Supplementary Material
Shock Signatures at Maritime Chokepoints: A Descriptive Comparison of Three Disruptions Using Open Daily Data JEMS-87259, revision 1.

Contents

Supplementary File S1: Port-to-shock exposure mapping, all 2,065 retrieved ports, supplied as a comma-separated data file
Supplementary File S2: Cross-correlation lag profile, -30 to +30 days, all 61 lags, supplied as a comma-separated data file
Supplementary S3: Interrupted time-series specification and sensitivity
Supplementary S4: Uncertainty procedures
Supplementary S5: Port-level robustness
Supplementary S6: Rerouting test
Supplementary S7: Data provenance and vintage drift
Supplementary S8: Supplementary Methods

Supplementary S3: Interrupted time-series specification and sensitivity

Table S3.1. Window dates for the six reference constructions of §2.2. Seasonal windows are the shock window shifted to the reference year, so day-of-year composition is held fixed. The multi-year reference excludes 2020 throughout for COVID-19 and, for the Red Sea series, 2021, which contains the Ever Given closure analysed here; the all-years row restores them, so the effect of the exclusion is visible rather than assumed. December 2023 is set aside as a transition month for the Red Sea shock and appears in no reference window

Shock	Shock window	Submitted before-after	Equal-length pre-event	Seasonal, previous year	Seasonal, multi-year	Seasonal, all years
Red Sea	2024-01-01 to 2024-06-30	2023-01-01 to 2023-11-30	2023-06-01 to 2023-11-30	2023-01-01 to 2023-06-30	2019, 2022, 2023	2019, 2020, 2021, 2022, 2023
Panama	2023-11-03 to 2024-02-29	2023-01-01 to 2023-10-31	2023-07-03 to 2023-11-02	2022-11-03 to 2023-02-28	2019, 2021, 2022	2019, 2020, 2021, 2022
Ever Given	2021-03-23 to 2021-03-29	2021-03-01 to 2021-03-22	2021-03-16 to 2021-03-22	2020-03-23 to 2020-03-29	2019, 2020	2019, 2020, 2022

Table S3.2. Transited capacity against the immediate pre-event reference, deadweight tonnage, moving-block intervals at a block length of 14 days. At every chokepoint where traffic fell, capacity falls further than transit counts, which is what a shock that turns away the largest vessels first produces; at the Cape of Good Hope, where traffic rose, capacity rises less than counts. The Ever Given capacity change, omitted from the submitted version, is the largest capacity loss in the study. Table S3.7 gives capacity under all six references.

Chokepoint	Shock	Baseline (m DWT/day)	Shock (m DWT/day)	Change (%)	CI low	CI high
Suez Canal	redsea	3.4	1.5	-56.9	-59.3	-53.4
Bab el-Mandeb Strait	redsea	3.4	1.3	-62.9	-65.2	-60.1
Cape of Good Hope	redsea	3.0	5.0	65.0	49.7	80.7
Panama Canal	panama	0.9	0.6	-36.1	-38.8	-33.4
Suez Canal	evergiven	2.6	0.5	-82.1	-83.0	-81.1

Table S3.3. Level shift as a percentage of the 2023 pre-event mean, by trend estimation window and by how the contaminated periods are handled. OLS with HAC(14) throughout.

Chokepoint	full_2019_2023 / dummies	full_2019_2023 / excluded	none_2023 / none	short_2022_2023 / none
Bab el-Mandeb Strait	-56.9	-61.0	-53.8	-62.0
Cape of Good Hope	84.4	83.1	83.5	72.3
Panama Canal	-41.0	-40.3	-37.5	-31.5
Suez Canal	-46.1	-49.9	-43.5	-50.0

Table S3.4. The same estimate under three HAC bandwidths and an ARMA-error intervention model, full window with contamination indicators.

Chokepoint	OLS+HAC(7)	OLS+HAC(14)	OLS+HAC(28)	ARMA(1, 0, 1)	ARMA(2, 0, 0)
Bab el-Mandeb Strait	-56.8 (-59.5, -54.2)	-56.8 (-60.0, -53.7)	-56.8 (-60.6, -53.1)	-56.1 (-61.1, -51.1)	
Cape of Good Hope	84.4 (77.4, 91.3)	84.4 (76.1, 92.6)	84.4 (74.4, 94.3)		84.4 (74.7, 94.0)
Panama Canal	-41.0 (-44.8, -37.3)	-41.0 (-45.4, -36.7)	-41.0 (-46.1, -36.0)		-40.0 (-49.1, -30.8)
Suez Canal	-46.1 (-49.1, -43.0)	-46.1 (-49.8, -42.4)	-46.1 (-50.6, -41.6)	-46.1 (-51.4, -40.7)	

Table S3.5. Residual autocorrelation diagnostics, primary specification.

Chokepoint	Shock	Durbin-Watson	Breusch-Godfrey LM(14)	p
Suez Canal	redsea	2.052	44.780	0.000
Bab el-Mandeb Strait	redsea	1.852	24.560	0.039
Cape of Good Hope	redsea	1.521	69.890	0.000
Panama Canal	panama	1.798	55.790	0.000
Suez Canal	evergiven	1.154	30.600	0.006

The trend estimation window is the largest single lever in the paper: the Suez estimate moves by about six percentage points across the three windows in Table S3.3, while the choice between indicator variables and row deletion moves it by under four. Inference is far more stable than the point estimate, and the Breusch-Godfrey statistic rejects in every specification, which is why the intervals are described as approximate in §2.2.

Table S3.6. The primary specification with a post-intervention slope term added, so the effect may grow or shrink after the intervention date: $y_t = \beta_0 + \beta_1 t + \beta_2 D_t + \beta_3 s_t D_t +$ calendar terms $+ \varepsilon_t$, with s_t the calendar days since the intervention. The onset jump is β_2; the window mean is $\beta_2 + \beta_3$ times the mean of s_t over the shock window, and the last-day effect is evaluated at its end. All percentages use the 2023 pre-event mean of Table 1; the intervals are HAC(14) with the delta method. A negative AIC difference favours the slope model.

