibrium system is determinate, and under the gross-substitutes condition established below the equilibrium is unique. Evaluated at equilibrium, country n's indirect utility under the log tier is

$$u_n(d) = T_n + R_n(d) - E + E \ln E - E \ln P_n(d), \qquad R_n \equiv p_n y_n. \tag{13}$$

The disruptor's problem. The disruptor moves first, internalizing the competitive mapping $d \mapsto p(d)$, and maximizes the objective function

$$V_1(d) = u_1(d) - \lambda\, u_2(d) - \kappa(d), \qquad \lambda \geq 0, \tag{14}$$

where country 2 is a designated target within the bloc and λ weights the harm inflicted upon it. Positive values of λ represent geopolitical competition, whereby the disruptor gains from harming its target; $\lambda = 0$ implies that the disruptor has an exclusively greedy incentive to disrupt, and there is no target. An equilibrium of the full game is a subgame-perfect equilibrium: an intensity d^* maximizing (14) given the competitive mapping, together with the associated competitive equilibrium. There is a clear upper bound on d: for $d > \bar{d}$, every disrupted cost is capped, so prices and utilities are flat in d while the cost of disruption continues to rise. Hence $d > \bar{d}$ implies $V_1'(d) = -\kappa'(d) < 0$, such that $d^* \in [0, \bar{d}]$.

3.4 General-Equilibrium Price Effects

The disruptor's incentive depends on how equilibrium prices respond to its chosen level of disruption. Write excess demand $Z_i(p; d) \equiv X_i(p; d) - y_i$ and totally differentiate the equilibrium conditions $Z = 0$ along the equilibrium path:

$$J\,\dot{p} = -Z_d, \qquad J_{ij} \equiv \frac{\partial Z_i}{\partial p_j}, \quad \dot{p}_j \equiv \frac{dp_j}{dd}, \quad (Z_d)_i \equiv \left.\frac{\partial Z_i}{\partial d}\right|_p. \tag{15}$$

Here Z_d is the direct effect: the excess demand the disruption creates at fixed prices. The matrix J encodes the feedback through which price movements re-clear markets. It is analytically convenient to work with the elasticity Jacobian $\tilde{J}_{ij} \equiv \partial X_i / \partial \ln p_j = p_j J_{ij}$, which shares the sign

pattern of J. Summing the shipment responses (9)–(10) across destinations delivers its entries:

$$\tilde{J}_{ii} = -\sigma X_i + \gamma \sum_n m_{ni} s_{ni}, \qquad \tilde{J}_{ij} = \gamma \sum_n m_{ni} s_{nj} \quad (j \neq i). \tag{16}$$

A price increase in good j hurts good j directly (the $-\sigma$ term) and helps every good through the general-equilibrium price-index channel: each importer's index rises by s_{nj} log points, and bundle demand responds with net elasticity γ.

Lemma 1 (Sign structure)**.** *Suppose $\sigma > 1$ and $\gamma > 0$. Then $\tilde{J}$ has strictly positive off-diagonal entries, strictly negative diagonal entries, and row sums $\sum_j \tilde{J}_{ij} = -\varepsilon X_i \leq 0$. Consequently $-\tilde{J}$ is an M-matrix: $\tilde{J}$ is invertible, $-\tilde{J}^{-1} \geq 0$ entrywise, and the equilibrium is unique.*

The proof is a direct computation using $\sum_j s_{nj} = 1$ and is collected in Appendix A.2. Economically, the lemma is the gross-substitutes property: an increase in the price of a single good raises demand for every other good, and own-market effects dominate cross-market feedback. To obtain tractable closed forms, I now evaluate at a symmetric benchmark.

Assumption 1 (Symmetric benchmark)**.** $y_n = y$, $T_n = T$, and $\beta_i = 1$ for all i; $\tau_{ni} = \tau > 1$ for all $n \neq i$; and the log tier holds. Units of each good are normalized such that the $d = 0$ equilibrium has $p_i = 1$ and $P_n = P$ for all n, i.

Define $\theta \equiv \tau^{1-\sigma} \in (0, 1)$. Benchmark expenditure shares across countries are thus $s_0 = 1/(1 + N\theta)$ for the home good and $s_1 = \theta/(1 + N\theta)$ on each foreign good. Domestic shipments are $a \equiv P^\sigma C$ and each foreign shipment is a constant share of that, θa, so each bilateral trade value is $\bar{x} \equiv \theta a$ and differentiated trade is balanced: $R_n = E$ for every n. Evaluated here, the Jacobian is the circulant

$$\frac{\tilde{J}}{a} = (D - O)\, I + O\, \mathbf{1}\mathbf{1}^\top, \qquad D = -\sigma(1 + N\theta) + \gamma \frac{1 + N\theta^2}{1 + N\theta} < 0, \qquad O = \gamma \frac{\theta[2 + (N-1)\theta]}{1 + N\theta} > 0, \tag{17}$$

with row sum $D + NO = -\varepsilon(1 + N\theta) < 0$. The direct effect of the disruption is a pure reallocation: at fixed prices, bloc members' price indices rise by $(m-1)s_1\delta'$ each, pushing their demand toward every bystander good, while the disrupted links carry the own-cost pullback. Collecting terms (Appendix A.4),

$$-Z_d = (m-1)\,k\,a\,\big(\underbrace{-m,\ldots,-m}_{b\text{ bystanders}},\ \underbrace{b,\ldots,b}_{m\text{ bloc}}\big)^{\top}, \qquad k \equiv \gamma s_1 \theta \delta' = \frac{(\sigma-1)\theta^2}{1+N\theta}\,\delta', \tag{18}$$

whose entries sum to zero: the disruption creates no aggregate excess demand, only a transfer of demand from bloc goods onto bystander goods.

The zero-sum structure of the direct effect simplifies the general-equilibrium solution, because the symmetry of the benchmark collapses the entire $(N{+}1)$-dimensional system down to two equations. Symmetry implies that every bystander good's price responds identically including that of the disruptor—call the common response $\hat{p}_B \equiv \dot{p}_i$ for the b bystanders, the disruptor included—and every bloc good's price responds identically, call it $\hat{q} \equiv \dot{p}_i$ for the m bloc members. It therefore suffices to write down one representative bystander row and one representative bloc row of $\tilde{J}\dot{p} = -Z_d$. Dividing by a and reading the circulant (17), a bystander row carries the diagonal D on itself and O on each of the other $b-1$ bystanders and m bloc goods, and a bloc row analogously—the system reduces to

$$\big[D + (b-1)O\big]\,\hat{p}_B + m\,O\,\hat{q} = -m(m-1)\,k, \tag{19}$$

$$b\,O\,\hat{p}_B + \big[D + (m-1)O\big]\,\hat{q} = b(m-1)\,k. \tag{20}$$

The right-hand sides are the two distinct entries of $-Z_d/a$ from (18): every bystander good absorbs $-m(m-1)k$ of reallocated demand and every bloc good sheds $b(m-1)k$. The determinant of this 2×2 system, using $m+b=N+1$, factors cleanly:

$$\Delta = \big[D + (b-1)O\big]\big[D + (m-1)O\big] - mb\,O^2 = (D-O)\,(D+NO) > 0, \tag{21}$$

the product of two strictly negative terms. Solving (19)–(20) by Cramer's rule, the common factor $D+NO$ cancels and the solution collapses to a single division by $D-O$:

$$\dot{p} = \frac{-Z_d}{a\,(D-O)}. \tag{22}$$

The direct effect (18) supplies the entire cross-sectional incidence—who gains, who loses, and

in what proportion—while general equilibrium contributes exactly one scalar, $1/|D-O|$, which uniformly dampens all price movements: as bloc prices fall, cheap bloc goods pull demand back from bystander goods, moderating every response. Reading off (22) delivers the paper's central analytical result.

Proposition 1 (Terms-of-trade effects of disruption)**.** *Under Assumption 1, for $d < \bar{d}$,*

$$\hat{p}_1 = \hat{p}_B = \frac{m(m-1)\,k}{|D-O|} > 0, \qquad \hat{p}_i = \hat{q} = -\frac{b\,(m-1)\,k}{|D-O|} < 0 \quad (i \in \mathcal{M}), \qquad \frac{\hat{p}_B}{\hat{q}} = -\frac{m}{b}, \tag{23}$$

where $\hat{p}_j \equiv d\ln p_j/dd$ and $\hat{p}_B$ is common to all bystanders. Every bystander's terms of trade improve; every bloc member's deteriorate.

Each group's price response is proportional to the opposite group's size, because the reallocated demand (respectively, the displaced supply) is spread across it. A corollary of the $m(m-1)$ scaling deserves emphasis, as it disciplines which disruptions are worthwhile.

Corollary 1 (Corridors, not pairs)**.** *As the world grows arbitrarily large ($N \to \infty$, θ fixed) for a fixed bloc size m, $\hat{p}_B = O(N^{-2})$, such that disrupting a constant-size bloc yields a vanishing impact. For a bloc that is a fixed fraction of the world, $m = \alpha N$ with $\alpha \in (0,1)$,*

$$\hat{p}_B \;\longrightarrow\; \frac{\alpha^2(\sigma-1)}{\sigma}\,\delta' \;>\; 0, \tag{24}$$

so corridor-scale disruption remains first order in an arbitrarily large world.

Corollary 1 implies that the incentive to disrupt exists only when the disrupted set carries a meaningful share of world trade. It explains why observed disruptions target major global chokepoints, like the Bab el-Mandeb Strait, rather than individual trade relationships.

3.5 Welfare and the Incentive to Disrupt

Differentiating (13) along the equilibrium path and applying the envelope theorem (the induced variation in C_n cancels out because $v'(C_n) = P_n$),

$$u_n'(d) = R_n\hat{p}_n - E\hat{P}_n = \underbrace{\mathrm{NX}_n\,\hat{p}_n}_{\text{net-export ToT}} + \underbrace{E\,(\hat{p}_n - \hat{P}_n)}_{\text{real income per unit}}, \qquad \mathrm{NX}_n \equiv R_n - E, \tag{25}$$

where $\hat{P}_n$ is the total log change in n's price index, inclusive of the direct corridor-cost term $(m-1)s_1\delta'$ for bloc members. At the benchmark, trade is balanced ($\mathrm{NX}_n = 0$) and the closed forms are clear. For the disruptor, and identically for any bystander,

$$u_1'(d) = \bar{x}\; m\left(\hat{p}_B - \hat{q}\right) = \bar{x}\,\frac{m(m-1)(N+1)\,k}{|D-O|} \;>\; 0. \tag{26}$$

The disruptor collects the full price *wedge* $\hat{p}_B - \hat{q}$ on each of its m bilateral relationships with bloc members—its export price to them rose while its import price from them fell—and nothing on bystander-bystander relationships, whose prices co-move. The marginal payoff inherits the quadratic bloc scaling discussed earlier. For a bloc member $n \in \mathcal{M}$,

$$u_n'(d) = \bar{x}\Big[\, b\left(\hat{q} - \hat{p}_B\right) - (m-1)\,\delta' \Big] \;<\; 0, \tag{27}$$

both terms strictly negative: a terms-of-trade loss on its b bystander relationships, plus the direct real-income cost of the costlier corridor on its $m-1$ intra-bloc import relationships.

Proposition 2 (Disruption without animosity)**.** *At the symmetric benchmark, the marginal benefit of disruption is strictly positive for every $\lambda \geq 0$ and $d < \bar{d}$. In particular, at $\lambda = 0$: if $\kappa'(0) < u_1'(0)$, then $d^* > 0$—a purely self-interested disruptor strictly prefers positive disruption, driven entirely by the terms-of-trade realignment in* (26)*.*

The interior optimum solves the first-order condition

$$\underbrace{u_1'(d)}_{\text{economic gain}} + \lambda\;\underbrace{\big(-u_2'(d)\big)}_{\text{geopolitical gain}} \;=\; \kappa'(d), \qquad d^* = \min\{d^{\mathrm{FOC}}, \bar{d}\}, \tag{28}$$

and since $-u_2' > 0$ by (27), the geopolitical weight unambiguously raises d^*: spite reinforces greed's incentives as long as $d^* < \bar{d}$. Two other features are worth noting. First, every bystander enjoys the gain (26) while only the disruptor pays $\kappa(d)$, so disruption is a public good among bystanders, and the disruptor's distinguishing feature lies exclusively in its geographic capability

to disrupt and any geopolitical incentive to do so. Second, the comparative static of $u_1'(0)$ in σ is non-monotone: a higher elasticity strengthens substitution per unit of shock but shrinks the trade shares ($\theta = \tau^{1-\sigma}$) through which the shock operates. I document this, and the robustness of the mechanism to the maintained assumptions, numerically in Section 3.8.

3.6 Retaliation and Deterrence

Disruption may invite a retaliatory response. A response that imposes a direct cost on the disruptor—strikes or sanctions on its assets—is accounted for by $\kappa(d)$, whose convexity implies that larger disruptions provoke proportionally larger responses. However, a response that raises the disruptor's own *trade costs*—a blockade, an embargo, carriers refusing its ports—fundamentally differs, as it alters the equilibrium the disruptor seeks to exploit. I model retaliation as a rule to which the bloc commits before the disruptor moves: for every unit of disruption, the bloc raises the cost of every link between itself and the disruptor by a fraction $r \geq 0$ of the disruption itself,

$$\frac{\partial \ln \tau_{n1}}{\partial d} = \frac{\partial \ln \tau_{1n}}{\partial d} = r\,\delta' \quad \forall\, n \in \mathcal{M}, \tag{29}$$

with $r = 0$ reflecting the no-retaliation model above and $r = 1$ a one-for-one response that exposes the disruptor to the same cost increase it imposes. The timing is that of a deterrence game and motivates $r > 0$ as a possible subgame perfect Nash equilibrium: the bloc announces r, the disruptor chooses d taking r as given, and the competitive equilibrium follows. Because the direct effect of the combined shock still sums to zero across goods, general equilibrium again enters only through the scalar $1/|D - O|$ (Appendix A.9). The closed forms thus extend directly.

Proposition 3 (Retaliation and deterrence)**.** *Under Assumption 1 and the retaliation rule (29), for $d < \bar{d}$:*

(i) Prices respond according to

$$\hat{p}_1 = \frac{m(m-1) - m\,r\,(N-1)}{|D-O|}\,k, \qquad \hat{p}_B = \frac{m(m-1) + 2\,m\,r}{|D-O|}\,k, \qquad \hat{q} = -\frac{b(m-1) + r\,(b-m)}{|D-O|}\,k, \tag{30}$$

where $\hat{p}_B$ now denotes the common response of the $b-1$ bystanders other than the disruptor.

(ii) The disruptor's marginal gain is linear and strictly decreasing in r,

$$u_1'(d;r) = u_1'(d;0)\Big[1 - r\,\frac{N-1}{m-1}\Big] - \bar{x}\,m\,r\,\delta', \tag{31}$$

and is nonpositive if and only if $r \geq \bar{r}$, where

$$\bar{r} \;\equiv\; \frac{u_1'(d;0)}{u_1'(d;0)\,\dfrac{N-1}{m-1} + \bar{x}\,m\,\delta'} \;<\; \frac{m-1}{N-1} \;<\; 1. \tag{32}$$

(iii) For a bloc member $n \in \mathcal{M}$,

$$u_n'(d;r) = u_n'(d;0) - \bar{x}\,r\Big[\delta' - (N+1)(m-b)\frac{k}{|D-O|}\Big], \tag{33}$$

so retaliation deepens the bloc's own loss whenever $m \leq b$.

The proof is in Appendix A.9; Appendix B.5 verifies every expression numerically. Retaliation erodes the disruptor's gain through two channels, the two terms of (31). First, the disruptor's goods are pushed out of bloc markets just as bloc goods are, so the demand displaced from the corridor lands on the other bystanders rather than on the disruptor: its price rises by less than theirs, by (30), and at $r = (m-1)/(N-1)$ it does not rise at all. Second, the disruptor now pays the higher cost on its m import links from the bloc, the term $\bar{x}\,m\,r\,\delta'$, which is the direct real-income loss a bloc member suffers in (27). Because both channels operate, the threshold $\bar{r}$ at which the economic motive vanishes is strictly below $(m-1)/(N-1)$: a bloc that is a minority of the world can neutralize greed with a response well short of proportional. In the numerical benchmark of Appendix B ($N+1 = 12$, $m = 6$), $\bar{r} = 0.15$.

Retaliation discriminates between the two motives. With $\lambda = 0$, a rule $r \geq \bar{r}$ makes the marginal benefit of disruption nonpositive at $d = 0$, so a purely self-interested disruptor does not disrupt at all. With $\lambda > 0$, the geopolitical term $-u_2'$ in (28) is *larger* under retaliation whenever $m \leq b$, by (33): a response that hurts the retaliator is, to a spiteful disruptor, part of the harm it set out to inflict. Deterrence by trade-cost retaliation therefore works on greed and not on grievance, and an observed disruption that persists in the face of it is evidence of the latter. Two further implications follow. For the inversion of Section 5, any anticipated retaliation lowers $u_1'(d^*)$ and with it the

greed-only bound $\bar{\kappa}'$, so the implied geopolitical weight in (54) can only rise; the estimates reported there, which set $r = 0$, are conservative in this respect. And because retaliation is costly to the retaliators at the margin when $m \leq b$, the rule is not a best response once d has been chosen: deterrence requires commitment, which is why it is a question of institutions and coalitions rather than of incentives alone.

3.7 An Augmented Gravity System and ACR Decomposition

Structural gravity with an endogenous resistance block. The equilibrium admits a structural gravity representation. Defining exporter capability $S_i \equiv \beta_i^{\sigma} p_i^{1-\sigma}$ and inward multilateral resistance $\Phi_n \equiv \sum_k S_k \tau_{nk}^{1-\sigma} = P_n^{1-\sigma}$, bilateral trade values satisfy

$$x_{ni} = S_i \frac{E_n}{\Phi_n} \tau_{ni}^{1-\sigma}, \qquad \sum_n x_{ni} = R_i, \tag{34}$$

the canonical form with trade elasticity $\sigma - 1$. Under the log tier the importer term is simply the constant E. The departure from the canon is that the resistances on $\mathcal{D}$ are no longer taken as given. The full model is the *augmented gravity system*: the gravity block (34), closed by the chokepoint optimality condition (28) that determines $\tau_{ni}(d^*)$ for all disrupted links. This structure dictates the empirical division of labor pursued below: the gravity block is estimated from ordinary, non-strategic variation in trade costs, while (28) is *inverted*—not estimated—to recover the strategic primitives (κ', λ) that rationalize an observed disruption.

An augmented ACR decomposition. Welfare changes admit a sufficient-statistic representation. Since $s_{nn} = S_n/\Phi_n$, log-differentiating gives $\hat{p}_n - \hat{P}_n = -\hat{s}_{nn}/(\sigma - 1)$, so (25) becomes

$$u_n'(d) = \underbrace{\mathrm{NX}_n \, \hat{p}_n}_{\text{strategic terms of trade}} - \underbrace{\frac{E}{\sigma - 1} \hat{s}_{nn}}_{\text{ACR own-share term}} \tag{35}$$

The second term is the canonical own-trade-share statistic of Arkolakis et al. (2012); the first is a terms-of-trade rent proportional to net exports, similar to other extensions to unbalanced trade. What is new is the term's role: here it is not simply an accounting correction but the disruptor's *objective*, with model-predicted sign, incidence across bystanders and bloc members,

and quadratic scaling in corridor size. At the balanced-trade benchmark, the first term vanishes and the disruptor's gain loads entirely on the own-share channel ($\hat{s}_{11} < 0$: its export price rises relative to its basket). Away from balanced trade, in the empirically relevant case, both channels operate; in the quantitative analysis of Section 4.4, the own-share channel accounts for the bulk of every country's welfare change, including the disruptor's.

3.8 Numerical Robustness

The closed-form results above rest on a series of assumptions made in the economic environment and, especially, in the symmetric benchmark used to obtain them. This subsection establishes that the paper's central result—the terms-of-trade incentive to disrupt—is robust to these assumptions, while identifying the boundaries within which the sharper, secondary predictions hold. Because relaxing the assumptions renders the analytics intractable, I characterize the model numerically: I solve for equilibrium prices and welfare under each relaxation and compare them against the benchmark model's predictions, holding all other features constant. For simplicity, I first focus on the three-country setting, where indices 1, 2, and 3 correspond to the disruptor and the two members of the disrupted bloc $\mathcal{M} = \{2, 3\}$, with country 2 being the designated target. In this setting, the disruptor is the sole bystander to the shock. Unless stated otherwise, I fix $\sigma = 4$, $\tau_0 = 1.5$ for all $i \neq j$ (with $\tau = 1$ for $i = j$), and $\bar{\tau} = 1.4\,\tau_0$, such that the bypass route is 40 percent more costly than the disrupted route—a rough approximation for the Cape of Good Hope's distance premium over the Bab el-Mandeb passage. Endowments are calibrated so that $p_i = 1$ is the no-disruption equilibrium under each specification, ensuring that cross-specification differences reflect responses to disruption rather than differences in the initial allocation. Appendix C confirms that the numerical solver reproduces the closed forms of Proposition 1 to seven digits, so the deviations reported below reflect genuine economic effects rather than numerical error.

Endowment asymmetry. This exercise relaxes the size symmetry assumed in Assumption 1 and allows for meaningful relative size disparities between the three countries. It scales each country's endowment from 0.5× to 2× and computes price changes ($d \ln p / dd$ at $d = 0$) and welfare changes at $d = 0.25$ along that interval. Figure 2 displays the results: each column corresponds to resizing a different country, the top row displays price changes, and the bottom row displays welfare

changes. The crucial welfare result of (26) and (27) is robust: across all relative size combinations, the disruptor gains ($\Delta u > 0$) while bloc members lose ($\Delta u < 0$). Price signs, however, are not universal. At moderate asymmetries (roughly 0.7× to 1.3×) the predictions hold, but for a very small bloc member the price response turns positive, as shown by the initial positive target price change in the upper center panel. A small country's price index is fragile, so a sharply contracting bundle demand forces home prices up to clear the market; this is a terms-of-trade improvement driven by consumption collapse rather than genuine gain. Important for the Red Sea application, the price effects in the first column are all well behaved: as the disruptor becomes relatively small, the price effect of its disruption shrinks but retains its sign. Since Iran is relatively small in the world economy, this is the most empirically relevant configuration.

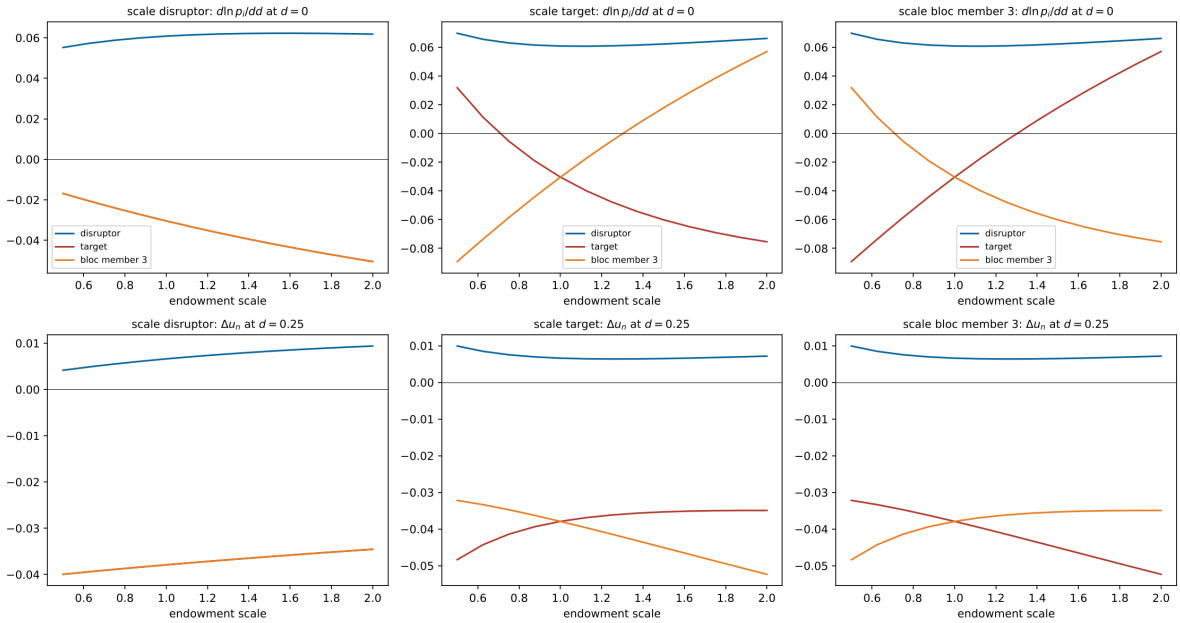


Figure 2: Price and welfare effects under endowment asymmetry. Price responses (top row) and welfare changes (bottom row) as each country's endowment is scaled from 0.5× to 2×, holding all else constant. Columns correspond to scaling the disruptor, the target, and the second bloc member respectively. Columns two and three are symmetric because, in the three-country setting, the target and the second bloc member are mathematically identical: the disruptor raises the cost of the trade link between them symmetrically.

