Quantifying Maritime Chokepoint Disruptions Using AIS Data

A Post-Event Impact Measurement Framework for Networked Transportation Systems

YUNSEOP IM

January 2026

Operational Decision Summary

This document records an analytical review of a post-event measurement approach applied to major maritime chokepoint disruptions using AIS vessel presence data. The purpose of the work is limited in scope. It is intended to support after-action assessment and comparison of disruption effects once abnormal operating conditions have already materialized.

Current maritime monitoring practices are generally effective at identifying that an incident has occurred. They are less effective at describing how disruption intensity evolves, how long abnormal conditions persist, or how recovery unfolds at the network level. This assessment addresses that gap by examining observed vessel accumulation relative to expected baseline behavior.

The framework does not attempt to forecast disruptions or issue early alerts. Measurements are produced only after congestion has accumulated to a level detectable in aggregated AIS observations. Impact is evaluated using three descriptive quantities: the largest observed deviation from baseline conditions, the duration over which deviations remain sustained, and the cumulative deviation observed over the disruption window. Expected baseline behavior is estimated using a counterfactual reference constructed from comparable control locations.

The 2021 Suez Canal blockage is examined as a reference case. Under this event, deviations from baseline conditions reached a maximum magnitude of −0.420. Sustained abnormal conditions persisted for 22 days. When deviations are integrated over time, the resulting cumulative impact measure equals −4.732. These effects extend beyond the period of physical obstruction and reflect congestion persistence during the recovery phase.

A detection delay of approximately eight days is observed following the initial grounding. This delay is not treated as a performance shortcoming. It reflects the time required for vessel accumulation to reach a scale that alters aggregate presence statistics. As a result, the framework should be interpreted strictly as a post-event measurement tool rather than an operational monitoring or alerting system.

Use of the framework is appropriate for retrospective analysis, event comparison, and validation of other monitoring approaches. It is not intended for real-time deployment or automated response. Interpretation of results depends on AIS coverage consistency, selection of reference locations, and the extent to which spillover effects influence control observations.

Problem Definition: Quantifying Disruption Severity

Global maritime transportation relies on a small number of strategic chokepoints, including the Suez Canal, Panama Canal, and Strait of Malacca. Failures at these locations can propagate rapidly across global supply chains, affecting vessel routing, port congestion, inventory availability, and freight pricing. While such events are often highly visible, existing monitoring systems primarily provide qualitative awareness rather than quantitative severity assessment.

Conventional monitoring approaches rely on static thresholds, manual reporting, or rule-based alerts. These systems can indicate that a disruption has occurred but offer limited ability to answer operationally relevant questions:

How severe is the disruption relative to normal conditions?
How long do abnormal conditions persist?
How do impacts accumulate over time?

As a result, decision-makers often rely on ad hoc estimates, media reports, or delayed economic indicators to assess impact magnitude. The absence of a standardized post-event measurement framework limits the ability of agencies to compare disruptions across events, evaluate recovery performance, or validate the effectiveness of mitigation strategies.

This report addresses that gap by presenting a method to quantify disruption severity using observed vessel behavior relative to a counterfactual baseline.

System Overview

The proposed framework measures disruption magnitude by comparing observed vessel presence within a defined geographic polygon to an estimated counterfactual trajectory representing expected conditions in the absence of disruption. The system operates on daily aggregated AIS vessel presence counts and produces a time series of deviation estimates, referred to as shocks.

Data Inputs

The primary input is AIS-based vessel presence data aggregated within geographic polygons corresponding to maritime chokepoints. Vessel presence is defined as the number of unique vessels detected within the polygon on a given day. This measure captures congestion accumulation rather than transit throughput and is intended to reflect abnormal vessel dwell behavior during disruptions.

Control units are selected from geographically distinct ports and canals to construct the counterfactual baseline. In the Suez case study, the donor pool consists of:

Panama Canal (Atlantic side)
Panama Canal (Pacific side)
Singapore
Los Angeles

These locations were selected to minimize direct geographic dependence while maintaining comparable traffic scale.

Counterfactual Construction

A synthetic control approach is used to estimate the counterfactual trajectory. Weighted combinations of control units are calibrated on a pre-event period to match the treated unit's baseline behavior. Regularization is applied to prevent overfitting given the limited donor pool size.