Chokepoint	Level shift only (%)	Onset jump (%)	Slope, vessels/day per 30 days	Slope p	Window mean, with slope (%)	Last day of window (%)	AIC, slope minus level
Suez Canal	-46.1	-37.7	-2.04 (-3.30, -0.78)	0.002	-46.1 (-49.1, -43.1)	-54.4	-24.5
Bab el-Mandeb Strait	-56.9	-49.8	-1.74 (-2.86, -0.62)	0.002	-56.8 (-59.5, -54.2)	-63.8	-19.9
Cape of Good Hope	84.4	64.5	3.22 (1.16, 5.27)	0.002	84.3 (77.8, 90.9)	104.2	-26.5
Panama Canal	-41.0	-35.6	-0.88 (-1.91, 0.15)	0.094	-41.1 (-45.0, -37.1)	-46.5	-2.4

Adding the slope changes how the effect is distributed inside the shock window and not its average over the window. At Panama the slope is small and its interval includes zero, and the window mean is −41.1% against −41.0% for the level shift. At the three Red Sea chokepoints the slope is clearly non-zero, so the effect is smaller at onset and larger at the end of the window, but the window mean again matches the level shift to within 0.1 percentage points. The level-shift coefficient therefore reports the mean departure from the counterfactual over the shock window, which is the quantity Table 1 states, and not the effect on any particular day.

Table S3.7. Change in daily transited capacity, by reference, laid out as Table 1 is for transit counts. ITS is the primary specification fitted to the capacity series, as a percentage of the 2023 pre-event mean; the other columns are window-mean percent changes with moving-block intervals at a block length of 14 days. The Ever Given ITS cell is empty for the reason given in §3.1.

Chokepoint / shock	ITS (primary)	Submitted before-after	Equal-length pre-event	Seasonal, previous year	Seasonal, multi-year	Seasonal, all years
Suez Canal / Red Sea	-56.9 (-62.0, -51.8)	-56.9 (-59.3, -53.4)	-56.9 (-59.4, -53.2)	-57.0 (-59.6, -53.6)	-48.7 (-51.6, -44.4)	-46.2 (-49.2, -41.7)
Bab el-Mandeb Strait / Red Sea	-63.8 (-68.4, -59.1)	-62.9 (-65.2, -60.1)	-62.9 (-65.1, -59.9)	-63.2 (-65.4, -60.1)	-56.4 (-58.8, -53.1)	-54.3 (-56.8, -50.8)
Cape of Good Hope / Red Sea	74.4 (66.4, 82.3)	65.0 (49.7, 80.7)	49.8 (37.4, 61.8)	81.8 (64.3, 101.7)	89.1 (74.3, 103.5)	85.7 (70.7, 99.9)
Panama Canal / drought	-48.0 (-51.8, -44.3)	-36.1 (-38.8, -33.4)	-33.0 (-35.8, -30.2)	-40.1 (-42.5, -37.6)	-40.2 (-42.8, -37.8)	-40.9 (-43.3, -38.7)
Suez Canal / Ever Given		-82.1 (-83.0, -81.1)	-81.5 (-82.5, -80.5)	-79.4 (-81.3, -77.3)	-80.2 (-81.3, -79.0)	-80.5 (-81.3, -79.7)

Capacity changes in the same direction under all six references at every chokepoint, and falls further than counts wherever traffic fell. The size of the change depends on the reference, as it does for transit counts: the Suez capacity decline ranges from −46.2% to −57.0% and the Cape of Good Hope rise from +49.8% to +89.1%. The four robustness axes of Tables S3.3 and S3.4 were run on transit counts only.

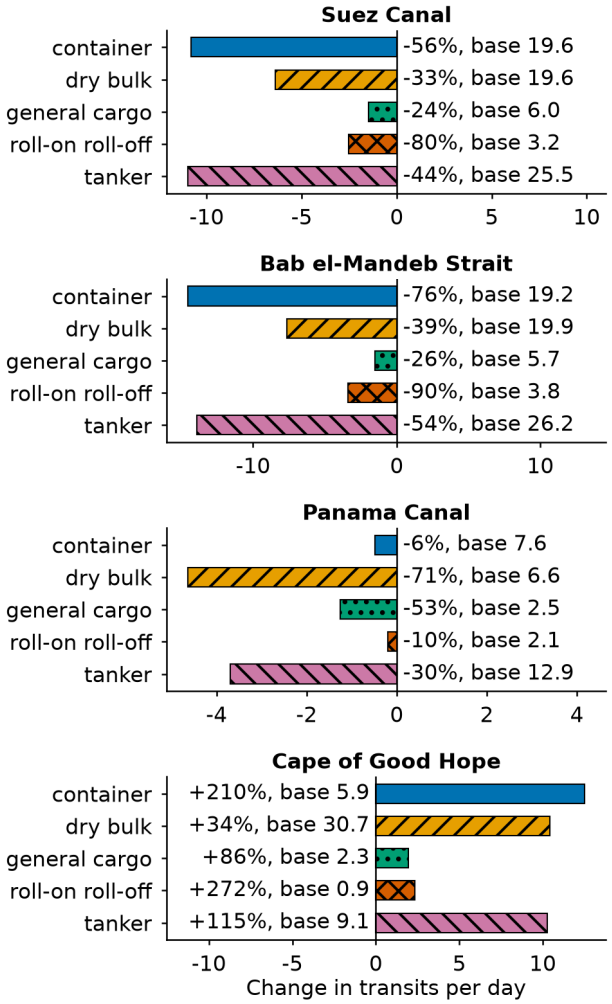


Figure S3.1. Change in daily transits by vessel class. Bars give the absolute change in vessels per day; each label gives the percentage change and the baseline it is measured against, so both rankings can be read together. Classes differ by colour and fill pattern. Horizontal scales differ: the Panama panel spans about a third of the others. Read with Table 2, which adds the shock-window levels and the intervals.