Utility specification. A major simplifying assumption is the utility form chosen in (1). A freely traded numeraire and a log upper tier together shut down income effects and fix the bundle-size elasticity at one, so that price-level increases are exactly offset by bundle-quantity decreases. This exercise generalizes the upper tier and, separately, removes the numeraire altogether by moving to

an Armington model. I first allow for three alternative upper tiers of the form

$$v(C) = A\frac{C^{\rho}}{\rho}, \quad \rho < 1,\ \rho \neq 0, \tag{36}$$

with $\rho \in \{-1, 0.5, 0.8\}$, so that the bundle-size elasticity is $\varepsilon \in \{0.5, 2, 5\}$ and the net index feedback $\gamma = \sigma - \varepsilon$ is $\gamma \in \{3.5, 2, -1\}$, respectively. The benchmark log upper tier corresponds to $\rho \to 0$, where $\varepsilon = 1$ and $\gamma = 3$. Note that $\gamma < 0$ implies that differentiated goods are gross complements.

Holding the calibration fixed at the benchmark and re-benchmarking each specification so that all five share the same $d = 0$ allocation, I trace out the disruption path $d \in [0, \ln 1.4]$ under each upper tier and, for comparison, under a pure Armington model in which all income is spent on the differentiated bundle ($C_n = I_n/P_n$, $I_n = p_n y_n$) and income effects are at work. Because the Armington setup is homothetic, the price system is determined only up to scale. I impose $p_1 = 1$ as the normalization and solve the remaining market-clearing conditions, so every other price reported for that specification is implicitly a price relative to that of the disruptor.

Figure 3 plots, for all five specifications, the disruptor's price path, the target's price path, and the disruptor's welfare gain. Two results stand out. First, the level of *nominal* prices is not robust to the upper tier, and the direction of the departure is pinned down by whether $\varepsilon < 1$. The result that the shock's direct effect sums to zero across countries—so that general equilibrium collapses to the single relative-price scalar $1/|D - O|$—relies on differentiated spending $P_n C_n$ being constant, a property of the log tier alone. Away from that particular case, spending on the bundle moves with the price index. With inelastic bundle demand ($\varepsilon = 0.5$), a costlier corridor increases total differentiated spending, so every nominal price is pushed upward, including that of the target, whose price path rises even though it is the country being hurt. With elastic bundle demand the same mechanism runs in reverse and reinforces the relative-price decline. Nominal price predictions are therefore an assumption-dependent feature of the log tier and should not be over-interpreted. By contrast, and importantly for my results, relative price changes are robust: the target's price relative to the disruptor's falls across all five cases.

Second, the welfare result is robust to the utility specification. The disruptor's welfare gain is positive across every specification, including the Armington economy in which income effects

are fully operative. The welfare mechanism—the augmented-ACR own-trade-share channel—does not depend on the curvature of the upper tier or on the absence of income effects; it depends only on the target's own trade share falling, which happens regardless of ε. Taken together, these two results clarify what the log-tier and numeraire assumptions buy analytically: tractability and a clean characterization of nominal price levels, without which the qualitative predictions nonetheless survive.

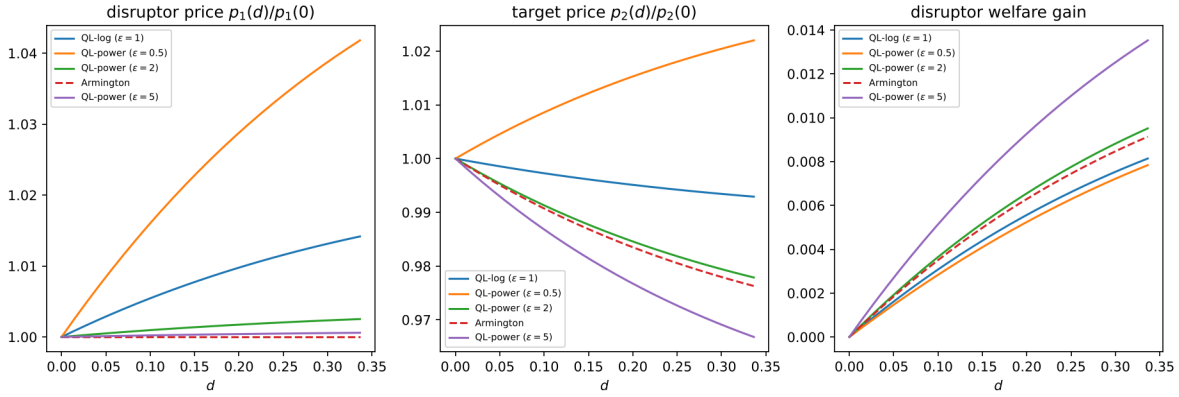


Figure 3: Robustness to the utility specification. Price paths (left and center panels, normalized to $d = 0$) and the disruptor's welfare gain (right panel) across disruption intensity $d \in [0, \ln 1.4]$ for five upper-tier specifications: the log benchmark ($\varepsilon = 1$), two power specifications bracketing it ($\varepsilon = 0.5$ and $\varepsilon = 2$), a gross-complements case ($\varepsilon = 5$), and Armington (dashed; income effects active, prices normalized to the disruptor's own price). Nominal price paths diverge in sign depending on whether ε is greater or less than one, while relative prices and the welfare ranking match the benchmark across all specifications.

Realistic exposure. The model assumes that exposure to the shock is purely intra-bloc: only trade flows between two bloc members are disrupted—no flow involving a non-member is affected. This is clearly not true of realistic transportation-network disruptions like the Red Sea crisis, where exposure at the country level is far more dispersed, with nearly every country having at least one trade route subject to the shock. Using the earlier example, the United States' trade with Mexico is unexposed, but its trade with India is exposed. Treatment in the data is therefore a continuum rather than a binary measure.[5] Modeling such treatment analytically would be difficult and intractable, but testing whether it changes the model's implications matters—especially for the prediction that bystanders benefit. I therefore replace the bloc structure with the empirically

[5]Strictly speaking there are no bystanders in the real scenario, since every country other than the disruptor experiences some increase in trade costs. There is, however, a continuum of treatment, with some countries highly exposed and others barely so.

measured dyadic exposure of Section 4.1, scaling each dyad's trade cost by its measured disparity and holding endowments symmetric so that the exercise isolates changes in the treatment structure.

Figure 4 reports the results: every qualitative prediction survives. The disruptor still gains, earning a welfare increase of roughly three percent of differentiated spending despite facing no direct cost shock of its own. Welfare is monotonically decreasing in exposure: the rank correlation between a country's mean log distance increase and its welfare change is -1.00, and every country above median exposure loses while roughly three quarters of those below it gain. The theory's sharp division between bystanders who gain via free-ridership and directly affected bloc members who suffer therefore reappears as a continuous gradient, with the sharp division now occurring below median exposure rather than at an observable bloc boundary. Geographic insulation, not the act of disrupting, is what earns the terms-of-trade gain. The disruptor's only distinction is that it pays $\kappa(d)$ that actually elicits the disruption.

The bloc abstraction also survives as a broad description of incidence. Constructing a matched clean-bloc economy whose uniform within-bloc shock delivers the same aggregate log-cost increase as the measured exposure, the two economies agree on the sign of the welfare change for 226 of 235 countries, and their welfare changes have a rank correlation of 0.86. The nine exceptions are low-exposure economies, among them Greenland, New Zealand, and Mauritania, that are bystanders in the bloc economy but lose slightly under measured exposure, none by more than about one percent of differentiated spending. The bloc structure of Section 3 therefore recovers the pattern of who wins and who loses; the continuous version refines the magnitudes of the winners and losers. The counterfactual exercises of Section 4.4 adopt this specification.

4 Empirical Analysis

In this section, I introduce the data and apply the model developed in Section 3 to quantify the effects of the Red Sea crisis. Section 4.1 describes the data and uses geospatial transportation network data to construct a novel treatment variable. Section 4.2 estimates the reduced-form effect of the crisis and relates the estimates to the model's objects. Section 4.4 computes the general-equilibrium welfare incidence of the disruption under three trade-cost shock specifications.

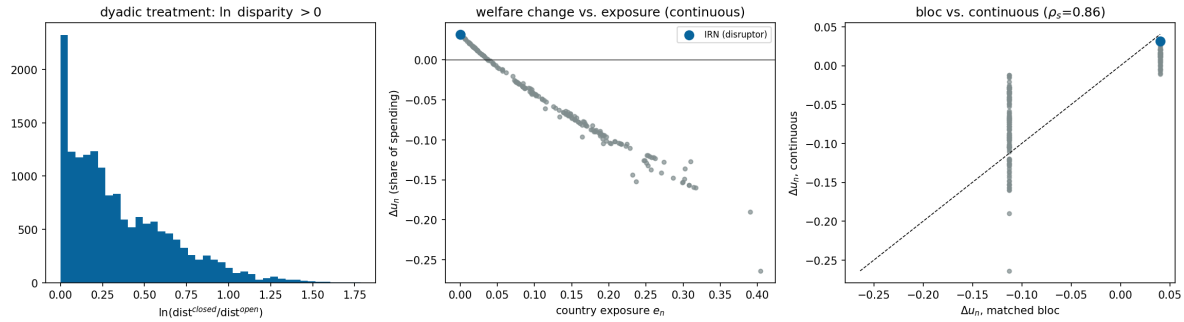


Figure 4: Welfare incidence under empirically measured exposure. The compact bloc structure of the theory is replaced with the dyad-specific rerouting exposure measured in Section 4.1, holding endowments symmetric. The left panel shows the distribution of positive dyadic log-distance increases; the center panel plots each country's welfare change against its exposure, with the disruptor highlighted; the right panel compares welfare changes under the continuous exposure structure to those under a matched clean bloc calibrated to the same aggregate shock, with the 45-degree line superimposed.

4.1 Data and the Red Sea Treatment

The Red Sea crisis entailed a large number of geospatially clustered attacks on shipping vessels around the Bab el-Mandeb Strait beginning in late October 2023. A country pair's ex-ante reliance on the strait implicitly measures its exposure to the attacks. Country pairs that relied heavily on the strait with poor alternative routes were more exposed to the crisis, while those whose routes completely avoid it were not exposed at all. Since a country pair's dependence on the strait was fixed by longstanding geography, it is plausibly exogenous to the attacks and contemporaneous shocks to trade demand, conditional on the fixed effects described in Section 4.2. On the other hand, targeting intensity among affected country pairs was not geographic: in line with the Houthis' aforementioned declaration, the data confirm that Western-linked shipping bore the brunt of the campaign. Such heterogeneity is a source of the differential pass-through documented and explored below.

Trade data. Partial-equilibrium estimation uses monthly, importer-reported bilateral trade values from UN Comtrade over the seventeen-month window surrounding the onset of the attacks. The general-equilibrium counterfactuals require baseline expenditure shares inclusive of internal trade, which monthly Comtrade does not report. For those, I use the 2022 annual cross-section of the International Trade and Production Database for Simulation (Borchert et al., 2024), the last full pre-crisis year. The frequency difference is deliberate: the partial estimates require monthly

variation to identify the dynamics of the response, whereas the counterfactual is a comparative-statics exercise over a pre-crisis baseline, for which annual data are the appropriate frequency and avoid problematically injecting high-frequency noise into the structural shares.

Why standard distance measures fail. In the gravity literature, bilateral distance is conventionally measured as the sum of city-to-city distances between two countries, weighted by the product of each city's population share, so that the result approximates the expected distance between two randomly chosen residents. It is generally referred to as population-weighted distance. The city-to-city distances fed into population-weighted distance are great-circle, or as-the-crow-flies. In this regard, traditional population-weighted distance is critically restricted for my analysis: in measuring distance solely from coordinates, it is unresponsive to changes in transportation infrastructure. For example, a great-circle arc, and its corresponding length, between Mumbai and London is identical regardless of the presence of the Suez Canal. It is thus virtually impossible to assess the impact of transportation infrastructure using a distance variable that eschews it. Additionally, great-circle distance systematically understates the true distance of freight, since arcs cut freely across realistic transportation impasses such as landmasses, ice, and coastlines that block vessels. Across all country pairs in my sample, the ratio of network-routed to great-circle distance averages 1.57 and never falls below 1.01: every pair's true freight path is, unsurprisingly, longer than the theoretical straight line connecting them.[6]

Constructing logistical distance. I thus replace great-circle distance with logistical distance: the length of the least-cost freight route between two cities over a global transportation network. The measure is developed in Giangiulio et al. (2025). I summarize its construction here. The initial network combines a global highway graph with a global maritime shipping-lane graph, comprising 560,282 road nodes and 11,062 maritime nodes. I then append the graph with 6,239 cities spanning 235 countries that are included in the population data, connecting each to its nearest network node. I also add 50 port nodes that bridge the road and maritime layers where coastal cities would otherwise lack maritime access. The resulting network contains 578,298 nodes and 852,167 edges. Least-cost route calculation is implemented by Dijkstra's algorithm over a cost-weighted version

[6]The aforementioned understatement is proportionally largest for relatively *close* pairs, where a strait, mountain range, or lack of direct highway route leads to more circuitous paths.

of this network, in which per-kilometer maritime movement is cheap, road movement is penalized, and port use (onboarding and offloading) carries a fixed cost, reflecting the difference in freight costs across modes. These cost penalties determine which route is selected: the dyadic distance reported below is the true kilometer length of that route. Logistical distance thus measures the distance along a realistic freight path.

Measuring exposure to the closure. Since logistical distance is constructed from a transportation-informed network, it *is* responsive to infrastructure changes. Accordingly, to model the effect of the Bab el-Mandeb Strait's closure, I compute the full set of 38,925,121 city-pair routes twice: once over the complete and unaltered network and once after severing the Strait, which essentially detours any route that relied on it around the Cape of Good Hope. Aggregating each to the country level with population weights yields two distances per country pair, $\text{dist}_{ni}^{\text{open}}$ and $\text{dist}_{ni}^{\text{closed}}$, over a total of 55,085 pairs. The treatment intensity is simply their log difference,

$$T_{ni} \equiv \sum_{c \in \mathcal{C}_n} \sum_{c' \in \mathcal{C}_i} w_{cc'} \left[\ln \text{dist}_{cc'}^{\text{closed}} - \ln \text{dist}_{cc'}^{\text{open}} \right] \geq 0, \tag{37}$$

computed at the city-pair level before aggregation, such that a country pair whose least-cost route never ex-ante transits the strait is coded as unexposed. To illustrate, the least-cost freight route from Mumbai to London (one city pair between the United Kingdom and India) runs 11,751 kilometers through the Red Sea and Suez; with Bab el-Mandeb severed it runs 20,135 kilometers around southern Africa, an increase of 71 percent. Figure 1 in the introduction plots both. By contrast, any city-to-city route between Mexico and the United States is identical under both scenarios, such that T_{ni} is exactly zero.

Properties of the measure. Of the 55,085 country pairs, 16,692 (30 percent) experience distance increases when the strait closes. Among these, the median detour is 29 percent and the ninetieth percentile 119 percent. The maximum experiences a nearly six-fold increase in distance. Of the exposed pairs, 8,237 (15 percent of all pairs) lengthen by at least 30 percent. As explained below, this substantially exposed subset defines the treatment group. Exposure is symmetric by

construction, so a pair's importer and exporter experience the same proportional detour.[7] By proportional distance gain, the most exposed dyads are Eritrea–Oman (+495 percent), the United Arab Emirates–Eritrea (+481 percent), and Egypt–Djibouti (+423 percent). The most exposed countries, aggregated uniformly across trade partners, are Saudi Arabia, Djibouti, Eritrea, Oman, and the United Arab Emirates, as they are clustered around the Strait and thus require large proportional detours for many trading routes. The least exposed are the economies of South America and Central and West Africa, whose trade routes are less likely to require passage through the Bab el-Mandeb and, in the event that they ex-ante do, require only relatively short detours. Importantly, none of this geography is imposed: it is simply the product of least-cost routing over a network that is otherwise naive to the Red Sea crisis.

I employ two external checks on the measure, the first of which also selects the treatment threshold. First, it is important to note that a route that lengthens when the strait closes is not necessarily a route that ex-ante transited it: the Bab el-Mandeb funnels most of its traffic through the Suez Canal, which is associated with costly fees and queue times that go unobserved. Accordingly, a pair whose Cape of Good Hope route is only marginally longer than its Red Sea route plausibly took the Cape before the crisis and is thus unexposed to the shock under my definition. Trade value intuitively illustrates this. Weighting dyads by their bilateral trade, every pair with any distance increase accounts for 28 percent of world trade, more than twice the 10 to 12 percent that independent estimates attribute to the Suez corridor: a naive definition of exposure would therefore assign treatment to far more trade than ever passed through the strait. Restricting exposure to dyads whose routes lengthen by at least 30 percent brings the treated share to 11.1 percent of bilateral trade, squarely within the independent range.[8] I thus set the threshold to conform with this external benchmark and use it to delineate treatment throughout the paper. Second, the treatment is essentially invariant to the relative freight costs used to select routes. Doubling the penalty on road movement leaves the correlation of T_{ni} at 0.998 and the

[7] A small number of pairs (748, or 1.4 percent, accounting for 0.5 percent of trade value) register a measured distance decrease. This is a consequence of recording true-kilometer distance along the cost-minimizing route: when the strait closes, the newly optimal route can occasionally be shorter in kilometers while more expensive in freight cost, usually due to more costly land-route rerouting. These are set to zero exposure throughout.

[8] The treated share falls monotonically in the threshold: 28 percent of trade at any increase, 15.6 percent at 20 percent, and 11.1 percent at 30 percent. The 30 percent cut-off is the one at which the measure reproduces the corridor's independently estimated share, and Section 4.2 shows that the estimates are stable across cut-offs from 0 to 50 percent.

rank correlation of country-level exposure at 0.999; halving it leaves them at 0.956 and 0.992. This invariance is expected: intercontinental pairs route to the nearest port and along the least-cost sea lane under any cost structure that makes maritime freight cheaper than road, and it is exactly these pairs whose distance responds to closure of the strait. The cost weights bind only on short routes with viable land substitutes, which by construction carry no Red Sea exposure.

4.2 Partial-Equilibrium Estimates

The structure of the Red Sea crisis lends itself well to identifying the impact of an armed trade-cost disruption on trade flows. Because exposure to the shock is a product of ex-ante reliance on the Red Sea corridor, determined by longstanding geography, it is plausibly orthogonal to contemporaneous political and economic conditions at the dyad level. I thus treat it as pseudo-exogenous to country pairs. Accordingly, as a proof of concept, I first estimate a simple difference-in-differences design within a saturated structural gravity specification. My simplest specification takes a seventeen-month window around the onset of the attacks in November 2023, allowing roughly eight months each of pre- and post-treatment observation once the October transition month is excluded. I estimate Equation (38) by Poisson pseudo-maximum likelihood (Santos Silva and Tenreyro, 2006), which accommodates zero trade flows and is robust to the heteroskedasticity endemic to trade data:

$$X_{ni,t} = \exp\left[\chi_{n,t} + \chi_{i,t} + \mu_{ni} + \beta\left(\text{Post}_t \times \text{Exposed}_{ni}\right)\right] \nu_{ni,t}. \tag{38}$$

Throughout, n indexes the importer and i the exporter, as in Section 3, so that $X_{ni,t}$ is the value of goods shipped from i to n in month t. Exposed_{ni} is simply a binary variable equal to one for dyads whose logistical distance rises by at least 30 percent when the strait closes, the set whose trade share matches independent estimates of the Suez corridor's (Section 4.1). Dyads with smaller detours are pooled with unexposed dyads in the control group: for such pairs, the Cape of Good Hope route is only marginally longer, so, given canal tolls and transit delays, much of their trade would plausibly have avoided the canal regardless and is unlikely to have been meaningfully treated. The importer-time and exporter-time effects ($\chi_{n,t}$, $\chi_{i,t}$) control for all country-level variation over time: demand and supply shocks, aggregate price levels, and structural gravity's multilateral resistance terms (Anderson and van Wincoop, 2003). The directional bilateral pair fixed effects (μ_{ni}) control

for all time-invariant determinants of trade between countries, including historical relationships, cultural similarity, patterns of specialization, and active trade agreements (Baier and Bergstrand, 2007; Larch et al., 2025).[9] The coefficient β therefore isolates the differential, time-varying effect of the disruption on exposed dyads, identified from variation *within* importer-time and exporter-time cells across country pairs with different geographic exposure. All specifications cluster standard errors by directional pair, allowing for arbitrary serial correlation within a dyad (Larch et al., 2025). Because the attacks began in mid-October 2023, I designate November 2023 as the first post-treatment period and treat October as a transition month, omitted from the specifications below.

Exploiting continuous exposure. My next specification better exploits the continuous nature of the exposure measure. Even among country pairs that ex-ante relied on the Red Sea, exposure to the crisis is heterogeneous: some experience close to a 30 percent increase in distance when the strait closes, while others face significantly more material detours. In the model, this corresponds to how close a pair's bypass trade cost is to its original cost. To capture this heterogeneity, I replace the binary treatment in Equation (38) with a time-varying logistical distance, $\ln \text{dist}_{ni,t}$, whose variation comes from an exposed pair's distance switching from its open-strait to its closed-strait value from November 2023 onward. For pairs below the 30 percent threshold this variable is held time-invariant and is therefore perfectly absorbed by the directional pair fixed effects in Equation (39):

$$X_{ni,t} = \exp\left[\chi_{n,t} + \chi_{i,t} + \mu_{ni} + \beta \ln \text{dist}_{ni,t}\right] \nu_{ni,t}. \tag{39}$$

The coefficient β is now a distance elasticity of trade, identified off the closure-induced change in logistical distance among exposed pairs.

Western heterogeneity. My final specification allows the effect to differ across countries. It is reasonable to expect that Western countries responded more strongly to the attacks, both because

[9] Because of the short sample period, the pair fixed effect is collinear with any trade agreement not enacted or terminated during the window, hence the exclusion of separate agreement controls. I additionally remove all country pairs involving Israel, Jordan, Egypt, or Syria (Israel and its neighbors) or Yemen (the Houthis' base) from the estimation, as their trade is likely affected by the concurrent conflict.

the Houthis declared them the target and because their carriers were more risk-averse in transiting the Red Sea. To separate the Western response from the general effect, I interact the time-varying distance with an indicator W equal to one when either the importer or the exporter is a Western country.[10] The coefficient β_W in Equation (40) is then the *additional* Western response over the general effect β_G:

$$X_{ni,t} = \exp\left[\chi_{n,t} + \chi_{i,t} + \mu_{ni} + \beta_G \ln \text{dist}_{ni,t} + \beta_W \left(\ln \text{dist}_{ni,t} \times W\right)\right] \nu_{ni,t}. \tag{40}$$

Table 1 reports all three specifications. Per the DiD results, the average disruption reduced trade on exposed dyads by roughly five percent (column 1). The continuous distance elasticity is negative and significant (column 2), and the interaction in column 3 shows this response is concentrated among Western dyads: the general distance elasticity is statistically indistinguishable from zero ($\beta_G = -0.016$), while the Western interaction is large and significant ($\beta_W = -0.147$). This heterogeneity motivates the distinction between the homogeneous and heterogeneous general-equilibrium counterfactuals in Section 4.4. These estimates are conservative for two reasons. First, Comtrade values aggregate across modes of transport, while the disruption only raised the cost of seaborne trade. Air freight, which carries a small share of tonnage but a disproportionate share of value, is mechanically untreated such that the coefficients are attenuated relative to the response of maritime trade alone. Second, any substitution from sea to air on exposed dyads props up aggregate trade values, even as seaborne trade falls. I therefore read the estimates as lower bounds on the maritime response. The estimates are also robust to the treatment cut-off itself. Re-estimating column 3 with exposure defined at every threshold from any increase to a 50 percent detour, in steps of 10 percentage points, yields a Western coefficient between -0.13 and -0.17, significant at the five percent level in every case and at the one percent level at every cut-off up to 30 percent. Over the same range, the general coefficient never differs significantly from zero, and the difference-in-differences estimate in column 1 ranges from -0.03 to -0.05. The Western concentration of the response is therefore a feature of the data, not of the cut-off. The 30 percent threshold matters most for the general-equilibrium analysis, where it restricts the shock to trade flows that actually

[10] The set of Western countries comprises the nineteen largest economies, by 2022 GDP, in the union of NATO and the G7, excluding Turkey, whose relationship with Israel at the time made it an implausible target of a campaign framed as punishing Israel's supporters. Membership is dated to 2022, the baseline year; every member's accession predates the crisis, so target status cannot be a response to the economic damage the attacks caused.

ex-ante relied on the Bab el-Mandeb Strait.

Table 1: The Red Sea Disruption and Bilateral Trade

Dependent variable: trade $X_{ij,t}$	(1) DiD (exposed × post)	(2) Distance elasticity	(3) Distance elasticity
Exposed × Post	-0.049^{***}		
	(0.014)		
ln distance		-0.083^{***}	-0.016
		(0.029)	(0.042)
ln distance × Western			-0.147^{***}
			(0.052)
Exporter × time FE	Yes	Yes	Yes
Importer × time FE	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes
Observations	210,803	210,803	210,803
Clusters (pairs)	17,932	17,932	17,932
Pseudo R^2	0.994	0.994	0.994

Notes. PPML estimates (Santos Silva and Tenreyro, 2006) of the Bab el-Mandeb disruption on monthly bilateral trade, February 2023–June 2024. Column (1) is the difference-in-differences specification of Equation (38): Exposed equals one for pairs whose least-cost shipping route lengthens by at least 30 percent when the strait closes (exposure ≥ 1.30); pairs with smaller detours are controls; Post equals one from November 2023. Columns (2)–(3) replace the binary treatment with a continuous measure: the time-varying log distance of treated pairs switches from the open-strait to the closed-strait routing at the November 2023 cutoff, recovering a distance elasticity of trade; column (3) interacts it with an indicator for pairs in which either country is Western. All specifications absorb exporter-time, importer-time, and directional pair fixed effects and cluster standard errors by pair. The October 2023 transition month is excluded. Countries on or adjacent to the corridor (Yemen, Israel, Jordan, Egypt, Syria) are dropped. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Relating the estimates to the model. In the model, disruption enters through the bilateral trade cost term $\tau_{ni}^{1-\sigma}$, while my empirical treatment enters cost itself through the pass-through relation $\Delta \ln \tau_{ni,t} = \varphi \, \Delta \ln \text{dist}_{ni,t}$. Accordingly, the distance elasticity estimated in Equations (39)–(40) is not the trade elasticity itself, but the composite

$$\beta = (1 - \sigma)\, \varphi, \tag{41}$$

the product of the trade elasticity $1 - \sigma$ and the distance-to-trade-cost pass-through φ. These two objects play distinct roles and are separately sourced. Because the crisis is a short-run event and my seventeen-month window is poorly suited to identifying a structural substitution elasticity, I calibrate σ from the range of short-run trade elasticities supported by the literature rather than estimating it, considering $\sigma \in \{2.0, 3.5, 5.0\}$ throughout.[11] The pass-through φ is then the transportation-cost object that my regression identifies conditional on that calibration, $\hat{\varphi} = \hat{\beta}/(1 - \sigma)$: the pooled estimate implies $\hat{\varphi} = 0.083$, 0.033, and 0.021 at $\sigma = 2$, 3.5, and 5, respectively. The inferred trade-cost change therefore varies by a factor of four across the calibrated band and should not be read off any single elasticity. This matters because the estimates are short-horizon objects: trade responds far less to a cost shock over months than over years (Boehm et al., 2023), so a given trade response maps into a larger underlying cost change the lower is the relevant elasticity. As Section 4.4 shows, the shock entering the counterfactual depends only on the estimated composite and not on this decomposition; σ governs how that shock translates into prices and welfare, and I therefore report all general-equilibrium results across the full σ range rather than at a single point.

While the binary difference-in-differences result in column 1 establishes that the disruption reduced trade on exposed dyads, the continuous specifications in columns 2 and 3 discipline the general-equilibrium analysis that follows. In each case, I take the geography-based treatment, scaled by the recovered pass-through, as the trade-cost shock entering the model directly, and solve for the resulting changes in prices, welfare, and trade flows.

[11] The literature's "trade elasticity" corresponds to $\sigma - 1$ in my notation. Head and Mayer (2014) report a median of roughly 5 across structural gravity estimates, Simonovska and Waugh (2014) obtain about 4, and Fontagné et al. (2022) a mean near 5 at the product level, while Boehm et al. (2023) find a short-run (one-year) elasticity below 1 and a long-run value near 2. The band $\sigma \in \{2, 3.5, 5\}$ spans the short-run and long-run values relevant to a crisis measured over months.

4.3 Robustness of the Partial Estimates

Equation (38)'s identifying assumption is that, absent the disruption, trade of exposed and unexposed dyads would have followed parallel paths. That assumption is untestable, but its observable implication is the absence of differential pre-trends. I test for pre-trends with an event-study specification that replaces the single post indicator with a set of event-time dummies, interacted with the binary exposure measure:

$$X_{ni,t} = \exp\Big[\chi_{n,t} + \chi_{i,t} + \mu_{ni} + \sum_{e \neq -1} \beta_e \left(\mathbf{1}\{\text{event time} = e\} \times \text{Exposed}_{ni}\right)\Big] \nu_{ni,t}, \tag{42}$$

Event time is measured relative to November 2023, the first month of the campaign. October 2023 is the omitted reference, so each β_e is measured relative to that month. Unlike the specifications above, which exclude October as a transition month, the event study retains it as the reference: although the attacks began on 19 October, transit times and customs-reporting lags mean that October arrivals largely reflect shipments dispatched before the campaign, and any residual contamination would push towards attenuation bias of the post-treatment coefficients, against finding an effect. Figure 5 reports the results; Appendix Table A3 gives the coefficients and test statistics.

The post-treatment path has a distinctive shape. The effect appears immediately. Trade on exposed dyads is 5.1 log points lower in November 2023 ($p = 0.018$), and the decline deepens through December to a trough in January 2024 of -0.083 ($p = 0.013$). The effect remains significant in February 2024 (-0.060, $p = 0.022$) and marginally so in March (-0.053, $p = 0.097$), before attenuating to insignificance in April and May 2024, with a renewed decline in June 2024 (-0.088, $p = 0.009$). This timing profile, an immediate response concentrated in the first four months of the campaign that then weakens while treatment remains in force, is consistent with the transitory rerouting adjustment cost discussed in Section 4.4. Five of the eight post-treatment coefficients are negative and significant at the five percent level, while none is significantly positive, and the lags are jointly significant ($\chi^2(8) = 32.56$, $p < 0.001$). The leads, however, are also jointly significant ($\chi^2(8) = 38.21$, $p < 0.001$), so I cannot claim strictly parallel pre-trends. Three pieces of evidence indicate that this reflects month-to-month variation rather than a differential trend capable of generating the estimated effect. First, and most directly, there is no linear pre-trend: regressing trade on the interaction of exposure with event time over pre-treatment months alone yields a

slope of −0.0017 (s.e. 0.0036, $p = 0.64$), a precisely estimated zero. Second, only a single lead is individually significant at the five percent level: September 2023 (−0.048, $p = 0.022$). Excluding it, the remaining leads yield $\chi^2(7) = 22.84$ ($p = 0.002$), so this one month accounts for a large share of the rejection. Diagnostics indicate it is a genuine one-month level dip rather than a coverage artifact: the number of reporting treated pairs that month (3,066) is in line with adjacent months, while mean treated trade is roughly four percent below the average of August and October. Third, the pre-period coefficients show no systematic level difference: relative to October 2023 they average −0.002 and are mixed in sign, whereas the post-period coefficients average −0.049, reflecting the difference-in-differences estimate of Table 1. Any residual month-to-month variation in the pre-period is thus an order of magnitude smaller than, and uncorrelated in direction with, the estimated effect.[12] As a final, direct check on whether any residual pre-trend drives the results, Table 2 re-estimates all three main specifications allowing exposed dyads their own linear time trend. The difference-in-differences estimate is essentially unchanged, moving from −0.049 to −0.050 ($p = 0.063$), and—most importantly for the general-equilibrium exercises that follow—the Western distance elasticity is unchanged at −0.147 ($p = 0.006$). The pooled distance elasticity falls from −0.083 to −0.049 and loses significance once the trend is included, consistent with the pooled response being driven by the Western dyads that column 3 isolates. In each case the trend term is itself statistically insignificant. Important for later analysis, the heterogeneous specification that disciplines Section 4.4's preferred counterfactual is therefore essentially unaffected by detrending.

Spillovers through the Panama Canal. The control group naturally contains country pairs whose trade relies on the Panama Canal, the principal maritime alternative to Suez for trade between Asia and the Atlantic coast of the Americas. If the crisis changed the cost of transiting the Panama Canal, those pairs' trade would move for reasons unrelated to their own exposure to the Bab el-Mandeb, possibly contaminating the comparison. During the sample window, the canal was itself under strain: Gatun Lake, which feeds its locks, sat at its lowest recorded level every day from mid-July 2023, daily transits were cut from 38 to as few as 20 before settling at 24—a single transit

[12] The joint test of the leads is invariant to the choice of base month, since testing that all leads equal zero relative to any reference is equivalent to testing that all pre-treatment months are equal to one another. Re-estimating with September, August, July, or June 2023 as the reference therefore returns an identical statistic. Extending the window back to November 2022 adds three further leads and raises the statistic to $\chi^2(11) = 52.74$, driven by November and December 2022, which lie nearly a year before the attacks; I therefore retain the February 2023 start.

slot sold at auction for $4 million in early January 2024, as operators weighed Panama against an increasingly risky Suez route (Griffith, 2024; Ulrichsen and Krane, 2026). Since the importer-time and exporter-time effects absorb reallocation of a country's purchases toward or away from Panama-routed suppliers, pair-specific changes in the cost of Panama-routed trade would survive the fixed effects. I measure each pair's reliance on the canal exactly as I measure its reliance on the strait: I sever the canal in the routing network and recompute logistical distance, holding the routing costs and population weights fixed. Of the pairs in the estimation sample, 558, carrying 8.9 percent of sample trade, lengthen by at least 30 percent when the canal is severed; no pair is also exposed to the Bab el-Mandeb at that threshold, such that all of them are defined as controls. Table 3 reports four checks. Panama-exposed pairs' trade did fall after November 2023 relative to other control pairs (column 3), although a continuous measure of the Panama detour enters with the opposite sign and is only marginally significant (column 4), so the evidence for a Panama-specific shift is mixed. Importantly, accounting for it leaves the estimates essentially unchanged. Dropping the Panama-exposed pairs, controlling for their post-period shift with either measure, or restricting the control group to pairs exposed to neither canal, which removes 31 percent of sample trade, keeps the difference-in-differences estimate between -0.046 and -0.052 and the Western distance elasticity between -0.132 and -0.147, with all results remaining significant at the five percent level or better; the general distance elasticity remains indistinguishable from zero. Spillovers through the Panama Canal thus do not drive the estimates in Table 1.

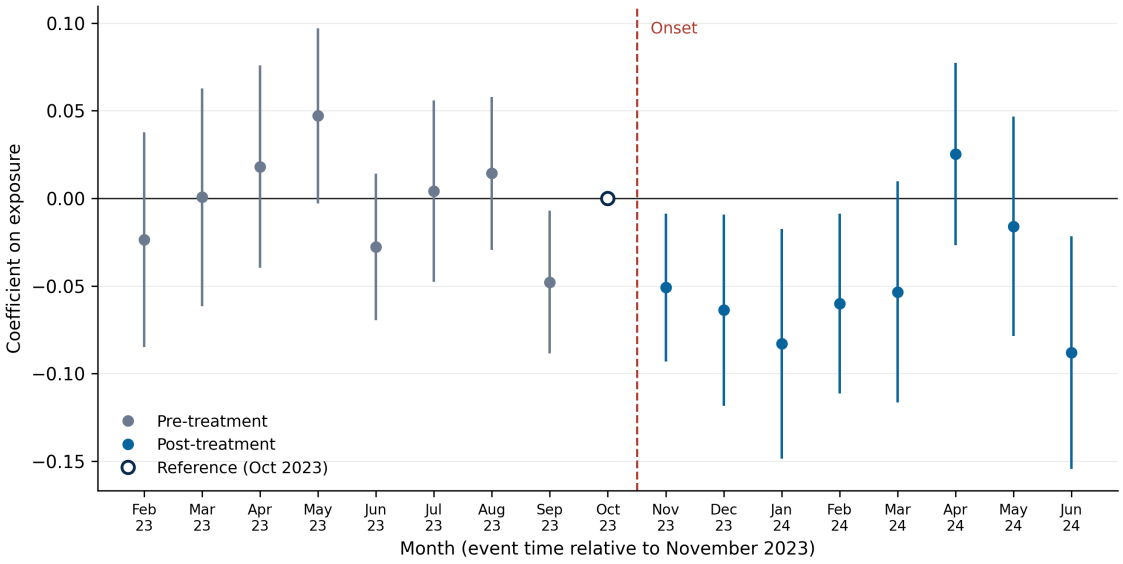


Figure 5: Event-study estimates around the Bab el-Mandeb closure. Coefficients on the binary exposure indicator interacted with month-of-event-time dummies, from Equation (42), with 95 percent confidence intervals clustered by directional pair. October 2023 is the omitted reference and is plotted at zero. The dashed line marks the November 2023 onset of the campaign. The effect emerges immediately, deepens to a trough in January 2024, and weakens by the spring of 2024 despite the rerouting continuing through the end of the sample. Point estimates, standard errors, and joint tests of the leads and lags are reported in Appendix Table A3.

Table 2: Robustness to a treated-specific linear trend. Each specification from Table 1 is re-estimated allowing exposed dyads their own linear time trend.

	DiD		Distance elasticity		Western split	
	(1)	(2)	(3)	(4)	(5)	(6)
Exposed × Post	−0.049***	−0.050*				
	(0.014)	(0.027)				
ln distance			−0.083***	−0.049	−0.016	−0.017
			(0.029)	(0.051)	(0.042)	(0.056)
ln distance × Western					−0.147***	−0.147***
					(0.052)	(0.053)
Exposed × trend		0.0002		−0.0023		0.0000
		(0.0031)		(0.0027)		(0.0028)
Treated-specific trend	No	Yes	No	Yes	No	Yes
Exporter × time FE	Yes	Yes	Yes	Yes	Yes	Yes
Importer × time FE	Yes	Yes	Yes	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	210,803					
Clusters (pairs)	17,932					

Notes. PPML estimates on monthly bilateral trade, February 2023–June 2024, with October 2023 excluded as a transition month. Exposed: Exposure ≥ 1.30. The treated-specific trend is exposure interacted with event time, where event time is measured in months relative to November 2023. Standard errors, clustered by directional pair, in parentheses. Countries on or adjacent to the corridor (Yemen, Israel, Jordan, Egypt, Syria) are dropped. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Table 3: Robustness to Panama Canal exposure. Columns 1 and 3 of Table 1 re-estimated after accounting for country pairs whose trade relies on the Panama Canal, measured by severing the canal in the routing network.

	Baseline	Drop Panama-exposed pairs	Control: Post × Panama-exposed	Control: Post × ln Panama detour	Control group exposed to neither canal
	(1)	(2)	(3)	(4)	(5)
Panel A: Difference-in-differences					
Exposed × Post	−0.049***	−0.046***	−0.046***	−0.051***	−0.052***
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Panama-exposed × Post			−0.041**		
			(0.020)		
ln Panama detour × Post				0.098*	
				(0.056)	
Observations	210,803	204,992	210,803	210,803	177,145
Panel B: Western split					
ln distance	−0.016	−0.012	−0.013	−0.021	−0.034
	(0.042)	(0.042)	(0.042)	(0.042)	(0.044)
ln distance × Western	−0.147***	−0.146***	−0.145***	−0.147***	−0.132**
	(0.052)	(0.052)	(0.052)	(0.052)	(0.053)
Panama-exposed × Post			−0.041**		
			(0.020)		
ln Panama detour × Post				0.098*	
				(0.055)	
Observations	210,803	204,992	210,803	210,803	177,145
Clusters (pairs)	17,932	17,432	17,932	17,932	15,186

Notes. PPML estimates with exporter-time, importer-time, and directional-pair fixed effects; sample and treatment as in Table 1. Panama exposure is the ratio of a pair's logistical distance with the Panama Canal severed to its distance on the complete network, computed on the same network, routing costs, and population weights as the Red Sea measure. Panama-exposed pairs are those whose distance rises by at least 30 percent (558 pairs, 8.9 percent of sample trade, all in the control group); column 5 drops every control pair whose distance rises at all when the canal is severed (2,953 pairs, 31 percent of sample trade). No pair is exposed to both chokepoints at the 30 percent threshold. Standard errors, clustered by directional pair, in parentheses. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

4.4 General-Equilibrium Counterfactuals

While partial-equilibrium estimates provide novel insight into how the Red Sea crisis impacted trade for treated country pairs, they cannot speak to how some countries might benefit from such a disruption. As demonstrated in Section 3, price-index adjustments generate terms-of-trade improvements for countries that are relatively less affected, and in certain cases these improvements offset the direct cost of the disruption and raise welfare. These are precisely the price effects absorbed by the importer- and exporter-time fixed effects in Section 4.2's estimating equations. The partial estimates thus recover the direct cost shock while differencing out the general-equilibrium response that generates the paper's central prediction. To recover that response, I solve the model's counterfactual equilibrium—taking Section 4.2's estimates as pseudo-exogenous trade-cost shocks—using the exact hat algebra of Dekle et al. (2008).

Equilibrium in changes. Exact hat algebra requires no knowledge of the model's unobservable primitives. Denote by $\hat{z} \equiv z'/z$ the ratio of any variable's counterfactual value to its baseline value. Since the unobserved taste parameters β_i and trade cost *levels* τ_{ni} enter the equilibrium conditions only through observed expenditure shares, expressing the system in changes eliminates them entirely, and the counterfactual can be computed from baseline trade shares, revenues, expenditures, and the elasticity σ alone. Taking the ratio of counterfactual to baseline trade shares in (5) and applying the price index definition,

$$\hat{s}_{ni} = \frac{\left(\hat{\tau}_{ni}\,\hat{p}_i\right)^{1-\sigma}}{\hat{P}_n^{1-\sigma}}, \qquad \hat{P}_n^{1-\sigma} = \sum_i s_{ni}\left(\hat{\tau}_{ni}\,\hat{p}_i\right)^{1-\sigma}, \tag{43}$$

where the second expression follows from dividing the counterfactual price index by its baseline counterpart, recognizing that baseline trade shares are the weights. Equation (43) is the model's cost-of-living channel: country n's price index rises with the trade-cost shocks on its import links, weighted by how large a share of its expenditure each supplier provided *before* the disruption. A country that ex-ante sourced little through the strait is barely affected; one that sourced heavily through it absorbs a large shock.

Market clearing (12) requires that each country's sales equal world expenditure on its good. In the endowment economy of Section 3, supply is fixed so revenue moves one-for-one with the mill

price, $R'_i = \hat{p}_i R_i$, and the counterfactual system is

$$\hat{p}_i R_i = \sum_n s'_{ni} E'_n, \qquad E'_n = \hat{p}_n R_n + D_n, \tag{44}$$

where $D_n \equiv E_n - R_n$ is country n's aggregate trade imbalance, held constant in the counterfactual.[13] Equations (43)–(44) constitute a system of $N+1$ equations in the $N+1$ unknown price changes, with everything else observed. I solve this system by damped fixed-point iteration: starting from $\hat{p} = \mathbf{1}$, I compute price indices and shares from (43), evaluate the excess demand implied by (44), and update prices where demand and supply do not balance, iterating until the maximum relative excess demand falls below 10^{-10}. The gross-substitutes property established in Lemma 1 guarantees a unique solution and convergence of the iteration. Because the system determines only relative prices, I normalize world revenue, $\sum_i \hat{p}_i R_i = \sum_i R_i$; welfare changes are invariant to this choice.

Welfare. Using counterfactual prices, welfare changes follow directly from the indirect utility function (13). Unlike the marginal expressions of Section 3.5, these are exact for a discrete shock:

$$\Delta u_n = \underbrace{R_n\big(\hat{p}_n - 1\big)}_{\text{terms of trade}} - \underbrace{E_n \ln \hat{P}_n}_{\text{cost of living}} . \tag{45}$$

The first term captures the country's position as a seller: with its endowment fixed, a higher mill price raises the value of its sales. The second captures its position as a buyer: a higher price index means a costlier consumption bundle. Countries whose import links are relatively unperturbed face a smaller cost-of-living increase and may, through demand reallocation and the mill-price response, experience a welfare gain, as Proposition 1 predicts. Applying the identity $\hat{p}_n - \hat{P}_n = -\hat{s}_{nn}/(\sigma - 1)$ from Section 3.7, (45) decomposes into the augmented ACR form of (35): a strategic terms-of-trade term proportional to net exports, and the canonical own-trade-share term.

Constructing the shock. I construct the trade cost shock using the logistical distances of Section 4.1. Let $\text{disp}_{ni} \equiv \text{dist}_{ni}^{\text{closed}}/\text{dist}_{ni}^{\text{open}} \geq 1$ denote the proportional increase in logistical distance

[13] Holding imbalances fixed in levels is standard convention; it isolates the reallocation induced by the trade-cost shock from any change in aggregate borrowing or lending.

a dyad experiences when the strait closes. Given a pass-through relation $\Delta \ln \tau_{ni} = \varphi_{ni} \ln \text{disp}_{ni}$, the object entering (43) is

$$\hat{\tau}_{ni}^{1-\sigma} = \text{disp}_{ni}^{(1-\sigma)\varphi_{ni}} = \text{disp}_{ni}^{e_{ni}}, \qquad e_{ni} \equiv \eta_{\text{dist}} + b_{ni}, \tag{46}$$

where e_{ni} is the distance elasticity of trade applied to dyad (n, i). It has two components: η_{dist}, the ordinary distance elasticity that prices the added kilometers, and b_{ni}, the crisis-specific composite estimated in Section 4.2, which prices the transitory adjustment premium. The three exercises that follow differ *only* in b_{ni}, so they nest and are directly comparable: Exercise 1 sets $b_{ni} = 0$; Exercise 2 sets $b_{ni} = \hat{\beta}$ for every dyad; and Exercise 3 sets $b_{ni} = \hat{\beta}_G + \hat{\beta}_W W_{ni}$, where W_{ni} indicates that either country is Western. Exercise 3 is the paper's benchmark and the specification carried into Section 5; Exercises 1 and 2 are building blocks that isolate, respectively, the permanent distance channel and a homogeneous adjustment premium. Two features deserve emphasis. First, the shock is a function of distance elasticities of trade alone, so the direct trade response on each dyad does not depend on how those elasticities decompose into σ and φ; the calibrated σ governs only the general-equilibrium price response and the translation of trade changes into welfare. Second, since $e_{ni} < 0$ and $\text{disp}_{ni} \geq 1$, the shock term satisfies $\hat{\tau}_{ni}^{1-\sigma} \leq 1$: exposed dyads see their trade-cost term shrink and their trade fall, as intended.

Four conventions govern the shock matrix. First, consistent with the treatment definition of Section 4.2, only dyads whose logistical distance rises by at least 30 percent are shocked; all other dyads are set to $\text{disp}_{ni} = 1$, so the shocked set carries 11.1 percent of baseline international trade, the Suez corridor's share. The smaller permanent detours of pairs below the threshold are therefore excluded from every exercise, including Exercise 1, which makes the counterfactuals conservative in this respect as well. Second, all trade dyads involving the disruptor are set to $\text{disp}_{ni} = 1$, in line with the model's assumption that the disruptor degrades the corridor without inflicting higher trade costs on itself. Third, a relatively small number of dyads register logistical distance *decreases* when the Red Sea is closed, an artifact of the land-route component of the distance measure; these are set to $\text{disp}_{ni} = 1$, as are dyads absent from the routing data. Finally, I restrict attention to the 232 countries appearing in both the trade and distance data; those dropped account for less than 0.01 percent of world trade. I further drop seven territories with no recorded trade in either

direction and 23 small economies, predominantly island territories and free ports, whose baseline trade imbalance exceeds three times their gross output, which I treat as data artifacts rather than real imbalances.[14] The final sample contains 202 countries.