Supplementary S4 — Uncertainty procedures

Table S4.1. Automatic block-length selection, Politis-White with the Patton-Politis-White correction, for every window carrying an interval.

Series and window	n	m-hat	Optimal block b*
Suez Canal	redsea_base	334	1.00
Suez Canal	redsea_shock	182	4.00
Suez Canal	evergiven_base	22	1.00
Suez Canal	evergiven_shock	7	
Bab el-Mandeb Strait	redsea_base	334	1.00
Bab el-Mandeb Strait	redsea_shock	182	5.00
Panama Canal	panama_base	304	8.00
Panama Canal	panama_shock	119	4.00
Cape of Good Hope	redsea_base	334	1.00
Cape of Good Hope	redsea_shock	182	11.00

Table S4.2. Trough location and value by search horizon.

Chokepoint	Shock	90 days	180 days	365 days
Bab el-Mandeb Strait	redsea	31.71 (2024-01-27)	28.86 (2024-04-05)	25.29 (2024-07-26)
Panama Canal	panama	17.57 (2024-01-22)	17.57 (2024-01-22)	17.57 (2024-01-22)
Suez Canal	evergiven	16.43 (2021-03-26)	16.43 (2021-03-26)	16.43 (2021-03-26)
Suez Canal	redsea	37.71 (2024-02-17)	35.57 (2024-04-30)	31.57 (2024-11-18)

Table S4.3. The Ever Given closure under four procedures.

Procedure	Estimate	CI low	CI high
asymmetric, baseline MBB block=3	-71.1	-72.5	-69.7
asymmetric, baseline MBB block=7	-71.2	-72.6	-69.4
stationary bootstrap both windows, mean block=2	-71.1	-86.8	-53.1
stationary bootstrap both windows, mean block=3	-71.3	-86.4	-54.0

Table S4.4. Range of the headline transit estimate across block lengths 3, 7, 14, 21, 28.

Chokepoint	Shock	Min over blocks	Max over blocks	Span (pp)
Bab el-Mandeb Strait	redsea	-55.6	-54.8	0.8
Cape of Good Hope	redsea	76.1	76.9	0.8
Panama Canal	panama	-33.4	-32.5	0.9
Suez Canal	evergiven	-78.5	-71.1	7.4
Suez Canal	redsea	-44.5	-43.7	0.8

Table S4.5. Window-boundary check: intervention date shifted by -7 to +7 days, primary interrupted time-series specification.

Chokepoint	Shock	Min	Max	Span (pp)
Bab el-Mandeb Strait	redsea	-56.9	-56.6	0.3
Cape of Good Hope	redsea	83.9	84.4	0.5
Panama Canal	panama	-41.1	-39.3	1.8
Suez Canal	redsea	-46.2	-46.0	0.2

The block length is not a material choice: across the ladder no estimate moves by more than the span in Table S4.4, and shifting the intervention date by a week moves the primary estimate by between 0.2 and 1.8 percentage points (Table S4.5). This is why the boundary check is reported here rather than as a headline robustness result: shifting a window boundary was never the thing that needed testing. The objection is that the windows are not seasonally comparable, and a shift of a few days cannot repair that.

Table S4.6. Temporal signature quantities. Baseline is the pre-event mean of the seven-day centred moving average. Trough quantities are located within 180 days of onset. Moved here from the main text in revision so that the article meets the journal's page limit; §3.3 states the findings it carries.

Chokepoint / shock	Baseline	Raw daily min (date)	Smoothed trough (date)	Trough-window mean	vs baseline (%)	Trough interval	Half-recovery	Trough stable across horizons
Suez Canal / Red Sea	73.92	28 (2024-02-15)	35.57 (2024-04-30)	35.57	-51.90	34.0 to 37.71	right-censored at 2026-09-10	no
Bab el-Mandeb Strait / Red Sea	74.81	19 (2024-04-23)	28.86 (2024-04-05)	28.86	-61.40	27.43 to 30.14	right-censored at 2026-09-10	no
Panama Canal / drought	31.68	15 (2023-12-26)	17.57 (2024-01-22)	17.57	-44.50	17.43 to 19.72	126 days	yes
Suez Canal / Ever Given	55.62	2 (2021-03-25)	16.43 (2021-03-26)	16.43	-70.50	16.43 to 50.29	3 days	yes

The three trough quantities are reported in separate columns because they are not interchangeable: for the Ever Given closure the raw daily minimum is an integer count of two transits and the smoothed trough is a rate roughly eight times larger. Figure S4.1 draws the four quantities on that closure week.

Table S4.7. Which signature dimensions are measurable, by event. Moved here from the main text in revision; §3.3 states the finding.			
Dimension	**Red Sea (Suez, Bab el-Mandeb)**	**Panama drought**	**Ever Given**
Depth	yes	yes	yes
Breadth across vessel classes	yes	yes	yes
Capacity intensity	yes	yes	yes
Duration	no - trough horizon-dependent, recovery censored	yes, 126 days	yes, 3 days
Substitution	yes (Cape of Good Hope)	no - no undisturbed alternative	no - event shorter than the sailing time
Onset profile	yes	yes	yes

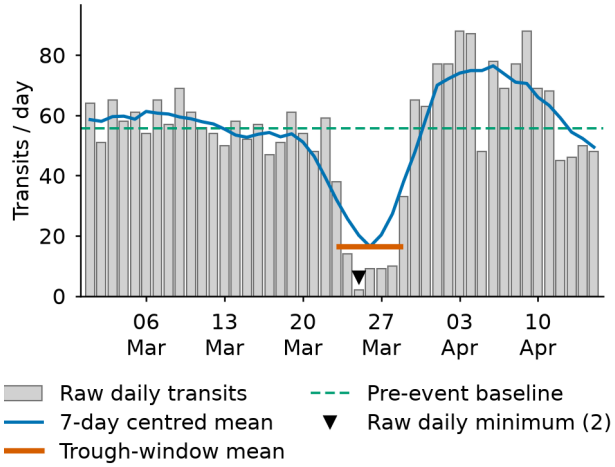


Figure S4.1. The Ever Given closure at Suez, 1 March–15 April 2021, with the quantities §2.4.1 distinguishes. Grey bars: raw daily transits. Blue: seven-day centred mean. Orange segment: trough-window mean, 16.43. Green dashed: pre-event baseline, 55.62. Triangle: raw daily minimum, two transits on 25 March. The average reaches its minimum exactly at the trough-window mean because it averages the same seven days. Read with Supplementary S4, Table S4.6.

Supplementary S5 — Port-level robustness

Table S5.1. Share of ports removed by the activity threshold.			
Shock	**Threshold (calls/day)**	**Control removed %**	**High removed %**
evergiven	0.0	0.0	0.0
evergiven	0.5	43.2	25.8
evergiven	1.0	65.3	42.9
evergiven	2.0	81.0	61.2
evergiven	5.0	95.8	82.4
panama	0.0	0.0	0.0
panama	0.5	32.5	46.7
panama	1.0	51.1	66.8
panama	2.0	70.4	81.1
panama	5.0	89.9	94.6
redsea	0.0	0.0	0.0
redsea	0.5	44.0	28.0
redsea	1.0	66.2	44.6
redsea	2.0	82.5	62.0
redsea	5.0	95.8	81.8

Table S5.2. Group composition and baseline balance, before the filter.

Shock	Group	Ports	Countries	Median base	p90 base	Share < 1/day
redsea	HIGH	1,139	76	1.210	8.931	0.446
redsea	CONTROL	361	70	0.593	3.437	0.662
redsea	MIXED	261	31	0.458	3.162	0.686
panama	HIGH	259	41	0.530	3.485	0.668
panama	CONTROL	808	93	0.970	5.020	0.511
panama	MIXED	692	43	1.054	10.658	0.484
evergiven	HIGH	1,091	76	1.273	9.136	0.429
evergiven	CONTROL	331	70	0.591	3.409	0.653
evergiven	MIXED	229	30	0.500	3.109	0.659

Table S5.3. Difference in median port-call deviation, high exposure minus control, under five specifications. An asterisk marks an interval excluding zero.

Specification	Red Sea	Panama	Ever Given
large_ports_only	-1.26 (-6.44, +3.22)	+5.60 (+3.72, +8.68) *	+2.78 (-5.56, +10.28)
mixed_as_control	-1.50 (-4.17, +0.14)	+2.14 (+0.44, +4.01) *	-1.74 (-5.86, +3.45)
mixed_as_high	-2.22 (-5.46, -0.43) *	+3.06 (+0.92, +5.19) *	-1.08 (-4.93, +4.71)
seasonal_reference	-1.58 (-6.15, +2.78)	+6.10 (-0.18, +8.95)	+8.23 (+2.27, +14.26) *
submitted_country	-2.47 (-5.63, -0.31) *	+3.63 (+1.54, +5.75) *	-1.35 (-5.51, +4.08)

The threshold removes the two groups at different rates in every design, so it is a feature of the comparison rather than a data-quality step. Only the Panama contrast survives more than two specifications, which is the basis for the statement in §3.4 that the port-level result is not robust to defensible choices about group membership.

Table S5.4. Group medians after the one-call-per-day filter, with the clustered interval beside the width an independent-port resample would have given. Moved here from the main text in revision; §3.4 states what it shows.

Shock	Group	Ports (countries)	Median deviation	Clustered interval	Clustered width	i.i.d. width
Red Sea	high exposure	631 (67)	-1.14%	-1.89 to -0.43	1.46	1.07
Red Sea	control	122 (53)	+1.15%	-0.71 to +4.29	5.00	4.58
Panama	high exposure	86 (30)	-0.83%	-2.01 to +0.43	2.44	2.88
Panama	control	395 (78)	-4.56%	-6.20 to -2.76	3.44	2.80
Ever Given	high exposure	623 (67)	-0.07%	-1.57 to +1.49	3.07	2.97
Ever Given	control	115 (47)	+1.54%	-4.14 to +4.76	8.90	7.48

Clustering on the country widens the Red Sea high-exposure interval from 1.07 to 1.46 percentage points, a factor of 1.4, so the port-level resample in the submitted version was mildly anti-conservative rather than badly wrong. The remaining five rows move by less.

Supplementary S6 — Rerouting test

Table S6.1. Correlation of Cape of Good Hope daily transit with the affected chokepoints, four series.

Pair	Series	n	r	CI low	CI high	Strongest lag	r at lag
Cape of Good Hope vs Suez Canal	levels	458	-0.548	-0.689	-0.317	-6	-0.588
Cape of Good Hope vs Suez Canal	first_difference	457	0.061	-0.038	0.155	-6	-0.120
Cape of Good Hope vs Suez Canal	calendar_adjusted	458	-0.434	-0.583	-0.255	-2	-0.477
Cape of Good Hope vs Suez Canal	its_residual	458	0.038	-0.058	0.119	7	0.089
Cape of	levels	458	-0.593	-0.707	-0.390	-2	-0.608

Table S6.1. Correlation of Cape of Good Hope daily transit with the affected chokepoints, four series.