Exercise 1: the distance channel. I begin with a simple counterfactual in which the closure's sole consequence is that exposed dyads become farther apart. I translate the distance increase into a trade-cost shock using a standard estimate of the distance elasticity of trade, η_{dist}, drawn from the gravity literature rather than from the crisis itself:

$$\hat{\tau}_{ni}^{1-\sigma} = \text{disp}_{ni}^{\eta_{\text{dist}}}, \qquad \eta_{\text{dist}} = -0.6. \tag{47}$$

This exercise isolates the *permanent* footprint of the closure. The event-study estimates in Section 4.3 imply that trade's heightened responsiveness to distance among exposed dyads is concentrated in the first months of the campaign: the response is largest between November 2023 and March 2024 and becomes insignificant by the spring of 2024, despite the routes remaining detoured. This suggests that the estimated coefficient on time-varying distance, in the presence of pair fixed effects, captures an additional and transitory per-kilometer *adjustment cost*—scarce vessel capacity, war-risk insurance premia, and disrupted schedules—beyond the simple increase in shipping distance. Once shipping has adjusted to the new, longer routes around the Cape of Good Hope, the affected dyads remain farther apart but face only ordinary costs per increment of distance. Equation (47) is therefore the counterfactual describing where the world settles once treated pairs fully adjust. Additionally, because η_{dist} is the product of gravity estimates identified from years of cross-sectional variation in distance, it is the appropriate parameter for valuing a permanent distance change, a question a several-month event window is poorly suited to answer. This exercise also validates the counterfactual machinery: since it uses no crisis-specific estimate, its reproducing the qualitative predictions of Section 3 on observed trade shares confirms the model's incidence predictions before any estimated shock is layered on.

[14] Gross output is measured as total baseline sales, inclusive of domestic sales. The largest economies removed by this filter are Macao, Gibraltar, and the Marshall Islands.

Exercise 2: adding the adjustment premium. The previous exercise deliberately abstracted from the transitory adjustment cost; I now reintroduce it. Because the pair fixed effects already absorb each dyad's ordinary trade costs, the estimated coefficient on time-varying distance measures the *additional* responsiveness of exposed dyads, and therefore layers on top of the standard distance elasticity rather than replacing it. The trade-cost shock becomes

$$\hat{\tau}_{ni}^{1-\sigma} = \text{disp}_{ni}^{\eta_{\text{dist}}+\hat{\beta}}, \tag{48}$$

where $\hat{\beta} = -0.083$ is the pooled estimate from column 2 of Table 1. Treatment across exposed dyads is thus homogeneous: the disruption raises per-kilometer costs by the same proportion regardless of who is trading. Equation (48) corresponds to the height of the disruption, when both the increased distance and the route-unfamiliarity premium affect trade costs. Because $\hat{\beta} < 0$, the exponent is more negative than in (47), so every treated pair faces a strictly larger cost shock and the welfare consequences of Exercise 1 are amplified. Together the two exercises bracket the disruption's effect over its lifecycle, and their difference isolates the welfare cost attributable to the rerouting adjustment alone, separate from the permanent cost of the added distance. Both, and their difference, are reported for every country in Appendix Table A4.

Exercise 3: heterogeneous exposure. As documented in Section 4.2 and column 3 of Table 1, country pairs with at least one Western-aligned member drove the negative coefficient on time-varying distance: interacting the treatment with a Western indicator renders the general coefficient statistically insignificant. This is economically meaningful. The Houthi attacks sought to punish Western-aligned countries for their support of Israel, and those countries therefore constitute the target of the campaign, from which the disruptor seeks concessions; one should expect them to be more sensitive to the attacks. To recover this disaggregated effect, I split the homogeneous cost shock according to the disaggregated estimates. All treated dyads still experience the negative impact of longer distances according to their disparity measure, but Western-aligned pairs face an additional deterioration:

$$\hat{\tau}_{ni,W}^{1-\sigma} = \text{disp}_{ni}^{\eta_{\text{dist}}+\hat{\beta}_G+\hat{\beta}_W} \qquad \hat{\tau}_{ni,NW}^{1-\sigma} = \text{disp}_{ni}^{\eta_{\text{dist}}+\hat{\beta}_G}, \tag{49}$$

where $\hat{\beta}_G = -0.016$ is the general coefficient and $\hat{\beta}_W = -0.147$ the additional Western coefficient from column 3 of Table 1, so that Western-aligned pairs carry a total premium of $\hat{\beta}_G + \hat{\beta}_W = -0.163$ and all other pairs a premium of -0.016. As in Exercise 2, these shocks incorporate the transitory adjustment cost, so the results represent a short-lived effect on welfare and trade. Because the non-Western coefficient is close to zero, this specification implies that the realignment is driven almost entirely by Western rerouting: non-Western dyads experience little beyond the distance channel itself. Comparing these results against Exercise 1 again recovers the additional effect of short-lived adjustment, now disaggregated by Western alignment.

Results. For each exercise, I compute the general-equilibrium welfare changes across three feasible elasticity values $\sigma \in \{2, 3.5, 5\}$. Figure 6 compares the percentage welfare changes computed for each exercise across these three values and demonstrates two findings. First, welfare effects are strongly correlated across elasticities: for any exercise, the Spearman rank correlation (ρ) between percentage welfare changes at any two elasticity values exceeds 0.99. Second, there is clean monotonicity: higher elasticity yields attenuated welfare effects. This is expected—with higher elasticity, countries are more capable of recouping lost trade from disruption. Because of these two facts, the rest of this section focuses on the case where $\sigma = 3.5$, the intermediate of the three values. I organize the discussion around Exercise 3, the benchmark, and then use Exercises 1 and 2 to decompose it.

Benchmark results: Exercise 3. The benchmark welfare changes are mapped in Figure 7 and reported, for the twenty largest economies and Iran, in column 3 of Table 4.[15] Three findings stand out.

First, roughly two thirds of countries lose. This is unsurprising: to gain from the disruption, a country must be sufficiently unexposed that the price effect of demand reallocating toward its goods exceeds the direct cost of its own higher trade costs. Relatively few countries clear that bar. Losses are also far larger than gains. The gradient scale in Figure 7 shows that the biggest losers experience double-digit percentage welfare declines, while the biggest winners gain only fractions of a percent. Because these are proportional effects, large economies like China and Japan register

[15] Full country-level welfare results are reported in Table A4 in the appendix.

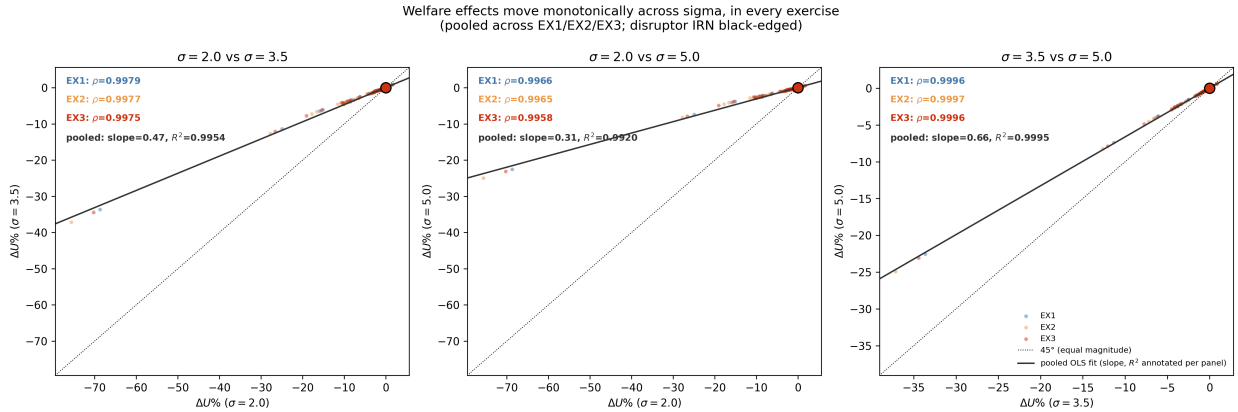


Figure 6: Welfare effects across elasticity. Each panel plots country-level percentage welfare changes ($\Delta U\%$) at one σ value against another, for each pairing of $\sigma \in \{2, 3.5, 5\}$, pooling across all three exercises (EX1, EX2, EX3, color-coded). The dotted line marks the 45-degree line, along which welfare changes would be identical across the two σ values; the solid line is the pooled OLS fit. The Spearman rank correlation coefficient (ρ), reported separately by exercise in each panel, exceeds 0.994 throughout, indicating that the cross-country ranking of winners and losers is essentially unaffected by the choice of σ. The pooled OLS slope is well below one in every panel (0.31–0.66), reflecting systematic attenuation of welfare effects as σ rises: greater substitutability lets countries more easily reroute trade around the disruption, softening its welfare consequences.

small percentage losses that are nonetheless substantial in absolute terms. This helps explain China's willingness to strike an agreement with the Houthis securing safe passage for its vessels through the Strait (Bloomberg News, 2024). At the other extreme, the United States and Canada register small welfare gains: less than one percent of their international trade is with partners whose routes lengthen by 30 percent or more, so they are essentially unshocked directly and collect the terms-of-trade gain from demand reallocation.[16]

Second, exposure rationalizes the geographic distribution of winners and losers. Countries clustered around the Strait, which rely on it heavily for shipping, tend to lose, as reflected in their dark red hue in Figure 7. Countries whose trade flows were ex-ante unlikely to transit the Strait

[16]The U.S. gain is small (0.008 to 0.028 percent of expenditure across exercises and elasticities) and follows from the geography of its trade: its trade with Asia crosses the Pacific and its trade with Europe the Atlantic, so the detour taxes others' trade rather than its own. In the benchmark, the detour depresses the factory-gate prices of Asian exporters whose sales to Europe it taxes (by 0.40 percent for China and 0.80 percent for Vietnam, relative to world income), so the import-weighted price of U.S. suppliers falls by 0.19 percent against 0.08 percent for U.S. output—the reallocation through which bystanders gained from the U.S.–China trade war (Fajgelbaum et al., 2024). Sector-level evidence points the same way on the export side: European buyers shifted toward Atlantic-basin energy suppliers "with no logistics constraints" (CNBC, 2024), and U.S. diesel shipments to Europe reached their highest level since 2018 in July 2024 (Perkins, 2024). The model omits costs that push the other way, including a roughly 70 percent rise in transpacific freight rates as capacity was diverted from routes that never transit the strait (Dunn and Leibovici, 2024) and close to $2 billion of U.S. naval ordnance expended against the Houthis in 2024 (Holzer and Cichon, 2025); the U.S. figure is best read as the absence of direct trade exposure rather than as a benefit from the crisis. It does not bear on the sign result of Section 5: the greed-only bound $\bar{\kappa}' = u_1'(d^*)$ does not involve the target's welfare, so $\lambda > 0$ requires only that the bloc as a whole is harmed at the margin, which holds under both weightings.

tend to experience moderate welfare increases. The winners are concentrated on the western coast of Africa, where the rerouting a closure requires is relatively small, and in South America, where routes to major trading partners avoid the Red Sea entirely. Figure 8 confirms this relationship directly, plotting each country's welfare change against its mean increase in log distance across all trading partners. The downward slope shows that welfare is decreasing in exposure: of the 63 countries that gain, all but two lie below median exposure, and the gains of those two are economically negligible.

Third, and central to the inversion exercise of Section 5, Iran's own welfare gain is very small. The final row of Table 4 reports Iran's percentage welfare change across all three exercises; the narrow distribution centered near 0.005 percent indicates that, from a purely economic standpoint, Iran gains little from the disruption it engineered. This result motivates Section 5's examination of the geopolitical motive and the conditions under which the model can rationalize a disruption of the observed magnitude.

Decomposing the benchmark: Exercises 1 and 2. Columns 1 and 2 of Table 4 decompose the benchmark. Column 1 (Exercise 1) retains only the permanent distance channel; column 2 (Exercise 2) adds a homogeneous adjustment premium on all treated dyads. Moving from the first to the second, shifts in relative welfare are small. Countries that were losing tend to lose slightly more, while countries that were winning tend to win slightly more.[17] The clear implication is that winners, due to stronger route deterioration between treated dyads, effectively benefit more from demand reallocation and its associated terms-of-trade improvement.

Exercise 3, the benchmark, then reallocates that homogeneous adjustment cost according to whether a dyad includes at least one Western country, as defined above. As Section 4.2 demonstrates, this results in a near-zero adjustment cost for non-Western pairs, and a relatively large adjustment cost for Western pairs. More concretely, relative to Exercise 2, the distance elasticity associated with the disruption becomes more severe for Western countries, increasing their trade's response, and less severe for non-Western pairs. Figure 9 plots welfare changes under Exercise 3 against those under Exercise 2 for the Western-aligned set (left) and the ten largest non-Western-aligned countries (right). Values above the 45-degree line indicate that welfare losses under Exercise

[17] Table A4 in the appendix provides a more complete picture of this.

2 exceeded those under Exercise 3; values below indicate that a country was made worse off by the disaggregation. The left panel demonstrates that modeling the explicitly larger adjustment cost for Western countries harms them more, as all points but one lie below the 45-degree line; the exception is the United States, whose welfare change is slightly positive and essentially identical under both exercises. Further, larger losses under Exercise 2 are amplified more, as shown by the greater distance of the countries in the bottom-left from the 45-degree line. The right panel shows the reverse. Non-Western-aligned countries are now less burdened with the homogeneous adjustment cost that was driven by their Western peers. In most cases, this slightly increases their welfare: seven of the ten points lie above the 45-degree line, including China, the largest. The exceptions, Brazil, Australia, and Indonesia, lie essentially on the line, as their welfare changes are close to zero under both exercises.

Limitations. Three features of the quantitative model bear on how its magnitudes should be read. The model has a single sector and no input-output linkages, so it cannot capture the amplification that occurs when a costlier imported input raises the price of downstream output, nor the fact that energy—the sector most dependent on the corridor—is among the least substitutable. Both omissions push the reported losses toward zero; a multi-sector extension in the spirit of Caliendo and Parro (2015) would enlarge them. The disruptor's own dyads are held at $\text{disp}_{ni} = 1$, consistent with the theory, although Iran's own trade with Europe transits the corridor; allowing that exposure would lower Iran's gain and strengthen the conclusions of Section 5. Finally, the endowment structure abstracts from supply responses and from the freight-rate spikes that the crisis produced, which fall on shippers everywhere rather than only on exposed dyads. Each of these would sharpen the incidence reported here without altering its direction. The following section applies the theory developed in Section 3 to exploit Exercise 3's results and understand the conditions under which Iran's observed disruption of the Red Sea is rationalizable.

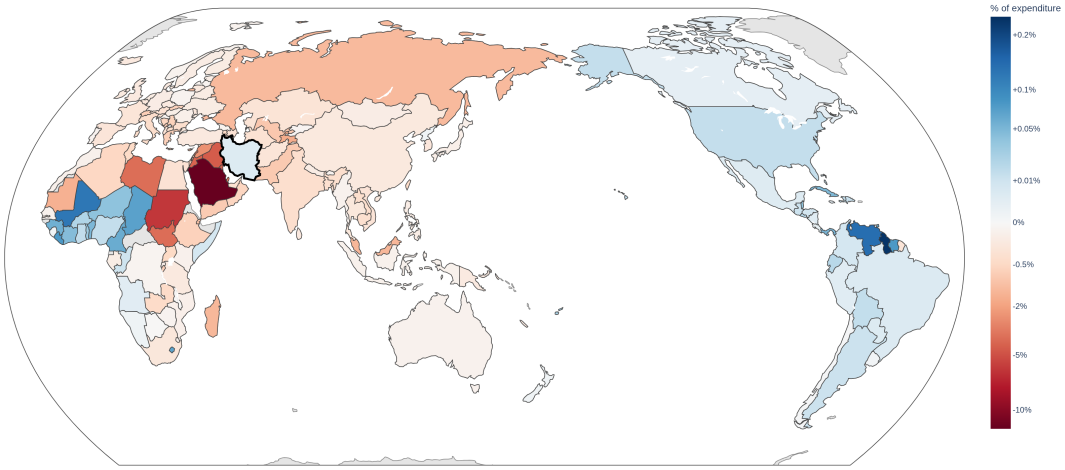


Figure 7: Welfare change ($\Delta U\%$ of baseline expenditure) by country under Exercise 3 (benchmark), $\sigma = 3.5$. Twenty-two countries are shown in neutral gray and excluded from the color scale: twenty-one with an unreliable baseline own-trade share (below 10 percent or exactly 100 percent), and Eritrea, the largest loser in every exercise, whose thin trade data make its magnitude unreliable and whose inclusion would compress the scale for every other country.

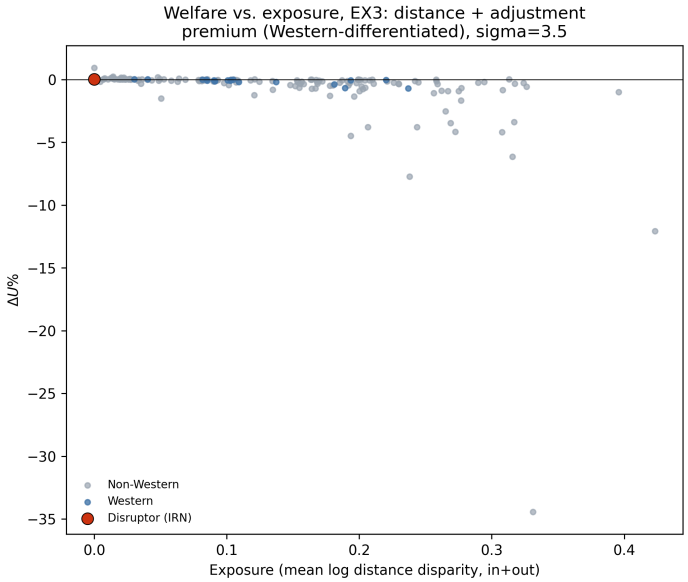


Figure 8: Welfare change ($\Delta U\%$ of baseline expenditure) on Exposure (measured as the mean increase in log distance across importers and exporters) under Exercise 3 (benchmark), $\sigma = 3.5$. Western countries are shown in blue. The general downward trend demonstrates that higher exposure leads to larger welfare losses.

Table 4: Welfare change by specification, 20 largest economies by 2022 nominal GDP (World Bank) and Iran, $\sigma = 3.5$

	% change in utility		
ISO	EX1	EX2	EX3
USA	0.01	0.01	0.01
CHN	-0.12	-0.13	-0.13
JPN	-0.03	-0.03	-0.04
DEU	-0.06	-0.06	-0.07
IND	-0.32	-0.36	-0.35
GBR	-0.09	-0.10	-0.12
FRA	-0.16	-0.18	-0.20
RUS	-1.23	-1.38	-1.25
CAN	0.00	0.00	0.00
ITA	-0.32	-0.36	-0.40
BRA	0.00	0.01	0.00
AUS	-0.02	-0.02	-0.02
KOR	-0.16	-0.17	-0.17
MEX	0.00	0.00	0.00
ESP	-0.16	-0.17	-0.19
IDN	-0.02	-0.03	-0.03
SAU	-11.34	-12.63	-12.08
NLD	-0.06	-0.07	-0.08
TUR	-0.12	-0.13	-0.12
CHE	-0.71	-0.79	-0.74
IRN†	0.0051	0.0056	0.0046

Rows are the 20 largest economies by 2022 nominal GDP (current US$, World Bank), in descending order. Iran (†), the disruptor, is not among the 20 largest by this measure – shown below the rule for reference, at higher precision since its welfare change is far smaller in magnitude. Countries with a degenerate baseline own-trade share are excluded throughout this paper's figures and tables (see data appendix); none of the countries above are affected.

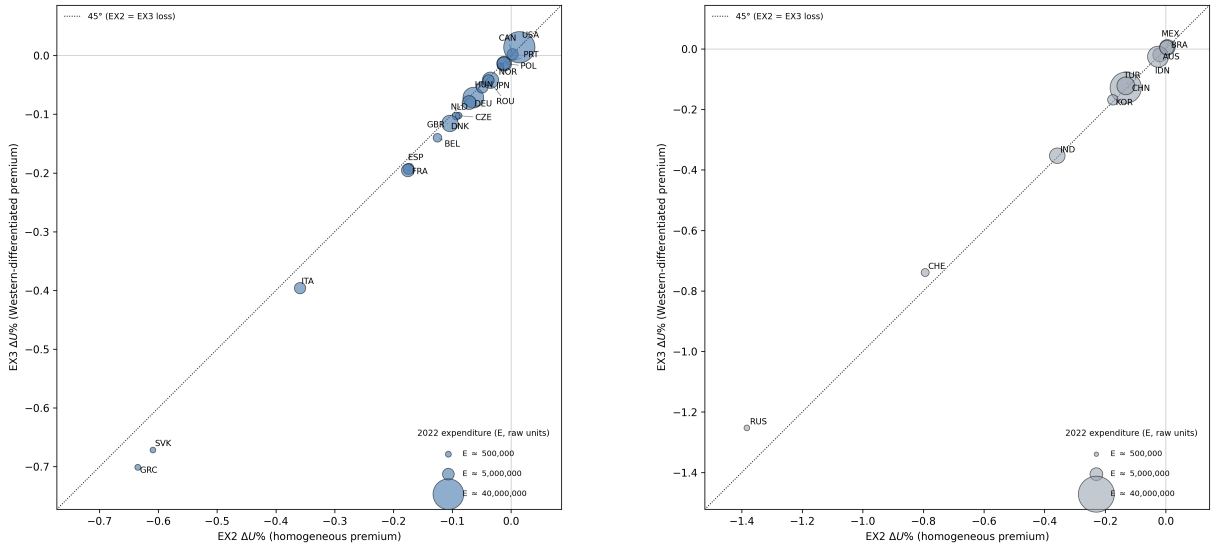


Figure 9: Exercise 3 (benchmark) versus Exercise 2. Welfare change ($\Delta U\%$ of baseline expenditure) under Exercise 3 plotted against Exercise 2, $\sigma = 3.5$. The left panel shows the nineteen Western-aligned countries; the right panel shows the ten largest non-Western-aligned economies. Points below the 45-degree line indicate a country made worse off by disaggregating the adjustment cost by Western alignment; points above indicate the reverse.

5 Recovering the Disruptor's Incentive

While the exercises above provide a thorough picture of how the Red Sea crisis impacted countries and their trade flows, they do little to explain the motivation behind the attacks. Without further context, it appears Iran had little incentive to disrupt the Red Sea at all. Its welfare gain is positive, owing to general-equilibrium price effects, but small: approximately +0.0046 percent of expenditure in the heterogeneous exercise. Under what circumstances could such a small gain motivate a costly, sustained military campaign? One possibility is that Iran holds a non-pecuniary incentive to harm its adversaries, captured by $\lambda > 0$ in (14). This section inverts the disruptor's first-order condition (28) to ask what the observed disruption reveals about λ. Proposition 2 establishes that greed alone can rationalize *some* disruption: whenever $\kappa'(0) < u_1'(0)$, the disruptor strictly prefers $d^* > 0$. It is silent, however, on whether greed can rationalize the *particular, large* intensity observed during the crisis. Iran's near-zero economic gains suggest it cannot, and confronting the model's implied cost of disruption with independent estimates of what Iran actually spends on the Houthi movement confirms this: no value of the trade elasticity in this paper's calibrated band reconciles the observed disruption with a purely self-interested disruptor.

5.1 Setup

An empirical target bloc. Section 3's model has a single target country that the disruptor wishes to harm. While the nineteen Western-aligned countries identified in Section 4.2 are geographically and economically disparate, their differential treatment when exposed—along with Iranian rhetoric surrounding the attacks—identifies them *collectively* as the likely target. I therefore work in the setting of Exercise 3 and generalize the single target's welfare to a weighted average over the Western set W:

$$\bar{u}_W(d) \equiv \sum_{n \in W} \omega_n \, u_n(d), \qquad \sum_{n \in W} \omega_n = 1. \tag{50}$$

My preferred weighting is by GDP, which reflects a disruptor concerned with aggregate rather than average harm and places the largest weight on the United States, the campaign's principal declared adversary. As a robustness check I also report equal weights ($\omega_n = 1/19$), the natural empirical analogue of the model's symmetric benchmark. Because the United States is itself only lightly exposed and gains slightly at the margin (Section 4.4), GDP weighting yields the *smaller* of the two harm measures, so the GDP-weighted results are the conservative ones. Both weightings are averages over the nineteen economies; the corresponding bloc totals are nineteen times the equal-weighted figures and are noted alongside the headline results below. Writing $H(d) \equiv -\bar{u}_W(d)$ for the resulting harm measure, sign-flipped so that $H \geq 0$ when W is hurt on net, the first-order condition (28) becomes

$$u_1'(d^*) + \lambda \, H'(d^*) = \kappa'(d^*), \tag{51}$$

which nests (28) exactly when $|W| = 1$, as in Section 3. Because λ is defined relative to a *weighted average* of Western welfare, its scale depends on the weighting scheme; GDP-weighted figures are the headline results throughout, with equal-weighted figures alongside them.

The identification problem. The marginal cost of disruption, $\kappa'(d^*)$, is not directly observed. Equation (51) is therefore one equation in two unknowns, and λ is not point-identified. Rather than manufacture a cost estimate, I proceed in two steps. I first report the knife-edge value $\bar{\kappa}' \equiv u_1'(d^*)$, the *unique* marginal cost at which a purely greedy disruptor ($\lambda = 0$) would have chosen exactly

the observed intensity, together with the locus of (κ', λ) pairs consistent with (51) across assumed costs. I then confront that knife edge with an independent, externally sourced estimate of Iranian expenditure on the Houthi movement, which is sufficient to sign λ even though it does not pin its exact value.

Parameterizing intensity. The disruption technology (11) applies a common $g(d)$ across every disrupted link, with the functional form left as a primitive. Empirically, the data supply dyad-specific baseline disparities rather than a single common intensity, so I replace the common $g(d)$ with a dyad-specific term, scaling Exercise 3's estimated exponent directly:

$$\hat{\tau}_{ni}(d)^{\,1-\sigma} = \begin{cases} \text{disp}_{ni}^{\,(\eta_{\text{dist}}+\hat{\beta}_G+\hat{\beta}_W)d}, & n \in W \text{ or } i \in W \\ \text{disp}_{ni}^{\,(\eta_{\text{dist}}+\hat{\beta}_G)d}, & \text{otherwise.} \end{cases} \tag{52}$$

Accordingly, $d = 0$ reproduces $\hat{\tau}_{ni} = 1$ for every pair (the undisrupted world) while $d = 1$ reproduces (49) exactly, which I confirm numerically by re-solving at $d = 1$ and matching every country's welfare change to floating-point precision. Because the intensity d has no real-world units—in the theory, only $d = 0$ and $d = \bar{d}$ are economically anchored—I extend that convention and define the observed closure to occur at $d^* = 1$. This is a normalization rather than a measurement. It treats the observed disruption as interior to the feasible range, which the data support: Suez transits fell by roughly 70 percent rather than to zero, indicating that the realized disruption stopped short of the complete closure that the bypass cap represents.

I numerically compute $u_n'(d^*)$ by a central difference,

$$u_n'(d^*) \;\approx\; \frac{u_n(1+\varepsilon_d) - u_n(1-\varepsilon_d)}{2\varepsilon_d}, \qquad \varepsilon_d = 0.05 \text{ (baseline)}, \tag{53}$$

re-solving the full equilibrium system (43)–(44) at $d = 1 \pm \varepsilon_d$ under (52), exactly as in the exercises above, and differencing the resulting welfare changes. Recomputing (53) for $\varepsilon_d \in \{0.10, 0.05, 0.025, 0.01\}$ confirms that both u_1' and H' move by less than 0.05 percent across a tenfold change in step size at every σ, so the reported derivatives are not simply artifacts of the finite-difference step.

Units and magnitudes. The function $u_n(d)$ gives how many units of the numeraire, at baseline prices, would leave country n exactly as well off as the disruption does, so $u_1'(d^*)$ is the marginal numeraire value to Iran of increasing the disruption. Because Iran's own trade dyads are held at $\text{disp}_{ni} = 1$, this is a pure general-equilibrium spillover: it arises entirely from how other countries' prices move, filtered back through Iran's baseline import shares and terms of trade. It is also computed only through trade that exists, and that trade is concentrated on heavily exposed partners. Weighting each partner's exposure (its mean increase in log distance across all of its trading partners, as in Figure 8) by its share of Iran's 2022 trade gives 0.19, against a median of 0.13 for the same statistic across all countries, which places Iran in the top decile; the United Arab Emirates, India, and Turkey, together 27 percent of Iran's trade, each rank in the top 15 percent of countries by exposure. Iran therefore has the access to disrupted markets that Appendix B shows the gain requires, and its gain is small because its share of those markets is small. The measure $H'(d^*)$ is the same object, sign-flipped and aggregated over W. Per the ITPD-S data, Iran's 2022 international trade is $51.5 billion in imports and $54.7 billion in exports, closely matching independent estimates of roughly $50–60 billion.[18] Throughout, I benchmark against Iran's international trade turnover of $106.2 billion per year, or $290.9 million per day. Because the counterfactual is calibrated to annual 2022 flows, every welfare magnitude in this section is an annual flow: $u_1'(d^*)$ and $H'(d^*)$ are dollars per year that the disruption is sustained, and the cost κ that the first-order condition weighs against them is accordingly the annual cost of sustaining it.

5.2 Findings

Table 5 delivers two results. First, the disruption is extraordinarily *leveraged* at the margin. Pushing it further gains Iran between $48 and $78 million while costing the GDP-weighted average Western economy between $1.6 and $8.1 billion: a ratio of roughly 34 to 103 dollars of harm inflicted per dollar of own gain (49 to 129 under equal weights). Summed rather than averaged over the nineteen economies, aggregate Western harm is $46 to $192 billion per year, or roughly 950 to 2,450 dollars per dollar of Iranian gain. Whatever Iran's motives, the corridor provides a

[18] Iran's *domestic* absorption in ITPD-S is considerably less reliable, implying a gross output roughly 4.5× Iran's actual 2022 GDP against multiples near 2.5× for the Western bloc and 3.1× for the world, likely reflecting imputation for a heavily sanctioned, non-reporting economy. I therefore benchmark Iranian magnitudes against international trade turnover, which validates well, rather than against gross output.

Table 5: Inversion results, by σ. Marginal welfare effects at the observed disruption intensity, computed under the heterogeneous specification of Exercise 3 with $\varepsilon_d = 0.05$.

σ	Iran's marginal gain $u_1'(d^*)$	Marginal harm per Western economy, $H'(d^*)$ GDP-wt.	Equal-wt.	Leverage H'/u_1'	Knife-edge $\bar{\kappa}'$
2.0	\$78.4M	\$8.1B	\$10.1B	103	\$78.4M
3.5	\$63.6M	\$2.8B	\$3.8B	43	\$63.6M
5.0	\$47.9M	\$1.6B	\$2.4B	34	\$47.9M

Notes. All figures are annual flows per unit of disruption intensity d, expressed in 2022 USD. Marginal harm is a weighted average over the nineteen Western economies, not a bloc total. Leverage is the ratio of GDP-weighted marginal harm per Western economy to Iran's own marginal gain. Full per-country detail is available in the replication files.

technology for inflicting damage that vastly overpowers what the disruptor collects via economic welfare augmentation.

Second, greed alone cannot rationalize the observed intensity. Because (51) holds with equality at the intensity Iran actually chose, every assumed marginal cost maps to exactly one value of λ:

$$\lambda = \frac{\kappa'(d^*) - \bar{\kappa}'}{H'(d^*)}. \tag{54}$$

Pure greed, $\lambda = 0$, is therefore consistent with the data at one precise cost: $\kappa'(d^*) = \bar{\kappa}'$, between \$48 and \$78 million, depending on the value of σ. Any higher cost requires $\lambda > 0$, since a purely greedy Iran would have ceased disrupting at a lower intensity than what was observed. Costs *below* that value are equally inconsistent: with κ convex and u_1' locally flat, a greedy Iran facing cheaper disruption would have found it optimal to push *harder* than it did.

That knife-edge cost is implausibly low. Like every magnitude in Table 5, $\bar{\kappa}'$ is an annual figure; at $\sigma = 3.5$ it amounts to 0.06 percent of Iran's annual trade turnover. Moreover, because the model maintains $\kappa(0) = 0$ and $\kappa'' > 0$, the *total* annual cost of sustaining the campaign at its observed intensity satisfies $\kappa(1) = \int_0^1 \kappa'(s)\,ds < \kappa'(1)$, providing an upper bound for total annual expenditure. Two clarifications are in order. The marginal cost $\kappa'(d^*)$ is a derivative with respect to the unitless intensity index d, not with respect to dollars spent: it is the dollar cost of a marginal proportional scale-up of the observed campaign. And nothing here requires κ to be linear or otherwise parameterized: the bound uses only $\kappa(0) = 0$ and convexity, and because it can be written $\kappa(d^*) < d^*\kappa'(d^*) = d^*u_1'(d^*)$, it is invariant to how d is normalized. Under the greed-only

scenario, then, the entire campaign (munitions, targeting infrastructure, and the diplomatic and military exposure it entailed) must have cost Iran less than $64 million per year at $\sigma = 3.5$, and less than $78 million at any σ in the calibrated band. Any evidence that the campaign cost more than this per year thus refutes $\lambda = 0$.

An external anchor for κ'. Such evidence exists. Thomas Juneau, an Iran specialist at the University of Ottawa whose work is the principal academic account of the Iran–Houthi relationship (Juneau, 2016, 2024), places Iran's total *annual* support to the Houthi movement in the "low hundreds of millions of dollars" (Juneau, 2025), a range usually reported as $100 to $300 million (Nada, 2024).[19] Because $\bar{\kappa}'$ is likewise an annual figure, the two are directly comparable. Even the lower end of Juneau's range exceeds $\bar{\kappa}'$ by a factor of 1.3 to 2.1 depending on σ, while the upper end exceeds it by a factor of 3.8 to 6.3. Greed alone cannot rationalize the observed disruption at *any* σ in this paper's calibrated band: a strictly positive geopolitical weight is required to reconcile Iran's revealed behavior with a marginal cost anywhere near what Iran is independently understood to spend. The comparison is sharpest in the case most favorable to greed. The lowest elasticity in the band, $\sigma = 2$, delivers the largest knife-edge cost, $78.4 million; the lowest external estimate is $100 million. Even this pairing leaves Iran's support 28 percent above the greed-only bound and implies $\lambda = 0.0027$ under GDP weights and 0.0021 under equal weights (Table 6). No combination of elasticity, cost estimate, and weighting scheme considered here delivers $\lambda \leq 0$, and the comparison uses only the pecuniary flow, before any of the in-kind support documented next.

What Iran spends, and has spent. The dollar anchor captures only part of Iran's commitment, which intelligence and United Nations sources document along four margins. *Duration.* A United Nations panel of experts found that Iranian arms shipments to the Houthis by sea could be traced back to at least 2009, anchored by Yemen's 2013 seizure of the Iranian vessel *Jihan* (Agence France-Presse, 2015; Ferragamo and Robinson, 2026), and Juneau (2016) documents small-scale transfers from the same period. The Defense Intelligence Agency dates the IRGC Quds Force's supply of "an increasing range of sophisticated arms and training" to 2014 (Defense Intelligence Agency, 2024),

[19] Juneau's figure measures Iran's annual support to the Houthi movement rather than the marginal cost of the Red Sea disruption itself. As I argue below, it is best read as a lower bound on the cost Iran bears for the capability the campaign draws upon.

and support scaled up sharply once the Saudi-led intervention began in 2015 (Johnston et al., 2020; Juneau, 2024). By 2020 the U.S. Department of State assessed that Iran had spent "hundreds of millions of dollars" on the group since 2012 (U.S. Department of State, Iran Action Group, 2020). *Money and fuel.* A principal financial channel has been oil: the Houthis raised "perhaps tens of millions of dollars per month" by taxing fuel donated by Iran (DeLozier, 2019), and the U.S. Treasury describes the Iranian commodity-smuggling network of the IRGC-backed financier Sa'id al-Jamal as providing "tens of millions of dollars in revenue to the Houthis" and "a funding stream underpinning the Houthis' ongoing attacks against commercial shipping in the Red Sea" (U.S. Department of the Treasury, 2024). *Weapons.* Cargoes bound for the Houthis and interdicted in January 2024, in the middle of the campaign, carried anti-ship cruise missile propulsion, guidance, and warhead sections and ballistic-missile components with the characteristics of Iranian designs (United Nations Security Council, 2024). *Personnel and training.* The UN Panel of Experts on Yemen reports that 200 Houthi combatants trained at an IRGC naval facility on the Caspian Sea in early 2020 and 54 more at an IRGC base west of Tehran in late 2022, that Houthi fighters have trained at Hezbollah grounds in southern Lebanon since at least 2020, and that a Quds Force general serves as the Houthis' "jihad assistant," supervising their aerospace and naval forces (United Nations Security Council, 2024); earlier accounts placed IRGC and Hezbollah advisers in Yemen in the dozens to hundreds (Juneau, 2016). For scale, the Department of State put Iran's backing of Hezbollah, its most developed partner, at roughly $700 million per year (U.S. Department of State, Iran Action Group, 2018): support to the Houthis in the low hundreds of millions is a sizable fraction of Iran's flagship proxy commitment, not an incidental expense.

The range understates the true cost. Juneau's range most likely understands the disruption's true cost to Iran for three reasons. First, the Red Sea campaign is not a one-year undertaking but the culmination of the investment just described. The capability that closed the Bab el-Mandeb in 2023 was assembled over the preceding decade, and it was assembled for that purpose: the State Department records that, with Iranian weapons, training, and advisors, the Houthis fired anti-ship cruise missiles at vessels just north of the strait as early as October 2016, extending Iran's reach "beyond the Persian Gulf to the Red Sea" (U.S. Department of State, Iran Action Group, 2020). The Defense Intelligence Agency counted at least twenty intercepted Iranian arms shipments to the

Table 6: Implied λ at the independently estimated range of Iran's annual support to the Houthis

	$\kappa' =$ \$100M (lower bound)		$\kappa' =$ \$300M (upper bound)	
σ	GDP-wt.	Equal-wt.	GDP-wt.	Equal-wt.
2.0	0.0027	0.0021	0.0274	0.0219
3.5	0.0132	0.0095	0.0857	0.0614
5.0	0.0320	0.0220	0.1547	0.1066

$\lambda = [\kappa' - \bar{\kappa}'(\sigma)]/H'(d^*;\sigma)$, evaluated at the lower (\$100M) and upper (\$300M) ends of the independently estimated range of Iran's annual support (Juneau, 2025; Nada, 2024), GDP- and equal-weighted H'. All values strictly positive at every σ, confirming $\lambda > 0$ is required under this cost range regardless of weighting scheme or σ.

Houthis by January 2024, and interdictions have continued throughout the campaign, including a 750-ton shipment in June 2025 carrying anti-ship ballistic-missile seekers and cruise missiles (Toomey and Ben Taleblu, 2025). A single year's flow of support therefore prices only the marginal replenishment of a stock of missiles, drones, and know-how whose accumulated cost is far larger, and the marginal cost relevant to (51) must carry the shadow value of that stock. Second, a dollar figure cannot fully price support delivered in kind: the time of IRGC and Hezbollah trainers and advisers, the use of Iranian military facilities, and weapons and components transferred rather than sold. To the extent the estimate omits them, it understates κ'. Third, the pecuniary flow omits costs that Iran plainly bears but does not pay in dollars: the exposure of its personnel and its proxy to the U.S.-led strikes on Houthi launch sites that began in January 2024, the diplomatic cost of being held responsible for a global shipping disruption, and the risk of losing a partner it spent more than a decade building. Every one of these considerations raises the true κ' above the anchor and, through (54), raises the implied λ with it. Retaliation that raised Iran's own trade costs would work in the same direction from the other side of (54), lowering $\bar{\kappa}'$ (Proposition 3); the figures here set that channel to zero. The values in Table 6 are best read as lower bounds on Iran's geopolitical weight. Iran spent over a decade and hundreds of millions of dollars developing the capacity to close the strait: a marginal return of \$64 million a year is too small to exclusively justify that investment.

Table 6 reports the implied geopolitical weight across the independently estimated range. The values remain small in absolute terms, reflecting the extreme leverage documented above: because a small marginal push destroys tens of dollars of Western welfare for every dollar Iran gains, even a modest geopolitical weight rationalizes a real-world-sized cost. At the upper bound, $\lambda = 0.0857$

(GDP-weighted, $\sigma = 3.5$, $\kappa' = \$300\text{M}$) says Iran would sacrifice one dollar of its own welfare to destroy roughly \$12 of GDP-weighted Western welfare; at the lower bound, $\lambda = 0.0132$ implies a considerably steeper exchange rate of roughly \$76 of Western harm per dollar of Iran's own welfare. Both are relatively small compared to the harm inflicted, and both endpoints come from independent estimates of Iran's spending on the Houthis. Even at the corner most favorable to greed, the lowest elasticity paired with the lowest cost, the required weight is strictly positive. Because that estimate omits the decade of investment behind the campaign and every cost Iran bears in kind rather than in dollars, the weights it implies are floors, not point estimates.

6 Conclusion

Chokepoint disruptions have been studied as exogenous shocks: accidents of geography or geopolitics that raise trade costs and reduce welfare. This paper treats a large chokepoint disruption as the choice it was, and asks what makes such a choice worthwhile. The $(N + 1)$-country model of Section 3 characterizes the conditions under which a country benefits from raising the cost of trade between others and delivers three analytical results. Disruption is profitable even for a purely self-interested actor, because the demand it displaces from the disrupted bloc reallocates toward every unharmed country and improves their terms of trade. The gain scales quadratically in the size of the disrupted corridor, so that severing a single trade relationship is negligible in a large world while disrupting an artery is not. And the disruptor's payoff is a terms-of-trade gain with a predicted sign, a predicted incidence across bystanders and bloc members, and a predicted scaling in corridor size, so that the incentive to disrupt is a testable object rather than an assumption.

Taking the model to the 2023–24 Red Sea crisis required a measure of distance that responds to infrastructure. Logistical distance, the true kilometer length of the least-cost freight route over a global network of 578,298 road and maritime nodes, delivers a continuous, geography-determined measure of each country pair's exposure to closure of the Bab el-Mandeb Strait, validated on two independent margins: the dyads whose routes lengthen materially carry a share of world trade matching outside estimates of the Suez corridor's share, and the treatment is essentially invariant to the relative freight costs used to select routes. The estimates that follow from it are, to my knowledge, the first to exploit continuous, geography-determined variation in exposure to the crisis.

Three empirical findings emerge. Trade fell on exposed country pairs, with the response concentrated among Western-aligned pairs and in the months immediately following the onset of the campaign, consistent with a transitory rerouting adjustment cost layered on the permanent cost of the longer route. In general equilibrium, welfare losses are increasing in exposure: roughly two thirds of countries lose, the corridor-dependent economies of the Gulf, the Horn of Africa, and South Asia lose the most, and the minority of countries whose trade never relied on the strait gain through the terms-of-trade channel the theory predicts. Iran, the disruptor, gains almost nothing: roughly five thousandths of a percent of expenditure.

That last finding is the key to the paper's final result. Because the disruptor's first-order condition binds with equality at the intensity Iran actually chose, pure self-interest rationalizes the observed campaign at exactly one marginal cost, and convexity of the cost of disruption bounds the annual cost of sustaining the campaign under pure self-interest at $64 million at the central trade elasticity, and at $78 million across the calibrated range. Independent estimates place Iran's annual support for the Houthi movement at $100 to $300 million, between roughly 1.3 and 6.3 times that bound, and that flow is only the latest installment of an investment in the Houthis' capacity to close the strait dating to at least 2009. Greed alone cannot rationalize the disruption even in the configuration most favorable to it, the lowest trade elasticity paired with the lowest cost estimate, and before counting the weapons, fuel, and training Iran supplies in kind. The geopolitical weight required to do so is nonetheless modest, precisely because the disruption is so highly leveraged: at the margin, further disruption costs the GDP-weighted average Western economy between roughly 34 and 103 times more than it gains Iran, so that even a slight taste for harming an adversary rationalizes a campaign of real-world cost. The Red Sea episode thus reveals an asymmetric technology. Access to a chokepoint confers an ability to harm that vastly exceeds the ability to gain, and the private return to exercising it is so small that the exercise itself is evidence of motive.

The framework is portable. The same machinery—a network-based exposure measure, a gravity block estimated from non-strategic variation, and an optimality condition inverted at the observed intensity—applies to any actor with access to any chokepoint. Severing the Strait of Hormuz, the Strait of Malacca, the Turkish Straits, or the Panama Canal in the routing network, recomputing exposure, and evaluating the adjacent state's marginal gain and leverage would rank the world's

chokepoints by their value as instruments of coercion, and would identify the actors for whom that instrument is cheapest to use. The corridor-scaling result implies that the ranking is dominated by a handful of arteries, which is where deterrence and insurance resources should concentrate. The same tools speak to the other side of the ledger: because bystanders share the disruptor's terms-of-trade gain without bearing its cost, the model explains why coalitions to reopen a corridor are hard to assemble and predicts which countries will free-ride. Logistical distance, finally, is applicable well beyond conflict. Any change to the world's transportation infrastructure—the opening of Arctic routes, the expansion of a canal, the closure of a port, the imposition of a sanctions regime on a shipping lane—alters least-cost routes in a way that great-circle distance cannot register, and the measure developed here provides the exposure variable such questions require. Extending the quantitative model to multiple sectors with input-output linkages is a natural next step; it would sharpen the magnitudes reported here without altering the mechanism at the paper's core. Retaliation, finally, cuts in the paper's favor: a target that raises the disruptor's own trade costs by a fraction of the disruption it suffers extinguishes the economic motive while leaving the geopolitical one intact, so a disruption sustained in the face of retaliation is further evidence of grievance over greed.

References

Agence France-Presse. Iran arming Yemen's Houthis since 2009: U.N. report. Reported by Al Arabiya English, May 1, https://english.alarabiya.net/News/middle-east/2015/05/01/Iran-arming-Yemen-s-Houthis-since-2009-U-N-report, 2015.

Treb Allen and Costas Arkolakis. Trade and the topography of the spatial economy. *Quarterly Journal of Economics*, 129(3):1085–1140, 2014.

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

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

James E. Anderson and Eric van Wincoop. Trade costs. *Journal of Economic Literature*, 42(3): 691–751, 2004.

Costas Arkolakis, Arnaud Costinot, and Andrés Rodríguez-Clare. New trade models, same old gains? *American Economic Review*, 102(1):94–130, 2012.

AXSMarine. Nearly 300 less ships through Suez canal through Red Sea crisis. AXSMarine Red Sea Crisis tracker, 2024. https://public.axsmarine.com/blog/nearly-300-less-ships-through-suez-canal-through-red-sea-crisis.

Scott L. Baier and Jeffrey H. Bergstrand. Do free trade agreements actually increase members' international trade? *Journal of International Economics*, 71(1):72–95, 2007.

Richard Baldwin and Daria Taglioni. Gravity for dummies and dummies for gravity equations. NBER Working Paper 12516, 2006.

John Sturm Becko. A theory of economic sanctions as terms-of-trade manipulation. *Journal of International Economics*, 150:103898, 2024.

John Sturm Becko, Gene M. Grossman, and Elhanan Helpman. Optimal tariffs with geopolitical alignment. NBER Working Paper 34108, National Bureau of Economic Research, 2025.

Tibor Besedeš, Jing Chu, and Antu Panini Murshid. Fly the unfriendly skies: The role of transport costs in gravity models of trade. *Journal of International Economics*, 152:103994, 2024.

Bloomberg News. Houthis tell China, Russia their ships won't be targeted. Bloomberg, March 21, 2024.

Christoph E. Boehm, Andrei A. Levchenko, and Nitya Pandalai-Nayar. The long and short (run) of trade elasticities. *American Economic Review*, 113(4):861–905, 2023.