Pair	Series	n	r	CI low	CI high	Strongest lag	r at lag
Good Hope vs Bab el-Mandeb Strait							
Cape of Good Hope vs Bab el-Mandeb Strait	first_difference	457	-0.017	-0.087	0.067	-28	0.127
Cape of Good Hope vs Bab el-Mandeb Strait	calendar_adjusted	458	-0.482	-0.613	-0.316	-1	-0.486
Cape of Good Hope vs Bab el-Mandeb Strait	its_residual	458	-0.072	-0.123	-0.010	11	0.085

Table S6.2. Candidate alternatives to a Panama transit over the shock window, against the seasonally matched reference.

Candidate	Class	Reference	Shock	Change %	CI low	CI high
Magellan Strait	all vessels	4.42	5.92	36.90	-4.10	82.70
Magellan Strait	container	0.46	0.29	-36.20	-62.50	7.00
Magellan Strait	dry bulk	2.25	3.28	49.00	10.60	91.10
Magellan Strait	tanker	1.09	1.45	34.50	-11.70	92.50
Suez Canal	all vessels	62.33	59.44	-4.90	-20.80	10.40
Suez Canal	container	17.54	12.92	-27.60	-43.30	-8.80
Suez Canal	dry bulk	16.67	19.52	17.60	1.00	32.40
Suez Canal	tanker	19.43	19.56	1.10	-16.10	18.40
Cape of Good Hope	all vessels	45.12	64.55	40.80	16.10	66.80
Cape of Good Hope	container	5.74	12.51	115.80	50.60	183.40
Cape of Good Hope	dry bulk	25.12	32.71	28.40	13.70	45.30
Cape of Good Hope	tanker	10.88	13.99	27.10	-0.70	56.60

The profile is flat once the regime change is removed, so the strongest lag in Table S6.1 is the maximum of a noisy surface rather than an estimated delay. Of the three candidates in Table S6.2, only the Magellan Strait is free of contamination from the concurrent Red Sea shock, and its base is too small for it to serve as a rerouting reference.

The complete lag profile over −30 to +30 days, all 61 lags, is provided as Supplementary File S2.

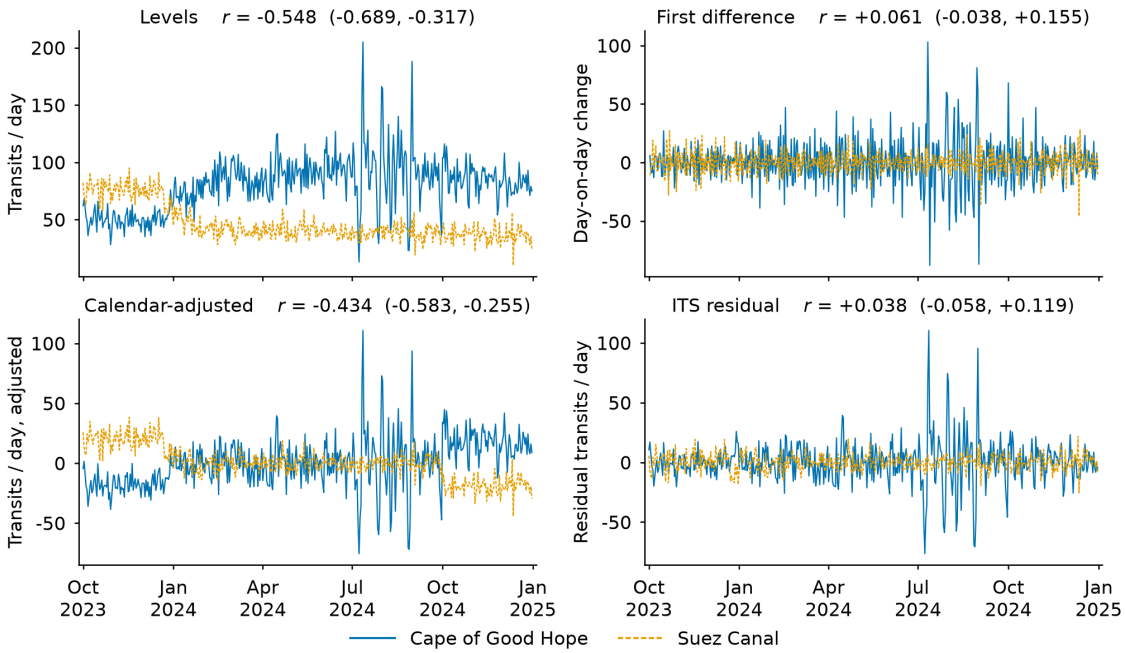


Figure S6.1. Daily transits at the Cape of Good Hope against the Suez Canal under four filters, 1 October 2023–31 December 2024, with the contemporaneous correlation and its bootstrap interval in each panel title. The level correlation survives calendar adjustment but not differencing or the removal of the interrupted time-series fit, so it reflects a shared level shift at a common date rather than day-to-day co-movement. Replaces the submitted Figure 3. Read with Supplementary S6, Table S6.1.

Supplementary S7 — Data provenance and vintage drift

Table S7.1. Retrieval log for the deposited vintage.

File	Retrieved (UTC)	Rows	SHA-256 (first 16)
a1	2026-09-18T08:54:31	11,252	c3bae69d7ea8b0aa
a2	2026-09-18T08:57:59	30,975	679cf2f01e345601
a3	2026-09-18T11:42:24	10,325	f8b2c477205ef7ea
e1	2026-09-18T12:16:09	2,813	fe21c29d078f1c33

Table S7.2. The same estimate on both retrievals.