Ingo Borchert, Mario Larch, Serge Shikher, and Yoto V. Yotov. Globalization, trade, and inequality: Evidence from a new database. USITC Working Paper 2024-11-A, United States International Trade Commission, 2024.

Giulia Brancaccio, Myrto Kalouptsidi, and Theodore Papageorgiou. Rigidities in transportation and supply chain disruptions. *AEA Papers and Proceedings*, 115:543–550, 2025.

Lorenzo Caliendo and Fernando Parro. Estimates of the trade and welfare effects of NAFTA. *Review of Economic Studies*, 82(1):1–44, 2015.

Thomas Chaney. Distorted gravity: The intensive and extensive margins of international trade. *American Economic Review*, 98(4):1707–1721, 2008.

Christopher Clayton, Matteo Maggiori, and Jesse Schreger. A framework for geoeconomics. *Econometrica*, 94(1):105–136, 2026.

CNBC. As tankers divert from Red Sea, booming US oil gets another boost. January 23, https://www.cnbc.com/2024/01/23/more-tankers-divert-from-red-sea-and-its-another-boost-for-us-oil.html, 2024.

Arnaud Costinot and Andrés Rodríguez-Clare. Trade theory with numbers: Quantifying the consequences of globalization. In *Handbook of International Economics*, volume 4, pages 197–261. Elsevier, 2014.

Defense Intelligence Agency. Iran: Enabling Houthi attacks across the Middle East. U.S. Department of Defense, https://www.dia.mil/Portals/110/Documents/News/Military_Power_Publications/Iran_Houthi_Final2.pdf, February 2024.

Robert Dekle, Jonathan Eaton, and Samuel Kortum. Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Papers*, 55(3):511–540, 2008.

Elana DeLozier. In damning report, UN panel details war economy in Yemen. Policy Analysis, Washington Institute for Near East Policy, https://www.washingtoninstitute.org/policy-analysis/damning-report-un-panel-details-war-economy-yemen, January 2019.

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

Jason Dunn and Fernando Leibovici. Shipping disruptions in the Red Sea: Ripples across the globe. On the Economy, Federal Reserve Bank of St. Louis, February 15, https://www.stlouisfed.org/on-the-economy/2024/feb/shipping-disruptions-red-sea-ripples-globe, 2024.

Jonathan Eaton and Samuel Kortum. Technology, geography, and trade. *Econometrica*, 70(5): 1741–1779, 2002.

Jonathan Eaton, Samuel Kortum, and Sebastian Sotelo. International trade: Linking micro and macro. *Advances in Economics and Econometrics: Tenth World Congress*, 2:329–370, 2013.

Pablo D. Fajgelbaum, Pinelopi K. Goldberg, Patrick J. Kennedy, Amit K. Khandelwal, and Daria Taglioni. The US–China trade war and global reallocations. *American Economic Review: Insights*, 6(2):295–312, 2024.

Mariel Ferragamo and Kali Robinson. Iran's support of the Houthis: What to know. Council on Foreign Relations, 2026. URL https://www.cfr.org/articles/irans-support-houthis-what-know. Updated April 10, 2026.

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

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

Blaize Giangiulio, Peter Herman, Serge Shikher, and Yoto V. Yotov. Logistical distance. Working paper, 2025.

Reuven Glick and Alan M. Taylor. Collateral damage: Trade disruption and the economic impact of war. *Review of Economics and Statistics*, 92(1):102–127, 2010.

Ross Griffith. Weight of freight: Panamax locks demand, auction prices jump. Argus Media, February 2024. URL https://www.argusmedia.com/en/news-and-insights/market-opinion-and-analysis-blog/weight-of-freight-panamax-locks-demand-auction-prices-jump. 24 February 2024.

Keith Head and Thierry Mayer. Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of International Economics*, volume 4, pages 131–195. Elsevier, 2014.

Elhanan Helpman, Marc Melitz, and Yona Rubinstein. Estimating trade flows: Trading partners and trading volumes. *Quarterly Journal of Economics*, 123(2):441–487, 2008.

Albert O. Hirschman. *National Power and the Structure of Foreign Trade*. University of California Press, Berkeley, 1945.

Robert Holzer and Andy Cichon. U.S. Navy year in review. *U.S. Naval Institute Proceedings*, 151 (3), March 2025.

International Energy Agency. Suez canal – factsheet. IEA, Paris, 2023. https://www.iea.org/.

Trevor Johnston, Matthew Lane, Abigail Casey, Heather J. Williams, Ashley L. Rhoades, James Sladden, Nathan Vest, Jordan R. Reimer, and Ryan Haberman. Could the Houthis be the next Hizballah? Iranian proxy development in Yemen and the future of the Houthi movement. Research Report RR-2551-A, RAND Corporation, 2020.

Thomas Juneau. Iran's policy towards the Houthis in Yemen: A limited return on a modest investment. *International Affairs*, 92(3):647–663, 2016.

Thomas Juneau. How war in Yemen transformed the Iran–Houthi partnership. *Studies in Conflict & Terrorism*, 47(3):278–300, 2024.

Thomas Juneau. What does a weakened Iran mean for the Houthis? The Yemen Review, April–June 2025, Sana'a Center for Strategic Studies, https://sanaacenter.org/the-yemen-review/april-june-2025/25031, July 2025.

Mario Larch, Serge Shikher, and Yoto V. Yotov. Estimating gravity equations: Theory implications, econometric developments, and practical recommendations. Working Paper 2025001, Center for Global Policy Analysis, LeBow College of Business, Drexel University, 2025.

Philippe Martin, Thierry Mayer, and Mathias Thoenig. Make trade not war? *Review of Economic Studies*, 75(3):865–900, 2008.

Thierry Mayer and Soledad Zignago. Notes on CEPII's distances measures: The GeoDist database. Working Paper 2011-25, CEPII, 2011.

Cathrin Mohr and Christoph Trebesch. Geoeconomics. *Annual Review of Economics*, 17:563–587, 2025.

Garrett Nada. Houthis & Iran. The Iran Primer, U.S. Institute of Peace and Woodrow Wilson International Center for Scholars, https://www.wilsoncenter.org/article/houthis-iran, 2024.

Robert Perkins. US diesel exports to Europe hit record high despite lackluster demand. S&P Global Commodity Insights, July 24, https://www.spglobal.com/energy/en/news-research/latest-news/crude-oil/072424-us-diesel-exports-to-europe-hit-record-high-despite-lackluster-demand, 2024.

Mahvash Saeed Qureshi. Trade and thy neighbor's war. *Journal of Development Economics*, 105: 178–195, 2013.

E. G. Ravenstein. The laws of migration. *Journal of the Statistical Society of London*, 48(2): 167–235, 1885.

J. M. C. Santos Silva and Silvana Tenreyro. The log of gravity. *Review of Economics and Statistics*, 88(4):641–658, 2006.

Ina Simonovska and Michael E. Waugh. The elasticity of trade: Estimates and evidence. *Journal of International Economics*, 92(1):34–50, 2014.

Bridget Toomey and Behnam Ben Taleblu. Post-12-day war, Iran continues to invest in the Houthis. Foundation for Defense of Democracies, July 2025. URL https://www.fdd.org/analysis/2025/07/22/post-12-day-war-iran-continues-to-invest-in-the-houthis/.

Kristian Coates Ulrichsen and Jim Krane. Maritime chokepoints and risks to global shipping and energy security. Working paper, Baker Institute for Public Policy, Rice University, March 2026. URL https://www.bakerinstitute.org/sites/default/files/2026-03/20260316-Maritime%20Chokepoints.pdf.

UNCTAD. Review of maritime transport 2024: Navigating maritime chokepoints. Technical report, United Nations Conference on Trade and Development, Geneva, 2024.

United Nations Security Council. Final report of the panel of experts on Yemen established pursuant to security council resolution 2140 (2014). UN Document S/2024/731, United Nations, October 2024.

U.S. Department of State, Iran Action Group. Outlaw regime: A chronicle of Iran's destructive activities. U.S. Department of State; summarized in The Iran Primer, U.S. Institute of Peace, https://iranprimer.usip.org/blog/2018/oct/04/state-department-report-1-iran%E2%80%99s-support-terrorism, September 2018.

U.S. Department of State, Iran Action Group. Outlaw regime: A chronicle of Iran's destabilizing activity. 2020 edition, U.S. Department of State, September 2020. URL https://www.state.gov/wp-content/uploads/2020/09/Outlaw-Regime-2020-A-Chronicle-of-Irans-Destabilizing-Activity.pdf.

U.S. Department of the Treasury. Treasury maintains pressure on Houthi illicit shipping and finance schemes. Press release JY2469, https://home.treasury.gov/news/press-releases/jy2469, July 2024.

U.S. Energy Information Administration. The Bab el-Mandeb strait is a strategic route for oil and natural gas shipments. Today in Energy, August 27, 2019. https://www.eia.gov/todayinenergy/detail.php?id=41073.

Yoto V. Yotov, Roberta Piermartini, José-Antonio Monteiro, and Mario Larch. *An Advanced Guide to Trade Policy Analysis: The Structural Gravity Model*. WTO and UNCTAD, Geneva, 2016.

A Mathematical Derivations

A.1 The Consumer's Problem and the CES Demand System

Fix a country n. Since quasilinear utility is linear in the numeraire, any true optimum yields the chosen aggregate C_n at minimum cost; otherwise, reallocating differentiated-good spending and pocketing the savings in numeraire would unambiguously increase utility. The first stage can thus be reworked as an expenditure minimization,

$$e(p_n, C_n) \equiv \min_{\{c_{ni}\}} \sum_i p_{ni} c_{ni} \quad \text{s.t.} \quad \Big[\sum_i \beta_i c_{ni}^{\frac{\sigma-1}{\sigma}}\Big]^{\frac{\sigma}{\sigma-1}} \geq C_n. \tag{55}$$

With multiplier μ on the constraint, the first-order condition for c_{ni} is $p_{ni} = \mu\, \beta_i\, C_n^{1/\sigma} c_{ni}^{-1/\sigma}$, which is equivalent to

$$c_{ni} = \Big(\frac{\mu \beta_i}{p_{ni}}\Big)^{\sigma} C_n. \tag{56}$$

Substituting (56) into the constraint,

$$C_n^{\frac{\sigma-1}{\sigma}} = \sum_i \beta_i \Big(\frac{\mu\beta_i}{p_{ni}}\Big)^{\sigma-1} C_n^{\frac{\sigma-1}{\sigma}} = \mu^{\sigma-1}\Big[\sum_i \beta_i^{\sigma} p_{ni}^{1-\sigma}\Big] C_n^{\frac{\sigma-1}{\sigma}} \quad \Longrightarrow \quad \mu = \Big[\sum_i \beta_i^{\sigma} p_{ni}^{1-\sigma}\Big]^{\frac{1}{1-\sigma}} = P_n,$$

the price index of the main text: the multiplier equals the unit cost of the bundle, so $e(p_n, C_n) = P_n C_n$. Substituting $\mu = P_n$ back into (56) gives the conditional demands and, multiplying by p_{ni}, the trade values x_{ni} of Equation (4); dividing by total spending $\sum_i x_{ni} = P_n^{\sigma} C_n \cdot P_n^{1-\sigma} = P_n C_n$ gives the shares (5). The upper tier (6) follows from substituting $x_0^n = I_n - P_n C_n$ into utility and maximizing over C_n. For the bundle-size elasticity (7), differentiating $v'(C_n) = P_n$ gives $v''(C_n)\, dC_n = dP_n$, so $\frac{d\ln C_n}{d\ln P_n} = \frac{P_n}{C_n}\frac{1}{v''(C_n)} < 0$ and $\varepsilon = -\frac{P_n}{C_n v''(C_n)} \geq 0$. Under the log tier $v'(C) = E/C$, so $C_n = E/P_n$, spending is $P_n C_n = E$, and $\varepsilon = 1$.

A.2 Proof of Lemma 1

From (9)–(10), the response of a single shipment to a single mill price is

$$\frac{\partial \ln m_{ni}}{\partial \ln p_j} = \begin{cases} -\sigma + \gamma\, s_{nj}, & j = i, \\ \gamma\, s_{nj}, & j \neq i. \end{cases} \tag{57}$$

Converting to levels ($\partial m_{ni}/\partial \ln p_j = m_{ni}\, \partial \ln m_{ni}/\partial \ln p_j$) and summing over destinations n yields the entries shown in (16).

Off-diagonals. $\tilde{J}_{ij} = \gamma \sum_n m_{ni} s_{nj} > 0$, since every factor is strictly positive.

Diagonal. Because each country purchases at least two goods, $s_{ni} < 1$ for every n. This implies both that $\sum_n m_{ni} s_{ni} < \sum_n m_{ni} = X_i$ and that $\tilde{J}_{ii} < -\sigma X_i + \gamma X_i = -\varepsilon X_i \leq 0$.

Row sums. Because the inner sum $\sum_j s_{nj} = 1$ does not depend on the outer index, it can be evaluated first:

$$\sum_j \tilde{J}_{ij} = -\sigma X_i + \gamma \sum_n m_{ni} \underbrace{\sum_j s_{nj}}_{=1} = -\sigma X_i + \gamma X_i = -(\sigma - \gamma) X_i = -\varepsilon X_i \leq 0. \tag{58}$$

A matrix with strictly negative diagonal, positive off-diagonals, and non-positive row sums is minus a row-diagonally-dominant M-matrix; invertibility and $-\tilde{J}^{-1} \geq 0$ entrywise follow from standard M-matrix theory, and uniqueness of equilibrium follows from gross substitutes. Under the log tier $\varepsilon = 1 > 0$, the dominance is strict, so no price normalization is required: the numeraire supplies the level. □

A.3 The Benchmark Jacobian Numbers D and O

Under Assumption 1, the price-index denominator is $\Phi = 1 + N\theta$, yielding the shares $s_0 = 1/(1+N\theta)$ and $s_1 = \theta/(1+N\theta)$. Shipments are $m_{ii} = a \equiv P^\sigma C$ at home and $m_{ni} = \theta a$ abroad, so $X_i = a(1+N\theta)$ for every good.

Diagonal. The destinations of good i weight good i by s_0 at home (shipment a) and s_1 at each of the N foreign destinations (shipment θa each):

$$\sum_n m_{ni} s_{ni} = a\, s_0 + N\, \theta a\, s_1 = a\, \frac{1 + N\theta^2}{1 + N\theta} \quad \Longrightarrow \quad \frac{\tilde{J}_{ii}}{a} = -\sigma(1 + N\theta) + \gamma\, \frac{1 + N\theta^2}{1 + N\theta} \equiv D.$$

Off-diagonal. For $j \neq i$, the destinations of good i weight good j as follows: country i itself

(shipment a; good j is foreign to it, share s_1); country j (shipment θa; good j is its home good, share s_0—the one and only s_0 term, since exactly one country calls good j home); and the remaining $N-1$ countries (shipment θa; share s_1):

$$\sum_n m_{ni} s_{nj} = a\, s_1 + \theta a\, s_0 + (N-1)\,\theta a\, s_1 = a\,\frac{\theta\big[2+(N-1)\theta\big]}{1+N\theta} \quad \Longrightarrow \quad \frac{\tilde{J}_{ij}}{a} = \gamma\,\frac{\theta\big[2+(N-1)\theta\big]}{1+N\theta} \equiv O.$$

Row-sum check. The row-summed numerator of $D+NO$ is $1+N\theta^2+N\theta[2+(N-1)\theta] = 1+2N\theta+N^2\theta^2 = (1+N\theta)^2$, so $D+NO = (\gamma-\sigma)(1+N\theta) = -\varepsilon(1+N\theta)$, consistent with (58) since $X_i = a(1+N\theta)$. Writing the resulting matrix as $(D-O)I + O\,\mathbf{1}\mathbf{1}^\top$ reproduces (17): the diagonal is $D-O+O=D$ and the off-diagonal is O.

A.4 The Direct-Effects Vector and the Constant k

Every disrupted link carries $\partial \ln \tau_{ni}/\partial d = \delta'$, so by (8), at fixed prices a shipment responds through two channels:

$$\left.\frac{\partial \ln m_{ni}}{\partial d}\right|_p = (1-\sigma)\,\delta' \cdot \mathbf{1}\{(n,i)\in\mathcal{D}\} + \gamma\left.\frac{\partial \ln P_n}{\partial d}\right|_p, \tag{59}$$

namely the own-link pullback and the price-index reallocation.

Step 1: index changes. At fixed prices, $\partial \ln P_n/\partial d|_p = \sum_{i:(n,i)\in\mathcal{D}} s_{ni}\,\delta'$. A bloc member has $m-1$ disrupted import links, each with share s_1; a bystander has none:

$$\left.\frac{\partial \ln P_n}{\partial d}\right|_p = \begin{cases} (m-1)\,s_1\,\delta', & n\in\mathcal{M}, \\ 0, & n\notin\mathcal{M}. \end{cases} \tag{60}$$

Step 2: bystander goods. For $i\notin\mathcal{M}$, no link into good i is disrupted, so only the index channel operates, at the m bloc destinations, each importing θa:

$$(Z_d)_i = \sum_{n\in\mathcal{M}} m_{ni}\,\gamma\,(m-1)s_1\delta' = m\,\theta a\,\gamma\,(m-1)s_1\delta' = m(m-1)\,k\,a, \qquad k \equiv \gamma\, s_1\,\theta\,\delta', \tag{61}$$

which defines the scale constant k: elasticity (γ) times share exposure (s_1) times cost sensitivity

of shipments (θ) times push intensity (δ'). Under the log tier ($\gamma = \sigma - 1$) and $s_1 = \theta/(1 + N\theta)$, $k = (\sigma - 1)\theta^2\delta'/(1 + N\theta)$, as in (18).

Step 3: bloc goods. For $i \in \mathcal{M}$, both channels operate. The pullback: good i travels on $m - 1$ disrupted links, each carrying θa, contributing $(1 - \sigma)\delta'(m - 1)\theta a$. The reallocation: the m raised indices apply to shipments $m_{ii} = a$ and θa from each of the other $m - 1$ members, so

$$(Z_d)_i = (m-1)\,\delta'\,a\left\{(1-\sigma)\theta + \gamma s_1\big[1 + (m-1)\theta\big]\right\}.$$

Simplifying the brace with $\gamma = \sigma - 1$ and $s_1(1 + N\theta) = \theta$,

$$(1-\sigma)\theta+\gamma s_1[1+(m-1)\theta] = (\sigma-1)\,\theta\,\frac{-(1+N\theta)+1+(m-1)\theta}{1+N\theta} = -(\sigma-1)\,\theta\,\frac{(N-m+1)\theta}{1+N\theta} = -\,b\,\frac{k}{\delta'},$$

using $N - m + 1 = b$. Hence $(Z_d)_i = -b(m - 1)\,k\,a$ for bloc goods.

Step 4: Walras check. Summing over all goods, $b \cdot m(m - 1)ka + m \cdot [-b(m - 1)ka] = 0$: at fixed prices, the disruption creates no aggregate excess demand and purely reallocates. This uses $\varepsilon = 1$; for general v the sum is proportional to $1 - \varepsilon$. Stacking with bystanders first and negating gives (18).

A.5 Solving the Collapsed System

By the two-group symmetry, $\dot{p}_i = \hat{p}_B$ for the b bystander goods and $\dot{p}_i = \hat{q}$ for the m bloc goods, so one representative row of each type of $\tilde{J}\dot{p} = -Z_d$ suffices. A bystander row of the circulant carries D on itself and O on the other $b - 1$ bystanders and m bloc goods; a bloc row carries D on itself and O on the b bystanders and $m - 1$ other bloc goods. Dividing by a and using the right-hand entries of (18) delivers (19)–(20).

Determinant. Expanding, with $m + b = N + 1$:

$$\begin{aligned}\Delta &= \big[D+(b-1)O\big]\big[D+(m-1)O\big] - mb\,O^2\\ &= D^2 + (m+b-2)\,DO + \big[(b-1)(m-1) - mb\big]O^2\\ &= D^2 + (N-1)\,DO + \big[\,bm - b - m + 1 - bm\,\big]O^2 = D^2 + (N-1)\,DO - NO^2\\ &= (D-O)\big(D+NO\big) \;>\; 0,\end{aligned}$$

the product of two strictly negative factors ($D-O<0$ by diagonal dominance; $D+NO<0$ by the row sum), which is (21).

Cramer's rule. For $\hat{p}_B$, replace the first column by the right-hand side:

$$\hat{p}_B = \frac{(m-1)k}{\Delta}\det\begin{pmatrix} -m & mO \\ b & D+(m-1)O \end{pmatrix} = \frac{(m-1)k\cdot(-m)\big[D+(m-1)O+bO\big]}{\Delta}$$
$$= \frac{-m(m-1)k\,(D+NO)}{(D-O)(D+NO)} = \frac{-m(m-1)\,k}{D-O},$$

using $(m-1)+b=N$. For $\hat{q}$, replace the second column:

$$\hat{q} = \frac{(m-1)k}{\Delta}\det\begin{pmatrix} D+(b-1)O & -m \\ bO & b \end{pmatrix} = \frac{(m-1)k\cdot b\big[D+(b-1)O+mO\big]}{\Delta} = \frac{b(m-1)\,k}{D-O}.$$

Since $D-O<0$, these are the signed expressions of Proposition 1.

A.6 Welfare: The Envelope Step and the Closed Forms

Differentiating (13) in d along the equilibrium path, $u_n'(d) = R_n'(d) - E\,\hat{P}_n = y_n p_n' - E\hat{P}_n = R_n\hat{p}_n - E\hat{P}_n$, using $p_n=1$ at the benchmark, where $\hat{P}_n$ includes the direct corridor-cost term (60) for bloc members. Adding and subtracting $E\hat{p}_n$ gives the two-term form (25). Equivalently, apply the envelope theorem to the general u_n: the induced variation in C_n cancels because $v'(C_n)=P_n$.

The disruptor, or any bystander. At the benchmark $\mathrm{NX}_n = 0$, so $u_1' = E(\hat{p}_1 - \hat{P}_1)$, with $\hat{P}_1 = s_0\hat{p}_B + (b-1)s_1\hat{p}_B + m\,s_1\hat{q}$ and no direct term. Using $s_0 + Ns_1 = 1$ and $N-(b-1)=m$,

$$\hat{p}_1 - \hat{P}_1 = \hat{p}_B\big[1 - s_0 - (b-1)s_1\big] - ms_1\hat{q} = m\,s_1\big(\hat{p}_B - \hat{q}\big).$$

Since $Es_1 = a(1+N\theta)\cdot\theta/(1+N\theta) = \theta a = \bar{x}$,

$$u_1'(d) = \bar{x}\, m\,(\hat{p}_B - \hat{q}) = \bar{x}\,\frac{m(m-1)(m+b)\,k}{|D-O|} = \bar{x}\,\frac{m(m-1)(N+1)\,k}{|D-O|} > 0,$$

where the middle step substitutes Proposition 1: $\hat{p}_B - \hat{q} = (m+b)(m-1)k/|D-O|$. This is (26).

A bloc member. For $n \in \mathcal{M}$, again $\mathrm{NX}_n = 0$, and now the direct term appears in $\hat{P}_n = s_0\hat{q} + (m-1)s_1\hat{q} + b\,s_1\hat{p}_B + (m-1)s_1\delta'$. Using $1 - s_0 - (m-1)s_1 = b\,s_1$,

$$\hat{p}_n - \hat{P}_n = b\,s_1(\hat{q} - \hat{p}_B) - (m-1)s_1\delta' \quad\Longrightarrow\quad u_n'(d) = \bar{x}\Big[b(\hat{q} - \hat{p}_B) - (m-1)\delta'\Big] < 0,$$

which is (27): both terms strictly negative.

A.7 Proof of Corollary 1

As $N \to \infty$ with θ fixed: from (18), $k = (\sigma-1)\theta^2\delta'/(1+N\theta) \sim (\sigma-1)\theta\delta'/N$; from (17), $O \to \gamma\theta$ (bounded) while $D \sim -\sigma N\theta$, so $|D-O| \sim \sigma N\theta$.

Fixed bloc. With m constant, $\hat{p}_B = m(m-1)k/|D-O| \sim m(m-1)\,\frac{(\sigma-1)\theta\delta'/N}{\sigma N\theta} = O(N^{-2})$.