Chokepoint	Shock	Quantity	Estimator	2026-06-30	2026-09-18	Drift
Suez Canal	redsea	transit	RD1 before-after, block 14	-44.18	-44.16	0.02
Suez Canal	redsea	level shift, vessels/day	ITS full 2019-2023 + dummies	-33.80	-34.05	-0.25
Suez Canal	redsea	capacity	RD1 before-after, block 14	-56.79	-56.82	-0.04
Suez Canal	redsea	transit, container	RD1 before-after, block 14	-55.62	-55.66	-0.05
Bab el-Mandeb Strait	redsea	transit	RD1 before-after, block 14	-55.35	-55.22	0.13
Bab el-Mandeb Strait	redsea	level shift, vessels/day	ITS full 2019-2023 + dummies	-42.45	-42.54	-0.09
Bab el-Mandeb Strait	redsea	capacity	RD1 before-after, block 14	-62.99	-62.94	0.05
Bab el-Mandeb Strait	redsea	transit, container	RD1 before-after, block 14	-75.86	-75.79	0.06
Cape of Good Hope	redsea	transit	RD1 before-after, block 14	77.28	76.95	-0.32
Cape of Good Hope	redsea	level shift, vessels/day	ITS full 2019-2023 +	41.50	41.26	-0.25

Table S7.2. The same estimate on both retrievals.						
Chokepoint	**Shock**	**Quantity**	**Estimator**	**2026-06-30**	**2026-09-18**	**Drift**
			dummies			
Cape of Good Hope	redsea	capacity	RD1 before-after, block 14	64.65	64.72	0.07
Cape of Good Hope	redsea	transit, container	RD1 before-after, block 14	209.55	209.47	-0.08
Panama Canal	panama	transit	RD1 before-after, block 14	-33.18	-33.07	0.12
Panama Canal	panama	level shift, vessels/day	ITS full 2019-2023 + dummies	-12.20	-13.00	-0.80
Panama Canal	panama	capacity	RD1 before-after, block 14	-36.28	-36.16	0.12
Panama Canal	panama	transit, container	RD1 before-after, block 14	-5.64	-5.60	0.04
Suez Canal	evergiven	transit	RD1 before-after, block 14	-70.76	-71.11	-0.35
Suez Canal	evergiven	capacity	RD1 before-after, block 14	-82.15	-82.17	-0.02
Suez Canal	evergiven	transit, container	RD1 before-after, block 14	-64.75	-64.80	-0.05

Between the two retrievals 38.9% of the 10,916 overlapping chokepoint-days changed, with a mean absolute change of 0.73 vessels per day, yet no percent-change estimate in Table S7.2 moves by more than 0.35 percentage points. The deposit therefore contains the frozen vintage rather than a query: the numbers in this paper reproduce from the deposited files, not from a fresh retrieval.

Supplementary S8 — Supplementary Methods

This section carries the procedural detail of Section 2. It is here rather than in the manuscript so that the article stays within the journal's length limit; nothing has been removed from the study. Each block below is the complete statement of a procedure that the main text names, specifies, cites and points to by number, and each heading says which question in the review the procedure answers.

S8.1 Uncertainty for chokepoint quantities: the moving-block bootstrap algorithm

Answers the request that the statistical procedures be written out and attributed. Referred to from §2.4.3.

Uncertainty for chokepoint quantities comes from a moving-block bootstrap [26], which preserves the short-range dependence that resampling individual days would destroy. For a window of n days and block length ℓ, one replication is formed by

1. drawing $\lceil n/\ell \rceil$ start positions uniformly from the $n - \ell + 1$ admissible starts,
2. concatenating the blocks of length ℓ that begin at those positions, and
3. truncating the concatenation to length n.

For a percent change, the reference mean and the shock mean are recomputed on every replication and the change is formed from the paired draws, so that both windows contribute to the interval rather than the shock window alone. Where a reference spans several calendar years, each year is resampled within itself and the yearly means are pooled by length, so that no block straddles a gap between years.

All intervals are 95 percent percentile intervals from 2,000 replications with the seed fixed at 42. The routine that performs the resampling is included in the deposit.

On the block length. The main text states the specification and its justification: 14 days for the reported tables, not the optimum for any single series, chosen for comparability and lying inside the data-driven range that automatic block-length selection [27], [28] returns. The per-window optima, ranging from 2.4 to 21.0 days, are in Supplementary S4, Table S4.1, and the ladder over 3, 7, 14, 21 and 28 days is in Table S4.4.

S8.2 The Ever Given closure: the asymmetric and stationary procedures

Answers the question of what the reported intervals cover for a seven-day event. Referred to from §2.4.4.

A conventional moving block cannot be sampled inside a seven-day window: with $n = 7$ and $\ell = 14$ the number of admissible starts is $\max(1, 7 - 14 + 1) = 1$, so every replication returns the original sample. The defect this produced in the submitted analysis is recorded in the footnote to §2.4.4.

The asymmetric procedure. The closure window holds seven complete daily observations of a unique, non-repeatable event, so no resampling of it can yield information about other closure weeks. The estimand is therefore the change of the observed closure-week mean against the baseline regime mean. The baseline is resampled by moving block as in S8.1; the closure week is treated as a complete enumeration and contributes no resampling variability. Because the procedure is asymmetric it is reported as such, and not as a symmetric interval.

The stationary-bootstrap sensitivity. For readers who prefer to treat the closure week as stochastic, a stationary bootstrap [29], whose geometric block lengths and circular wrapping are defined at $n = 7$ where a fixed 14-day block is not, is applied to both windows with a mean block length of 3 and reported alongside. The point estimate is stable under both procedures at about −71 percent; the interval is not, running 2.8 percentage points wide under the asymmetric procedure and 33.7 under the stationary bootstrap. Both are in Supplementary S4, Table S4.3.

S8.3 Selection-aware intervals around a trough

Answers the request that intervals account for the trough having been chosen by inspection. Referred to from §2.4.5.

A trough is located by inspecting the series, so an interval that conditions on the chosen date understates the uncertainty. The trough is therefore re-selected inside every replication: the event-period series is resampled by a circular block bootstrap [30], the seven-day moving average is recomputed on the resampled series, and its minimum is taken afresh. What is reported is the percentile interval of that re-selected minimum, not an interval around a fixed date.