Fixed fraction. With $m = \alpha N$, so that $m(m-1) \sim \alpha^2 N^2$,

$$\hat{p}_B \sim \frac{\alpha^2 N^2\cdot(\sigma-1)\theta\delta'/N}{\sigma N\theta} = \frac{\alpha^2(\sigma-1)}{\sigma}\,\delta' > 0,$$

the θ terms canceling exactly, so the limit is independent of the baseline trade cost. □

A.8 The Augmented ACR Decomposition

With $\beta_i = 1$, the own share is $s_{nn} = S_n/\Phi_n$ with $S_n = p_n^{1-\sigma}$ and $\Phi_n = P_n^{1-\sigma}$. Log-differentiating in d,

$$\hat{s}_{nn} = (1-\sigma)\hat{p}_n - (1-\sigma)\hat{P}_n \quad\Longrightarrow\quad \hat{p}_n - \hat{P}_n = -\frac{\hat{s}_{nn}}{\sigma-1}.$$

Substituting into the two-term form (25) yields (35). At the balanced-trade benchmark, $\mathrm{NX}_n = 0$ and the disruptor's gain loads entirely on the own-share term: from Appendix A.6, $\hat{p}_1 - \hat{P}_1 = ms_1(\hat{p}_B - \hat{q}) > 0$, so $\hat{s}_{11} = (1-\sigma)(\hat{p}_1 - \hat{P}_1) < 0$. The disruptor's own-trade share falls as its export

price rises relative to its basket—the sufficient-statistic signature of its gain.

A.9 Retaliation: Proof of Proposition 3

Step 0: the collapse holds for any zero-sum shock. Write the benchmark Jacobian as $\tilde{J} = a[(D - O)I + O\,\mathbf{1}\mathbf{1}^\top]$ and sum the rows of $\tilde{J}\dot{p} = -Z_d$: $(D + NO)\,\mathbf{1}^\top\dot{p} = -\mathbf{1}^\top Z_d/a$. Whenever the direct effect sums to zero, $\mathbf{1}^\top\dot{p} = 0$, the rank-one term drops out, and $\dot{p} = -Z_d/[a(D-O)]$ for *every* good, whatever the cross-sectional pattern of the shock. Equation (22) is the special case of a two-group shock. Under the log tier the combined disruption-and-retaliation shock is zero-sum (Step 2), so the retaliation closed forms require only the direct-effect vector.

Step 1: index changes at fixed prices. Under (11) and (29), every intra-bloc link carries δ' and every disruptor–bloc link carries $r\delta'$. By (60), with s_1 the share of each foreign good,

$$\left.\frac{\partial \ln P_n}{\partial d}\right|_p = \begin{cases} m\,s_1\,r\,\delta', & n = 1, \\ 0, & n \in \mathcal{B}\setminus\{1\}, \\ (m-1+r)\,s_1\,\delta', & n \in \mathcal{M}, \end{cases}$$

where $\mathcal{B}$ is the set of b bystanders. The disruptor's index now moves because its imports from the bloc are dearer.

Step 2: the direct-effect vector. Apply (59) good by good, with $k = \gamma s_1\theta\delta'$, $1/s_1 = 1/\theta + N$, and $\gamma = \sigma - 1$. For the disruptor's good, demand comes from itself (a, index channel only), the other bystanders (unaffected), and the m bloc members (θa each, own-link pullback $-(\sigma-1)r\delta'$ plus their index channel):

$$(Z_d)_1 = a\,\gamma\,ms_1r\delta' + m\theta a\big[(1-\sigma)r\delta' + \gamma(m-1+r)s_1\delta'\big] = ka\big[m(m-1) - m\,r\,(N-1)\big].$$

For another bystander's good, only index channels operate, at the disruptor (θa) and the m bloc members:

$$(Z_d)_B = \theta a\,\gamma\,ms_1r\delta' + m\theta a\,\gamma(m-1+r)s_1\delta' = ka\big[m(m-1) + 2\,m\,r\big].$$

For a bloc good, demand comes from itself (a), the disruptor (θa, pullback $-(\sigma-1)r\delta'$ plus its

index), the other bystanders (unaffected), and the other $m-1$ members (θa each, pullback $-(\sigma-1)\delta'$ plus their index):

$$(Z_d)_{\mathcal{M}} = ka\Big[\tfrac{m-1+r}{\theta} - \tfrac{r}{s_1} + mr + (m-1)\big(-\tfrac{1}{s_1} + m - 1 + r\big)\Big] = -ka\big[b(m-1) + r\,(b-m)\big],$$

using $N+1-2m = b-m$. At $r=0$ the three entries reduce to (18). Summing with multiplicities 1, $b-1$, and m, the $m(m-1)$ terms give $m(m-1)[1+(b-1)-b] = 0$ and the r terms give $mr\,[-(N-1)+2(b-1)-(b-m)] = mr\,[-N+(b-1+m)] = 0$, using $b-1+m = N$: the shock is zero-sum. Step 0 then delivers (30).

Step 3: the disruptor's welfare. By the envelope step of Appendix A.6, $u_1' = E(\hat{p}_1 - \hat{P}_1)$ at balanced trade, with $\hat{P}_1 = s_0\hat{p}_1 + (b-1)s_1\hat{p}_B + ms_1\hat{q} + ms_1 r\delta'$, the last term the direct cost of dearer bloc imports. Using $1-s_0 = Ns_1$ and $\bar{x} = Es_1$,

$$u_1' = \bar{x}\big[N\hat{p}_1 - (b-1)\hat{p}_B - m\hat{q} - m\,r\,\delta'\big].$$

Substituting (30), the $m(m-1)$ terms collect to $m(m-1)(N+1)\,k/|D-O|$, which is (26), and the r terms to $mr\,[-N(N-1) - 2(b-1) + (b-m)]\,k/|D-O| = -m\,r\,(N-1)(N+1)\,k/|D-O|$, using $b-1+m = N$ once more. Hence

$$u_1'(d;r) = \bar{x}\,m(N+1)\frac{k}{|D-O|}\big[(m-1) - r(N-1)\big] - \bar{x}\,m\,r\,\delta' = u_1'(d;0)\Big[1 - r\frac{N-1}{m-1}\Big] - \bar{x}\,m\,r\,\delta',$$

which is (31). Setting it to zero gives (32); since the second term in the denominator is positive, $\bar{r} < (m-1)/(N-1)$, and $(m-1)/(N-1) < 1$ because $m \le N$.

Step 4: a bloc member's welfare. For $n \in \mathcal{M}$, $\hat{P}_n = s_0\hat{q} + s_1\hat{p}_1 + (b-1)s_1\hat{p}_B + (m-1)s_1\hat{q} + (m-1+r)s_1\delta'$, so

$$u_n' = \bar{x}\big[b\,\hat{q} - \hat{p}_1 - (b-1)\hat{p}_B - (m-1+r)\,\delta'\big].$$

At $r=0$ this is (27). The r terms in the bracket are $r\,[-b(b-m) + m(N-1) - 2(b-1)m]\,k/|D-O| - r\delta'$; substituting $N = b+m-1$, the square bracket equals $m^2 - b^2 = (N+1)(m-b)$, which gives (33). □

B Ricardian Foundations, the Extensive Margin, and Market Access

Section 3 delivers the disruptor's terms-of-trade gain in an Armington economy, where every country consumes a positive quantity of every other country's good by assumption. A natural concern is that the gain is an artifact of that love-of-variety structure: if bloc members were not obliged to buy the disruptor's good, demand displaced from the corridor might never reach it. This appendix addresses the concern in two steps. First, I re-derive the model on Ricardian foundations in the manner of Eaton and Kortum (2002), in which trade adjusts entirely on the *extensive* margin—each good is bought from a single lowest-cost source, and a trade-cost shock changes *which* goods a country imports from each partner rather than how much of each. I show that the equilibrium system, the comparative statics, and every result of Section 3 carry over verbatim. Second, I show that what the gain actually requires is not that bloc members consume the disruptor's good by assumption, but that the disruptor have *market access* to the bloc: a disruptor whose trade with the bloc is zero not only fails to gain but loses. The quantitative exercises of Section 4.4 compute Iran's gain through its observed 2022 trade shares, so the condition is verified in the data rather than imposed.

B.1 Environment

There are $N+1$ countries. Each is endowed with $L_n > 0$ units of labor and, as in Section 3.1, $\tilde{T}_n > 0$ units of a freely traded numeraire with price one. Labor produces a continuum of goods $\omega \in [0,1]$ under constant returns and perfect competition. Country i draws its efficiency in good ω, $z_i(\omega)$, independently across goods and countries from the Fréchet distribution $F_i(z) = \exp(-A_i z^{-\theta})$, with $A_i > 0$ the country's absolute advantage and $\theta > \eta - 1$ the inverse dispersion of productivity, where $\eta > 1$ is the elasticity of substitution across goods within the continuum.[20] Iceberg costs $\tau_{ni} \geq 1$ ($\tau_{nn} = 1$) are as before, and the disruption technology (11) is unchanged. The wage in country i is w_i, so the cost of delivering one unit of good ω from i to n is $w_i \tau_{ni} / z_i(\omega)$, and country n buys each good from the cheapest source. Preferences are quasilinear over the numeraire and a CES

[20] I write A_i for the Fréchet location parameter, rather than the customary T_i, because T_n denotes the numeraire endowment in the main text.

aggregate of the continuum,

$$U_n = x_0^n + v(C_n), \qquad C_n = \Big[\int_0^1 q_n(\omega)^{\frac{\eta-1}{\eta}}\,\mathrm{d}\omega\Big]^{\frac{\eta}{\eta-1}}, \tag{62}$$

with the log tier $v(C_n) = E\ln C_n$ maintained for the closed forms. Income is $I_n = \tilde{T}_n + w_n L_n$.

B.2 The Eaton–Kortum Demand System

Lemma 2 (Ricardian demand system)**.** *Let $\Phi_n \equiv \sum_k A_k(w_k\tau_{nk})^{-\theta}$. Then:*

(i) The fraction of goods country n buys from i equals n's expenditure share on i,

$$\pi_{ni} = \frac{A_i(w_i\tau_{ni})^{-\theta}}{\Phi_n}, \qquad \sum_i \pi_{ni} = 1, \tag{63}$$

and the exact price index of the continuum is $P_n = \gamma_\eta\,\Phi_n^{-1/\theta}$, with $\gamma_\eta \equiv \Gamma\big(\frac{\theta+1-\eta}{\theta}\big)^{1/(1-\eta)}$.

(ii) Under the log tier, spending on the continuum is $P_nC_n = E$, so the value of n's imports from i is $x_{ni} = \pi_{ni}E$. Under a general upper tier with bundle-size elasticity ε as in (7), $\mathrm{d}\ln(P_nC_n) = (1-\varepsilon)\,\mathrm{d}\ln P_n$.

(iii) Holding trade costs fixed,

$$\frac{\partial\ln\pi_{ni}}{\partial\ln w_j} = -\theta\,\mathbf{1}\{j=i\} + \theta\,\pi_{nj}, \qquad \frac{\partial\ln P_n}{\partial\ln w_j} = \pi_{nj}, \tag{64}$$

and, holding wages fixed, a disruption of intensity d changes shares by

$$\left.\frac{\partial\ln\pi_{ni}}{\partial d}\right|_w = -\theta\,\delta'\,\mathbf{1}\{(n,i)\in\mathcal{D}\} + \theta\sum_{k:(n,k)\in\mathcal{D}}\pi_{nk}\,\delta', \qquad \left.\frac{\partial\ln P_n}{\partial d}\right|_w = \sum_{k:(n,k)\in\mathcal{D}}\pi_{nk}\,\delta'. \tag{65}$$

Proof. (i) is Eaton and Kortum (2002), Theorems 1–3: the distribution of the lowest delivered price in n is $G_n(p) = 1-\exp(-\Phi_n p^\theta)$; conditional on being the lowest-cost supplier, every source faces the same price distribution, so the fraction of goods sourced from i equals i's share of n's spending; and the CES index of a Fréchet price distribution is $\gamma_\eta\Phi_n^{-1/\theta}$, finite when $\theta > \eta-1$. (ii) follows from the upper-tier first-order condition $v'(C_n) = P_n$ exactly as in Appendix A.1; the continuum's

internal structure is irrelevant to that step because C_n enters utility only through v. (iii) Log-differentiate (63): $\partial \ln \pi_{ni}/\partial \ln w_j = -\theta \mathbf{1}\{j = i\} - \partial \ln \Phi_n/\partial \ln w_j$ and $\partial \ln \Phi_n/\partial \ln w_j = -\theta \pi_{nj}$; the index elasticity is $-\frac{1}{\theta}\partial \ln \Phi_n/\partial \ln w_j = \pi_{nj}$. The disruption derivatives follow identically with $\partial \ln \tau_{nk}/\partial d = \delta'$ on $\mathcal{D}$ in place of $\partial \ln w_j$. □

The content of (i) is the extensive margin: when the corridor becomes costlier, a bloc importer does not buy a little less of every bloc good and a little more of every bystander good; it stops buying a set of goods of measure $\theta \pi_{nk} \delta' \, \mathrm{d}d$ from each bloc partner k and starts buying a set of measure $\theta \pi_{nj} (\sum_{k \in \mathcal{D}(n)} \pi_{nk}) \delta' \, \mathrm{d}d$ from each bystander j, the disruptor included. Aggregated, however, the shift in the *measure* of goods is the shift in *expenditure*, and (64)–(65) have exactly the form of (9)–(10) and (59).

B.3 Isomorphism with the Armington Economy

Proposition 4 (Isomorphism)**.** *Define the map $\sigma - 1 \mapsto \theta$, $p_i \mapsto w_i$, $y_i \mapsto L_i$, $\beta_i^{\sigma} \mapsto A_i$, $s_{ni} \mapsto \pi_{ni}$. Under this map the Ricardian economy of Appendix B.1 has the same market-clearing system, the same elasticity Jacobian $\tilde{J}$, the same direct-effect vector Z_d, and the same indirect utility (13) as the Armington economy of Section 3. Consequently Lemma 1, Propositions 1 and 2, Corollary 1, the first-order condition (28), and the augmented ACR decomposition (35) hold verbatim, with θ in place of $\sigma - 1$ and wages in place of mill prices.*

Proof. Market clearing in the Ricardian economy equates the wage bill to world spending on country i's goods, $w_i L_i = \sum_n x_{ni} = \sum_n \pi_{ni} P_n C_n$. The Armington condition (12), $y_i = \sum_n m_{ni}$, multiplied by the mill price p_i and using $x_{ni} = p_{ni} c_{ni} = p_i m_{ni}$, is $p_i y_i = \sum_n x_{ni} = \sum_n s_{ni} P_n C_n$: the two systems coincide in value form under the map. For the Jacobian, log-differentiate the value of shipments. In the Armington economy $x_{ni} = p_i \tau_{ni} m_{ni}$, so by (9)–(10), $\partial \ln x_{ni}/\partial \ln p_j = (1 - \sigma)\mathbf{1}\{j = i\} + \gamma s_{nj}$ with $\gamma = \sigma - \varepsilon$. In the Ricardian economy, by Lemma 2(ii)–(iii), $\partial \ln x_{ni}/\partial \ln w_j = -\theta \mathbf{1}\{j = i\} + (\theta + 1 - \varepsilon)\pi_{nj}$. Under $\theta = \sigma - 1$ these are identical, so $\tilde{J}$ coincides entry by entry (Appendix A.2). The same comparison of (65) with (59) shows that Z_d coincides. Indirect utility is $u_n = \tilde{T}_n + w_n L_n - E + E \ln E - E \ln P_n$, which is (13) with $R_n = w_n L_n$ and P_n the Ricardian index; the constant γ_η drops out of every change. Since all results of Section 3 are derived from $\tilde{J}$, Z_d, and (13) alone, they carry over. For the ACR decomposition, $\pi_{nn} = A_n w_n^{-\theta}/\Phi_n$ gives

$\hat{\pi}_{nn} = -\theta(\hat{w}_n - \hat{P}_n)$, the Ricardian counterpart of Appendix A.8 with θ in place of $\sigma - 1$. □

This is the Arkolakis et al. (2012) equivalence applied to a strategic trade-cost shock: because the Armington and Eaton–Kortum models share the gravity structure (34) and the import-demand system, they share every comparative static that depends only on observed shares and the trade elasticity, and the incentive to disrupt is one of them. In particular, the proof shows exactly which assumptions do the work. The gain does not come from love of variety (there is none in the Ricardian economy); it comes from *substitution*: a costlier corridor raises bloc members' price indices and shifts their sourcing toward whichever suppliers remain cheap, and every supplier that enjoys such a shift sees its factor price rise. The extensive margin is not an alternative to this mechanism; it is the mechanism's microfoundation in a Ricardian world. Numerically, solving the Ricardian economy of Appendix B.1 at the symmetric benchmark ($N + 1 = 12$, $m = 6$, $\theta = 3$, $\tau = 1.5$) reproduces the closed forms of Proposition 1 and (26) to eight significant digits ($\hat{w}_B = 0.11116043$, $\hat{w}_{\mathcal{M}} = -0.11116043$, $u_1'(0) = 0.09279479$ in units of E).

B.4 What the Gain Requires: Market Access

The concern that motivates this appendix survives the isomorphism in a sharper form. In both economies every bilateral share is strictly positive whenever trade costs are finite, so neither can speak to a disruptor that trades *nothing* with the bloc. The relevant question is therefore: how does the disruptor's gain depend on its own access to bloc markets, and what happens in the limit where that access vanishes? The augmented ACR decomposition gives the answer in general, without symmetry.

Proposition 5 (Market access)**.** *Let trade costs be arbitrary (not necessarily symmetric) and let the disruptor's trade be balanced,* $\mathrm{NX}_1 = 0$*. Then in either economy*

$$u_1'(d) = E \sum_{k \neq 1} \pi_{1k} \left(\hat{w}_1 - \hat{w}_k \right) = -\frac{E}{\theta} \hat{\pi}_{11}, \tag{66}$$

so the disruptor gains if and only if its wage rises relative to the import-share-weighted wage of its suppliers, equivalently if and only if its own-trade share falls. Suppose further that the disruptor trades with no bloc member, $\pi_{n1} = \pi_{1n} = 0$ *for all* $n \in \mathcal{M}$*, and that the remaining bystanders are*

symmetric with common wage response $\hat{w}_B > 0$. *Then* $\hat{w}_1 = c\,\hat{w}_B$ *with* $0 < c < 1$, *and*

$$u_1'(d) = E\,(1 - \pi_{11})\,(\hat{w}_1 - \hat{w}_B) < 0 : \tag{67}$$

a disruptor without market access to the bloc loses from disrupting it.

Proof. By the envelope step of Appendix A.6, $u_1' = R_1\hat{w}_1 - E\hat{P}_1$. With $\mathrm{NX}_1 = 0$, $R_1 = E$, and by Lemma 2(iii) $\hat{P}_1 = \sum_k \pi_{1k}\hat{w}_k$ (the disruptor's own links are undisrupted, so there is no direct term). Subtracting gives the first equality in (66), and $\hat{\pi}_{11} = -\theta(\hat{w}_1 - \hat{P}_1)$ gives the second. For the second part, take the disruptor's row of the equilibrium system $\tilde{J}\hat{w} = -Z_d$. Since no bloc importer buys the disruptor's goods, $(Z_d)_1 = 0$: the disruption creates no direct excess demand for them. Hence $\tilde{J}_{11}\hat{w}_1 + \sum_{j\neq 1}\tilde{J}_{1j}\hat{w}_j = 0$, or $\hat{w}_1 = \sum_{j\neq 1}\omega_j\hat{w}_j$ with $\omega_j \equiv \tilde{J}_{1j}/|\tilde{J}_{11}| \geq 0$. By Lemma 1 the row sum of $\tilde{J}$ is $-\varepsilon X_1 < 0$, so $\sum_{j\neq 1}\omega_j = 1 - \varepsilon X_1/|\tilde{J}_{11}| < 1$. The weights are positive on every good that the disruptor's customers buy: on the other bystanders' goods, and, through those customers' price indices, on bloc goods as well (the disruptor's customers are bystanders, and bystanders buy from the bloc). Since $\hat{w}_j = \hat{w}_B > 0$ for the former and $\hat{w}_j = \hat{w}_{\mathcal{M}} < 0$ for the latter (the signs of Proposition 1; Table A1 confirms that they survive the removal of the disruptor's bloc links), $\hat{w}_1 \leq c\,\hat{w}_B$ with $c \equiv \sum_{j\in\mathcal{B}\setminus\{1\}}\omega_j < 1$; write $\hat{w}_1 = c\,\hat{w}_B$ with $c \in (0,1)$ after absorbing the (negative) bloc term. Substituting into (66) with $\pi_{1k} = 0$ for $k \in \mathcal{M}$ and $\hat{w}_k = \hat{w}_B$ for the disruptor's actual suppliers yields (67), negative because $\hat{w}_B > 0$ and $c < 1$. □

The economics is transparent. Demand displaced from the corridor lands on the bystanders that bloc members actually buy from. A disruptor with access to the bloc is one of them, and collects the wedge $\hat{w}_B - \hat{w}_{\mathcal{M}}$ on each of its m bloc relationships, as in (26). A disruptor without access collects nothing directly; its wage rises only at second hand, because other bystanders' spending leaks toward it as their own goods become dear, and it rises by strictly less than theirs. Meanwhile every good it imports has become more expensive. Its real income falls. The terms-of-trade motive therefore rests on a verifiable economic condition, not on a preference assumption: the disruptor must sell into, or buy from, the markets whose sourcing the disruption redirects.

Table A1 traces the transition numerically in the Ricardian benchmark, scaling the trade cost on the disruptor's links to the bloc by a factor $a \geq 1$ while holding all other costs at $\tau = 1.5$ (the

economy is re-benchmarked at each a so that $w = \mathbf{1}$ at $d = 0$). The disruptor's gain is strictly decreasing in a, falls by two thirds when its bloc links are only 25 percent costlier than everyone else's, turns negative once they are 50 percent costlier (its share of each bloc market having fallen from 7.0 to 2.2 percent), and converges to the loss of Proposition 5 as $a \to \infty$. Throughout, the sign of the gain tracks the sign of $-\hat{\pi}_{11}$ exactly, as (66) requires.

Table A1: Market access and the disruptor's marginal gain. Ricardian quasilinear-log benchmark, $N + 1 = 12$, $m = 6$, $\theta = 3$, $\tau = 1.5$, $E = 1$. The disruptor's trade costs with the bloc are $a\tau$; $a = \infty$ removes those links. All derivatives at $d = 0$, central differences with $\varepsilon_d = 10^{-4}$.

a	π_{n1}, $n \in \mathcal{M}$	$u_1'(0)$	$\hat{w}_1$	$\hat{w}_B$ (others)	$\hat{w}_{\mathcal{M}}$	$\hat{\pi}_{11}$
1	0.0696	+0.0928	+0.1112	+0.1112	−0.1112	−0.278
1.25	0.0369	+0.0260	+0.0817	+0.1155	−0.1117	−0.111
1.5	0.0217	−0.0070	+0.0652	+0.1176	−0.1129	−0.016
2	0.0093	−0.0353	+0.0511	+0.1193	−0.1147	+0.071
3	0.0028	−0.0509	+0.0440	+0.1202	−0.1161	+0.121
5	0.0006	−0.0562	+0.0418	+0.1205	−0.1167	+0.138
∞	0	−0.0577	+0.0412	+0.1206	−0.1169	+0.142

Zeros and the data. In the Fréchet economy, as in Armington, bilateral zeros arise only as the limit $\tau_{ni} \to \infty$. Models that generate zeros at finite trade costs—bounded productivity support (Eaton et al., 2013) or fixed export costs with heterogeneous firms (Helpman et al., 2008; Chaney, 2008)—replace the condition "$\pi_{n1} > 0$" with "a positive measure of disruptor firms find bloc market n profitable"; the logic of Proposition 5 is unchanged, with the added feature that a disruption can *create* access by raising bloc price indices enough to pull marginal disruptor firms over the entry threshold. Two features of the quantitative analysis make this foundation the one that actually matters for the paper's conclusions. First, the exact hat algebra of Section 4.4 takes the 2022 share matrix as data: a link that is zero in ITPD-S stays zero in every counterfactual, so Iran's gain is computed only through trade that exists. Second, Iran's access to disrupted markets is a matter of record. In the data every country is exposed to some degree, so the relevant question is not whether Iran's partners are exposed but how heavily. Weighting each partner's exposure (its mean increase in log distance across all of its trading partners) by its share of Iran's 2022 trade gives 0.19, against a median of 0.13 for the same statistic across all countries, which places Iran in the top decile. The United Arab Emirates (23 percent of Iran's imports), India (7 percent), and Turkey

(9 percent of imports and exports) each rank in the top 15 percent of countries by exposure, and Iraq (12 percent of exports) in the top quarter; China, Iran's largest partner, is near the median.[21] Iran's gain is small not because it lacks access but because it is small: its share of any bloc market is a fraction of a percent, and (66) scales the wedge by exactly those shares.

B.5 Numerical Check of Proposition 3

Table A2 evaluates the retaliation rule (29) in the Ricardian benchmark of Table A1 (full access, $a = 1$), solving the equilibrium at $d = \pm 10^{-4}$ and differencing. Every entry matches the closed forms of Proposition 3 to six digits: the disruptor's gain falls linearly in r, crosses zero at $\bar{r} = 0.154$ as (32) predicts, and its wage response vanishes at $r = (m-1)/(N-1) = 0.5$ and equals a bloc member's at $r = 1$. Because $m = b$ in this calibration, bloc prices are invariant to r by (30), and the bloc's additional loss is the direct term $\bar{x}\, r\, \delta'$ alone. Retaliation by the target alone, on its own links with the disruptor (not covered by the proposition, which assumes a bloc-wide rule), halves the disruptor's gain at $r = 0.5$ and eliminates it at $r = 1$ ($u_1' = -0.008$), at a marginal cost to the target of -0.541 against -0.441 without retaliation and no cost to the other bloc members.

Table A2: Retaliation and the disruptor's marginal gain. Ricardian quasilinear-log benchmark of Table A1, $a = 1$; the bloc raises every disruptor–bloc link by $r\,\delta'$ per unit of disruption. Derivatives at $d = 0$; closed forms from Proposition 3.

r	$u_1'(0)$ numerical	$u_1'(0)$ closed form	$u_{\mathcal{M}}'(0)$	$\hat{w}_1$	$\hat{w}_B$
0	+0.0928	+0.0928	−0.4406	+0.1112	+0.1112
0.10	+0.0325	+0.0325	−0.4476	+0.0889	+0.1156
$\bar{r} = 0.154$	−0.0000	0	−0.4513	+0.0770	+0.1180
0.25	−0.0580	−0.0580	−0.4580	+0.0556	+0.1223
0.50	−0.2087	−0.2087	−0.4754	0.0000	+0.1334
1.00	−0.5102	−0.5102	−0.5102	−0.1112	+0.1556

C Solver Validation

The numerical exercises of Section 3.8 rely on a solver for the full equilibrium system. This appendix verifies that its price changes coincide with the closed-form solutions of Section 3.4. For a grid of twelve calibrations ($\sigma \in \{2, 3, 4, 6\}$, $\tau_0 \in \{1.2, 1.5, 2\}$), it compares the closed-form values in

[21] Iran's trade with the nineteen Western-aligned economies is, by contrast, small: 11 percent of imports and 2 percent of exports. What matters for Iran's gain is its access to heavily exposed markets, not to the Western target.

Proposition 1 with those obtained via finite differences at $d = 0$. Figure A1 plots the numerically derived price changes for bystander and bloc countries against their analytically solved counterparts across all twelve calibrations. All points lie on the 45-degree line with a maximum error of roughly 9×10^{-8}, demonstrating that the two solution methods agree to seven digits. Any deviation reported in Section 3.8 therefore reflects a genuine economic effect rather than numerical error.

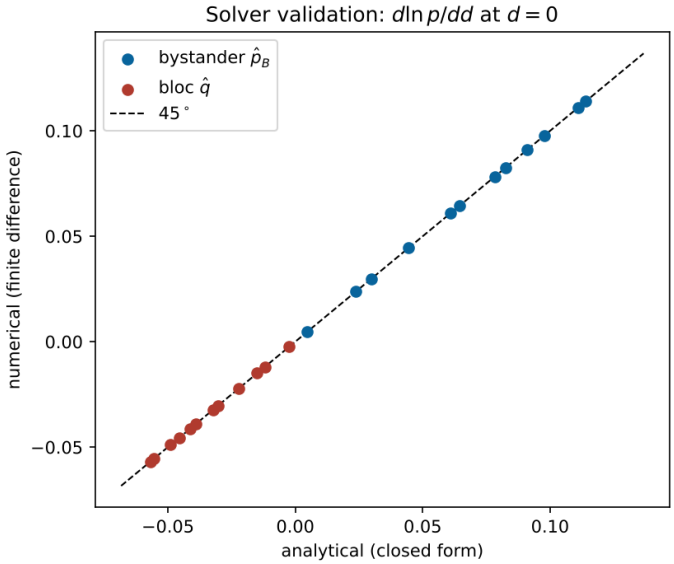


Figure A1: Solver validation. Numerically computed price responses plotted against their closed-form counterparts at $d = 0$, for bystander (disruptor) and bloc goods. Each dot corresponds to one of twelve (σ, τ_0) calibrations, with the 45-degree line superimposed. Maximum absolute error is 9×10^{-8}.

D Additional Tables

Table A3 reports the event-study coefficients plotted in Figure 5, together with the joint tests of the leads and lags discussed in Section 4.3. Table A4 reports the welfare change of every country under each of the three exercises of Section 4.4 at $\sigma = 3.5$.

Table A3: Event-study estimates around the Bab el-Mandeb closure. Event-study version of the DiD specification, interacting the binary exposure indicator with a dummy for each month.

Event time	Calendar month	Coefficient
Pre-treatment		
−9	February 2023	−0.024 (0.031)
−8	March 2023	0.001 (0.032)
−7	April 2023	0.018 (0.029)
−6	May 2023	0.047* (0.025)
−5	June 2023	−0.028 (0.021)
−4	July 2023	0.004 (0.026)
−3	August 2023	0.014 (0.022)
−2	September 2023	−0.048** (0.021)
−1	October 2023	— reference —
Post-treatment		
0	November 2023	−0.051** (0.022)
+1	December 2023	−0.064** (0.028)
+2	January 2024	−0.083** (0.033)
+3	February 2024	−0.060** (0.026)
+4	March 2024	−0.053* (0.032)
+5	April 2024	0.025 (0.027)
+6	May 2024	−0.016 (0.032)
+7	June 2024	−0.088*** (0.034)
Exporter × time FE		Yes
Importer × time FE		Yes
Pair FE		Yes
Observations		225,456
Clusters (pairs)		18,045
Pseudo R^2		0.994
Joint test of leads		$\chi^2(8) = 38.21$ [0.000]
excluding September 2023		$\chi^2(7) = 22.84$ [0.002]
Joint test of lags		$\chi^2(8) = 32.56$ [0.000]
Linear pre-trend slope		−0.0017 (0.0036)

Notes. PPML estimates (Santos Silva and Tenreyro, 2006) of monthly bilateral trade on the binary exposure indicator (Exposure ≥ 1.30) interacted with a set of event-time dummies. Event time is indexed relative to November 2023. October 2023 is the omitted reference category, to which each estimated coefficient is relative: although the attacks began on 19 October, transit times and customs-reporting lags mean that October arrivals reflect shipments dispatched before the campaign, and any residual contamination would bias the post-treatment coefficients toward zero. Unlike the specifications in Table 1, which exclude October 2023 as a transition month, the event study therefore retains it. All specifications absorb exporter-time, importer-time, and directional pair fixed effects. Standard errors, in parentheses, are clustered by directional pair. Countries involved or adjacent to the ongoing war are dropped. The linear pre-trend slope is the coefficient on exposure interacted with event time, estimated on pre-treatment months (February–October 2023) only. p-values in brackets. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A4: Welfare effects of the Red Sea closure by country and specification ($\sigma = 3.5$)

	% change in utility			Value change in utility		
ISO	EX1	EX2	EX3	EX1	EX2	EX3
AFG	-0.20	-0.22	-0.20	-40.7	-43.5	-40.0
AGO	0.00	0.00	0.00	1.3	1.5	2.8
ALB	-0.24	-0.27	-0.24	-85.7	-95.3	-86.3
ARE	-0.26	-0.29	-0.28	-14,307.0	-15,717.4	-15,476.8
ARG	0.01	0.01	0.01	102.2	114.8	107.0
ARM	-0.45	-0.49	-0.46	-529.7	-584.1	-543.5
ATG	-0.00	-0.00	-0.01	-0.0	-0.0	-0.1
AUS	-0.02	-0.02	-0.02	-1,038.9	-1,168.8	-1,186.7
AUT	-0.27	-0.30	-0.28	-3,589.5	-4,000.7	-3,685.9
AZE	-0.28	-0.31	-0.26	-302.7	-332.8	-277.9
BDI	-0.69	-0.77	-0.72	-29.0	-32.4	-30.0
BEL	-0.11	-0.13	-0.14	-2,507.2	-2,808.0	-3,142.3
BEN	0.05	0.05	0.05	5.2	5.9	5.6
BFA	0.02	0.02	0.02	5.5	6.2	7.2
BGD	-0.28	-0.32	-0.33	-1,741.4	-1,929.6	-1,997.7
BGR	-0.34	-0.38	-0.34	-732.9	-812.5	-746.1
BHR	-0.31	-0.33	-0.33	-1,237.6	-1,346.8	-1,318.7
BHS	0.01	0.01	0.01	2.5	2.8	2.6
BIH	-0.54	-0.60	-0.56	-225.5	-251.1	-231.8
BLR	-0.07	-0.08	-0.07	-60.3	-67.6	-63.3
BLZ	0.03	0.04	0.04	2.0	2.2	2.2
BMU	-0.06	-0.06	-0.06	-23.3	-26.5	-25.1
BOL	0.01	0.02	0.01	11.9	13.4	12.0
BRA	0.00	0.01	0.00	352.4	398.2	372.2

Continued on next page

Table A4 – continued from previous page

	% change in utility			Value change in utility		
ISO	EX1	EX2	EX3	EX1	EX2	EX3
BRB	0.01	0.01	0.01	0.3	0.3	0.2
BRN	-0.17	-0.19	-0.18	-43.1	-48.5	-46.8
BTN	-0.03	-0.03	-0.02	-0.9	-1.0	-0.6
BWA	-0.00	-0.00	-0.00	-9.7	-10.7	-9.0
CAF	0.05	0.05	0.05	0.4	0.5	0.5
CAN	0.00	0.00	0.00	105.2	118.4	99.4
CHE	-0.71	-0.79	-0.74	-13,404.5	-14,947.1	-13,910.9
CHL	0.00	0.00	0.00	1.8	2.1	0.6
CHN	-0.12	-0.13	-0.13	-35,392.0	-39,712.2	-38,664.1
CIV	0.04	0.04	0.05	25.9	29.3	29.6
CMR	0.05	0.06	0.06	12.4	14.0	15.2
COD	-0.01	-0.01	-0.01	-7.2	-9.0	-9.8
COG	0.02	0.02	0.01	3.2	3.4	2.1
COK	-0.14	-0.16	-0.16	-0.3	-0.3	-0.3
COL	0.01	0.01	0.01	84.2	94.8	93.7
CPV	-0.32	-0.36	-0.33	-15.8	-17.8	-16.7
CRI	0.00	0.00	0.00	9.9	11.1	12.3
CUB	0.05	0.06	0.05	5.0	5.7	4.9
CUW	0.03	0.03	0.03	1.0	1.1	0.9
CYM	0.03	0.03	0.03	3.9	4.4	4.5
CYP	-0.90	-1.00	-0.93	-554.5	-614.1	-571.6
CZE	-0.08	-0.09	-0.10	-771.7	-865.0	-991.5
DEU	-0.06	-0.06	-0.07	-10,075.6	-11,305.5	-12,639.4
DJI	-0.96	-1.05	-1.01	-308.6	-337.6	-321.9
DNK	-0.08	-0.09	-0.10	-1,282.5	-1,441.4	-1,596.1

Continued on next page

Table A4 – continued from previous page

	% change in utility			Value change in utility		
ISO	EX1	EX2	EX3	EX1	EX2	EX3
DOM	0.01	0.01	0.01	16.7	18.8	18.4
DZA	-0.53	-0.59	-0.53	-938.4	-1,046.2	-933.6
ECU	0.02	0.02	0.02	14.8	16.5	17.3
EGY	-0.26	-0.28	-0.26	-5,740.3	-6,242.0	-5,835.8
ERI	-33.63	-37.15	-34.43	-184.8	-204.1	-189.2
ESP	-0.16	-0.17	-0.19	-6,868.4	-7,687.5	-8,514.8
EST	-0.05	-0.06	-0.05	-73.1	-82.3	-72.1
ETH	-0.60	-0.67	-0.65	-240.7	-268.4	-260.3
FIN	-0.08	-0.09	-0.08	-411.5	-462.5	-419.1
FJI	0.02	0.02	0.03	2.2	2.5	2.7
FRA	-0.16	-0.18	-0.20	-9,688.6	-10,879.6	-12,118.2
FRO	0.01	0.01	0.01	0.3	0.3	0.5
FSM	-0.17	-0.19	-0.22	-0.3	-0.3	-0.3
GAB	-0.05	-0.06	-0.07	-3.3	-3.8	-4.2
GBR	-0.09	-0.10	-0.12	-9,974.9	-11,205.9	-12,465.0
GEO	-0.01	-0.01	-0.01	-200.4	-212.3	-202.6
GHA	0.01	0.01	0.01	27.6	31.0	31.0
GIN	0.06	0.07	0.06	8.4	9.3	7.9
GMB	-0.10	-0.12	-0.10	-4.8	-5.5	-4.8
GNB	0.18	0.21	0.14	0.8	0.9	0.6
GNQ	0.75	0.85	0.92	9.3	10.5	11.3
GRC	-0.57	-0.63	-0.70	-2,993.8	-3,320.4	-3,666.7
GRL	-0.08	-0.09	-0.10	-0.8	-0.9	-1.0
GTM	0.01	0.01	0.01	33.1	37.3	37.3
GUY	0.17	0.19	0.23	9.2	10.3	12.8

Continued on next page

Table A4 – continued from previous page

	% change in utility			Value change in utility		
ISO	EX1	EX2	EX3	EX1	EX2	EX3
HKG	-0.29	-0.32	-0.34	-5,377.3	-6,021.3	-6,265.5
HND	0.01	0.01	0.01	9.1	10.3	10.6
HRV	-0.01	-0.01	-0.01	-243.5	-264.3	-242.7
HTI	0.01	0.01	0.01	3.1	3.6	3.6
HUN	-0.04	-0.05	-0.05	-2,146.4	-2,385.6	-2,623.7
IDN	-0.02	-0.03	-0.03	-3,079.5	-3,438.7	-3,500.0
IND	-0.32	-0.36	-0.35	-23,951.6	-26,629.3	-26,268.5
IRL	-0.11	-0.13	-0.11	-1,898.6	-2,128.6	-1,876.9
IRN	0.01	0.01	0.00	82.6	92.0	75.7
IRQ	-3.99	-4.44	-4.50	-4,427.1	-4,927.0	-4,997.5
ISL	-0.12	-0.14	-0.12	-37.3	-41.9	-36.6
ISR	-0.66	-0.73	-0.68	-4,412.4	-4,877.3	-4,541.9
ITA	-0.32	-0.36	-0.40	-14,736.8	-16,424.7	-18,123.6
JAM	0.00	0.00	0.00	2.5	2.8	2.7
JOR	-4.08	-4.51	-4.17	-2,976.5	-3,289.9	-3,041.2
JPN	-0.03	-0.03	-0.04	-3,311.8	-3,734.3	-4,492.6
KAZ	-0.20	-0.22	-0.22	-3,265.6	-3,556.3	-3,632.6
KEN	-0.02	-0.03	-0.03	-400.9	-441.9	-437.4
KGZ	-0.01	-0.01	-0.01	-119.4	-130.5	-127.2
KHM	-0.29	-0.33	-0.34	-451.5	-505.2	-520.5
KOR	-0.16	-0.17	-0.17	-5,118.1	-5,747.5	-5,546.2
KWT	-3.76	-4.16	-4.19	-2,655.3	-2,933.1	-2,959.7
LAO	-0.32	-0.36	-0.37	-70.0	-78.5	-79.4
LBN	-3.35	-3.71	-3.46	-1,076.0	-1,190.1	-1,109.5
LBR	0.07	0.08	0.08	23.9	26.8	27.2

Continued on next page

Table A4 – continued from previous page

	% change in utility			Value change in utility		
ISO	EX1	EX2	EX3	EX1	EX2	EX3
LBY	-3.78	-4.19	-3.80	-1,182.2	-1,309.6	-1,187.5
LIE	-0.02	-0.02	-0.02	-1.4	-1.6	-1.8
LKA	-0.12	-0.13	-0.14	-457.8	-504.9	-509.7
LSO	0.06	0.07	0.07	1.7	1.9	2.0
LTU	-0.13	-0.14	-0.13	-179.3	-201.8	-177.8
LUX	-0.15	-0.17	-0.15	-534.5	-600.4	-520.6
LVA	-0.01	-0.01	-0.01	-53.1	-59.8	-55.2
MAR	-0.03	-0.03	-0.03	-1,127.9	-1,259.3	-1,143.7
MDA	-0.15	-0.17	-0.16	-157.7	-174.9	-163.3
MDG	-1.12	-1.25	-1.31	-125.2	-140.0	-146.1
MDV	-0.93	-1.03	-1.08	-54.9	-61.0	-63.7
MEX	0.00	0.00	0.00	220.3	245.7	264.1
MKD	-0.35	-0.39	-0.36	-155.8	-173.5	-161.2
MLI	0.11	0.12	0.13	15.8	17.6	18.5
MLT	-0.89	-1.00	-0.92	-523.6	-584.2	-539.9
MMR	-0.02	-0.02	-0.02	-280.6	-294.2	-297.1
MNE	-0.73	-0.81	-0.75	-79.6	-88.9	-82.0
MNG	-0.07	-0.07	-0.07	-97.5	-108.5	-103.1
MOZ	-0.03	-0.03	-0.03	-45.4	-51.0	-50.9
MRT	-1.48	-1.66	-1.52	-111.7	-125.5	-114.8
MUS	-0.18	-0.20	-0.20	-53.9	-60.0	-59.3
MWI	-0.00	-0.00	-0.00	-6.9	-7.7	-6.9
MYS	-1.17	-1.31	-1.34	-6,322.1	-7,077.1	-7,252.9
NAM	0.00	0.00	0.00	17.4	20.3	19.2
NCL	-0.08	-0.09	-0.09	-4.0	-4.5	-4.6

Continued on next page

Table A4 – continued from previous page

	% change in utility			Value change in utility		
ISO	EX1	EX2	EX3	EX1	EX2	EX3
NER	0.03	0.04	0.04	8.5	9.5	10.2
NGA	0.01	0.01	0.01	115.6	129.9	132.4
NIC	0.00	0.00	0.00	6.2	7.0	7.3
NLD	-0.06	-0.07	-0.08	-4,730.2	-5,312.6	-5,968.0
NOR	-0.02	-0.02	-0.02	-118.6	-129.9	-131.8
NPL	-0.07	-0.08	-0.07	-79.3	-87.9	-77.8
NZL	0.00	0.00	0.00	7.2	8.0	8.9
OMN	-0.54	-0.59	-0.58	-1,995.6	-2,188.9	-2,119.3
PAK	-0.76	-0.84	-0.83	-2,194.1	-2,424.9	-2,421.5
PAN	0.05	0.05	0.05	24.7	27.9	26.1
PER	0.00	0.00	0.00	19.0	21.6	22.6
PHL	-0.05	-0.06	-0.06	-1,415.0	-1,585.0	-1,634.8
PNG	-0.09	-0.10	-0.10	-24.1	-27.1	-28.4
POL	-0.01	-0.01	-0.01	-899.8	-1,010.7	-1,189.4
PRK	0.00	0.00	0.00	11.1	13.7	14.0
PRT	-0.01	-0.01	-0.01	-726.3	-810.2	-895.8
PRY	0.00	0.00	0.00	17.6	19.8	19.6
PSE	-0.90	-0.99	-0.92	-231.9	-257.5	-237.9
QAT	-2.86	-3.18	-3.39	-4,041.5	-4,485.0	-4,776.9
ROU	-0.03	-0.04	-0.04	-1,464.2	-1,623.4	-1,784.9
RUS	-1.23	-1.38	-1.25	-11,427.2	-12,837.0	-11,627.8
RWA	-0.04	-0.05	-0.04	-63.9	-70.5	-67.1
SAU	-11.34	-12.63	-12.08	-44,536.1	-49,593.5	-47,421.8
SDN	-6.04	-6.62	-6.16	-1,449.1	-1,588.1	-1,477.0
SEN	-0.02	-0.02	-0.02	-30.0	-33.8	-30.6

Continued on next page

Table A4 – continued from previous page

	% change in utility			Value change in utility		
ISO	EX1	EX2	EX3	EX1	EX2	EX3
SGP	-0.07	-0.08	-0.08	-9,726.9	-10,882.3	-11,212.1
SLB	0.15	0.17	0.18	0.8	1.0	1.0
SLE	0.00	0.00	0.00	0.0	0.1	0.3
SLV	0.02	0.02	0.02	16.8	18.9	18.7
SMR	-0.08	-0.09	-0.08	-2.3	-2.6	-2.5
SOM	0.01	0.01	0.01	226.4	284.9	248.0
SRB	-0.64	-0.72	-0.66	-766.4	-854.2	-786.8
SSD	-3.68	-4.09	-3.80	-96.7	-107.4	-100.0
SUR	0.06	0.06	0.08	1.9	2.1	2.7
SVK	-0.55	-0.61	-0.67	-2,225.4	-2,483.8	-2,738.8
SVN	-0.26	-0.29	-0.27	-980.5	-1,092.5	-998.1
SWE	-0.08	-0.09	-0.08	-869.7	-978.3	-885.6
SWZ	-0.03	-0.03	-0.03	-1.7	-1.9	-1.9
SYC	-6.47	-7.22	-7.74	-180.9	-201.9	-216.3
SYR	-2.48	-2.75	-2.54	-247.8	-274.4	-253.7
TCD	0.06	0.06	0.07	5.4	6.1	6.7
TGO	0.03	0.03	0.03	10.8	12.2	12.1
THA	-0.13	-0.15	-0.15	-4,444.6	-4,977.5	-5,045.8
TJK	-1.60	-1.78	-1.67	-166.6	-185.1	-173.9
TKM	-0.25	-0.29	-0.27	-89.0	-100.4	-94.4
TON	0.07	0.08	0.08	0.4	0.4	0.5
TTO	0.05	0.05	0.06	16.8	19.1	21.9
TUN	-0.45	-0.50	-0.46	-541.9	-603.9	-554.9
TUR	-0.12	-0.13	-0.12	-11,620.0	-12,786.6	-11,846.9
TWN	-0.37	-0.41	-0.43	-6,320.0	-7,102.0	-7,396.7

Continued on next page

Table A4 – continued from previous page

	% change in utility			Value change in utility		
ISO	EX1	EX2	EX3	EX1	EX2	EX3
TZA	-0.11	-0.13	-0.13	-206.6	-230.1	-225.2
UGA	-0.44	-0.49	-0.49	-151.5	-168.7	-168.0
UKR	-0.08	-0.09	-0.08	-1,400.1	-1,551.0	-1,428.4
URY	0.00	0.00	0.00	9.6	10.9	10.2
USA	0.01	0.01	0.01	5,191.5	5,839.1	5,846.9
UZB	-0.84	-0.93	-0.90	-905.6	-997.8	-968.9
VCT	0.01	0.01	0.01	0.2	0.2	0.2
VEN	0.12	0.14	0.14	17.3	19.5	19.6
VGB	0.05	0.05	0.05	1.3	1.5	1.3
VNM	-0.18	-0.20	-0.21	-4,337.9	-4,857.9	-5,052.1
VUT	0.04	0.04	0.05	0.2	0.2	0.2
WSM	0.12	0.13	0.14	0.8	0.9	0.9
YEM	-0.79	-0.86	-0.81	-123.6	-135.3	-127.0
ZAF	-0.16	-0.18	-0.16	-412.8	-467.0	-417.9
ZMB	-0.43	-0.49	-0.45	-221.7	-249.4	-229.3
ZWE	-0.02	-0.02	-0.01	-3.5	-3.8	-2.4