The bootstrap is not consistent for extreme order statistics. The main text states this limitation and its consequence: the interval is reported as approximate and asymmetric, and Supplementary S4, Table S4.6 gives the trough-window mean — an average, for which the bootstrap is well behaved — rather than the minimum itself.

S8.4 Exposure mapping: the coast split and the four exclusions

Answers the requirement that the port-to-shock mapping be published in full and its exclusions justified. Referred to from §2.5.1.

The mapping is country-level. Two countries have coastlines on both sides of the relevant divide and are split by port longitude rather than assigned whole: the United States and Canada, whose Pacific ports face Panama routings and whose Atlantic and Gulf ports do not.

Four groups of ports are excluded from both exposure classes rather than assigned:

1. Offshore pseudo-ports — floating storage and offloading facilities, which the source lists among the ports but which vessels do not call at in the ordinary sense.
2. Multi-coast countries that cannot be split cleanly, where a longitude rule would misassign a material share of traffic. Mexico is the principal case: its Pacific and Gulf coasts serve different lanes and its port records do not separate them reliably.
3. Caspian-only states, which have no maritime connection to either chokepoint.
4. The Russian Federation, whose port traffic over the study period is dominated by sanction-driven rerouting unrelated to either shock.

Ports whose dominant lane genuinely splits are labelled mixed and excluded from the binary contrast, and are reported separately. Ambiguity was resolved towards *mixed* deliberately: the aim is that the high-exposure and control groups are structurally clean, not that every port receives a label.

No port-call quantity, import or export volume, or any other shock-period outcome enters the assignment, and no modelled spillover output was used. The complete assignment for all 2,065 retrieved ports is Supplementary File S1.

S8.5 Group composition, the activity filter and the five alternative specifications

Answers the objection that the port comparison is misaligned and that its groups were never described. Referred to from §2.5.2 and §2.5.4.

Composition and balance. Group composition and baseline balance are reported before any group difference is interpreted, by group and by shock. In the Red Sea design the high-exposure group holds 1,139 ports in 76 countries with a median baseline of 1.21 calls per day, against 361 ports in 70 countries at 0.59. The Panama design has the imbalance reversed: 259 high-exposure ports in 41 countries at 0.53 against 808 controls in 93 countries at 0.97. Any group difference therefore partly reflects port size rather than exposure. The full table is Supplementary S5, Table S5.2.

The activity filter. The port-level analysis retains ports with a baseline mean of at least one call per day. This threshold is not neutral between groups: in the Red Sea design it removes 44.6 percent of high-exposure but 66.2 percent of control ports, so it changes the composition of the two groups differently. It is therefore reported as a result rather than as a data-quality footnote. The full ladder at thresholds of 0, 0.5, 1, 2 and 5 calls per day is Supplementary S5, Table S5.1.

The five specifications. Because the mapping is coarse, the contrast is recomputed under five specifications:

1. the country mapping of S8.4, as pre-specified;
2. mixed-lane ports assigned to high exposure;
3. mixed-lane ports assigned to control;
4. large ports only, at a threshold of five calls per day;
5. the seasonal reference of §2.2 in place of the immediate pre-event window.

A contrast that changes sign or loses its interval across these specifications is reported as not robust. Supplementary S5, Table S5.3 gives the difference and its clustered interval under each, and §3.4 states which survive.

S8.6 Minimum detectable effect for a difference in medians, by simulation

Answers the objection that the power calculation did not match the statistic reported. Referred to from §2.5.3.

The reported quantity is the median across ports of the percent deviation of shock-window mean daily calls from the reference mean, and the contrast is the difference between the two group medians. The two-sample formula for a difference in means does not describe this statistic under clustering, so its standard error is obtained by simulation:

1. Draw countries with replacement from the observed set, retaining every port in a drawn country, until the resampled set matches the original number of countries.
2. Compute the group medians and their difference on the resampled set.
3. Repeat for 2,000 replications and take the standard deviation of the resampled differences as the standard error of the median difference.

The minimum detectable effect is then $(z_{0.975} + z_{0.80}) \times \mathrm{SE}$, with $z_{0.975} \approx 1.96$ and $z_{0.80} \approx 0.84$, for a two-sided test at $\alpha = 0.05$ and 80 percent power. It is reported both in percentage points and standardised by the between-port standard deviation, so that it is comparable with the threshold of 0.20 pre-specified before the port-level analysis.

Clustering on the country rather than the port widens the Red Sea high-exposure interval by about a factor of 1.4 and leaves the others close to unchanged, which is why the submitted port-level resample was mildly anti-conservative [31].

S8.7 The rolling-recognition rule, in full

Answers the objection that the early-recognition claim was retrospective. Referred to from §2.7.

Let B be the baseline mean of the seven-day centred moving average, and let the window be [onset, onset + k]. Within that window:

- depth is the minimum of the moving average relative to B, as a percentage;
- the series counts as recovered if it stands at or above B + (trough − B)/2 at day k and has done so for the preceding seven days;
- slope is the ordinary least-squares slope of the moving average over the last min(k, 30) days.

The label follows from these three quantities alone:

Condition	Label
depth > −10%	no material shock
recovered	transient
not recovered and slope ≤ 0	persistent
not recovered and slope > 0	undetermined, recovering

No quantity computed after day k enters the label at day k. A label is useful only if it is both correct and stable, so the reported quantity is the first k at which the rule reaches the eventual label and never subsequently leaves it, not merely the first k at which the two happen to coincide.

Table S8.1. Label assigned from the first *k* days, against the eventual label. Depth is the minimum of the seven-day moving average within the window, relative to baseline. Moved here from the main text in revision so that the article meets the journal's page limit; §3.6 states the findings it carries and Figure 2 draws the same grid.

Chokepoint / shock	***k*=7**	***k*=14**	***k*=30**	***k*=60**	***k*=90**	***k*=120**	**Eventual**
Suez Canal / Red Sea	no material shock	no material shock	persistent	persistent	undetermined	persistent	persistent
Bab el-Mandeb / Red Sea	no material shock	no material shock	persistent	persistent	undetermined	undetermined	persistent
Panama Canal / drought	*persistent*	*persistent*	*persistent*	undetermined	undetermined	*persistent*	transient
Suez Canal / Ever Given	undetermined	transient	transient	transient	transient	transient	transient
Depth at k, Suez / Red Sea	−1.1%	−1.1%	−21.0%	−47.1%	−49.0%	−51.5%	
Depth at k, Panama	−24.2%	−28.8%	−30.6%	−36.9%	−44.5%	−44.5%	

The two Red Sea chokepoints are invisible for the first two weeks, correct in a middle band of horizons, and lost again at the end of the ladder. Only the Ever Given is recognised early and stably, which is the easiest possible case for a rule of this kind. At the two horizons an operator would most plausibly use, the rule does not distinguish the transient Panama shock from the persistent Suez one.

S8.8 The four filters behind the rerouting correlation

Answers the objection that a correlation on levels cannot isolate rerouting. Referred to from §2.3.

Concurrent movement is measured as the contemporaneous Pearson correlation of two daily transit series over the Red Sea transition and shock period, 1 October 2023 to 31 December 2024, 458 days. A correlation on the levels of two series that change regime at the same date is driven partly by that common regime change, so the correlation is computed on four series in parallel:

1. the levels, retained for comparison with the submitted version;
2. first differences, which remove any common level shift;
3. a calendar-adjusted series, with day-of-week and month means removed;
4. the residuals of an interrupted time-series model of the same specification as §2.2, refitted on the correlation window itself — 458 observations, a linear trend, day-of-week and month terms and a level shift at 1 January 2024 — from which trend, calendar composition and the level shift have all been removed.

Only the last two isolate movement that is not attributable to the regime change itself.

Intervals come from a moving-block bootstrap [26] at the block length of §2.4.3, with the blocks applied jointly to both series so that the day-by-day pairing is preserved; a correlation resampled with independent blocks would not estimate the same quantity.

A lead–lag profile is computed over lags of −30 to +30 days, all 61 of which are in Supplementary S6 and Supplementary File S2.

S8.9 The three trough quantities and the half-recovery rule

Answers the requirement that "trough" never be used unqualified and that the recovery rule be fully reproducible, including for series that never recover. Referred to from §2.4.1 and §2.4.2.

Smoothing. The seven-day centred moving average is computed on the full series with a minimum period equal to the window, so no truncated window can be returned at a segment boundary, and it is never applied to a segment that begins at the event onset. This is the defect that produced the "13.7 transits" figure in the submitted version: a centred seven-day mean with a minimum period of one, applied to a segment starting on the onset date, returned a six-day truncated mean of 13.667, which the minimum search then selected.

The three quantities. The raw daily minimum is the smallest daily transit count in the search window, an integer, reported with its date. The smoothed trough is the minimum of the seven-day centred moving average, a rate and not a count. The trough-window mean is the mean of the seven raw daily counts centred on the smoothed trough.

Half-recovery. Let B be the mean of the smoothed series over the pre-event window and T the smoothed trough located within 180 days of event onset. The half-recovery level is

$$M = T + (B - T) / 2$$

and half-recovery time is the number of days from the trough date to the first date on which the smoothed series reaches or exceeds M *and remains at or above M for 14 consecutive days*. The persistence requirement is what prevents a single crossing from counting as recovery; it is the reason the Panama half-recovery is 126 days rather than the 74 reported in the submitted version.

Where no such date exists on or before the last observation, half-recovery is right-censored and reported as "not reached by [date]" together with the length of the follow-up period. No value is imputed for a censored series.

S8.10 The vessel-class decomposition

Answers the requirement that every vessel class be reported rather than a container subset with the remainder inferred. Referred to from §2.6.

For each chokepoint and each class the decomposition reports the baseline level in vessels per day, the absolute change, the percent change and its confidence interval, under both the immediate pre-event reference and the seasonal reference of §2.2. The aggregate cargo series and the all-vessel series are included so that the parts reconcile with the whole.

Where a class records no traffic at a chokepoint in the reference window, the levels are reported and the percent change is left undefined rather than divided by zero.

Intervals use the procedure of S8.1 applied within each class; the Ever Given uses the asymmetric procedure of S8.2.

S8.11 The extract, the port counts and the two capacity measures

Answers the request that the analysed sample be visible and that the capacity measure be defined. Referred to from §2.1.

The main chokepoint extract holds 11,252 rows, that is 2,813 consecutive days for each of the four chokepoints of §2.1, from 1 January 2019 to 13 September 2026, with no missing values in any reported field. The Magellan Strait series, used only in the rerouting test, is retrieved separately over the same span.

Of the 2,065 ports PortWatch publishes, 1,996 record a positive baseline call rate and 878 meet the activity threshold of at least one call per day. The exposure mapping of §2.5 then removes offshore terminals, mixed-lane ports and excluded economies, which leaves the

contrasts resting on 753 ports for the Red Sea shock, 481 for Panama and 738 for the Ever Given. The threshold is reported in §2.5.2 as a result rather than as a data-quality step, because it removes the two groups at different rates.

Transited capacity is the summed deadweight tonnage of the vessels passing on a given day. Deadweight measures carrying capacity by weight, so a container ship, a tanker and a car carrier of equal deadweight carry very different value; capacity changes are therefore read within a vessel class and never across classes. The Suez Canal Authority's "net ton" [2] is a volumetric measure of a different kind again, and no quantity in this article compares the two.