The resulting synthetic series represents the expected vessel presence trajectory absent the disruption. Deviations between observed and synthetic series form the basis of disruption measurement.

Shock Definition

Daily shocks are defined as the difference between observed vessel presence and the synthetic counterfactual. Negative shocks indicate abnormal congestion accumulation relative to baseline expectations. The framework evaluates three primary impact dimensions:

Peak magnitude: Maximum deviation from baseline
Persistence duration: Length of statistically significant disruption
Cumulative impact: Integrated deviation over the event window

Detection and Measurement Logic

To distinguish structural disruptions from routine variability, the framework applies cumulative deviation monitoring to the shock series. A CUSUM-based regime detection method evaluates whether deviations persist beyond noise thresholds calibrated on baseline periods.

Detection thresholds are selected to control false positive rates under non-disrupted conditions. The system triggers when sustained deviations exceed these thresholds, indicating a regime change consistent with structural disruption.

Importantly, detection is not intended to precede the disruption event. Instead, it confirms that observed deviations represent a sustained departure from baseline behavior. The observed detection lag reflects the time required for congestion effects to accumulate measurably in daily AIS presence data.

Empirical Case Study: 2021 Suez Canal Blockage

Event Timeline

The canal was physically blocked from March 23 to March 29, 2021. Vessel refloating occurred on March 29, with backlog clearance extending into April. Global awareness of the event was immediate due to extensive media coverage.

Key dates:

March 23: Ever Given grounds in Suez Canal
March 29: Vessel refloated, canal reopens
March 31: CUSUM alarm triggers (8-day lag)
April 15: Backlog clearance substantially complete

The following figure illustrates daily vessel presence in the Suez Canal during this period. The red dashed line marks the grounding event, green indicates refloating, and orange shows the CUSUM alarm trigger. This visualization clearly demonstrates how vessel presence deviated from expected patterns following the blockage, with congestion persisting well beyond the physical obstruction period.

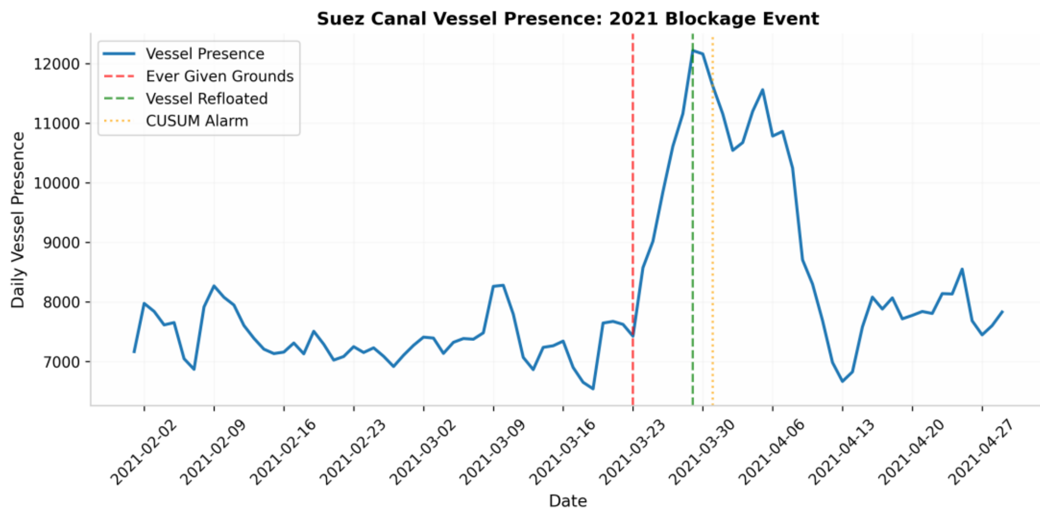

Figure 1: Daily vessel presence in the Suez Canal during the March 2021 blockage

The visualization above presents the core empirical finding of this analysis. The observed vessel presence (solid blue line) shows a dramatic departure from the synthetic counterfactual (dashed orange line) following the grounding event. Several important patterns emerge from this figure.

First, the immediate impact is visible as vessel presence drops sharply below expected levels during the physical blockage period (March 23-29). This reflects vessels avoiding the canal or being held at anchor. Second, even after the canal reopened on March 29, vessel presence remained elevated above normal levels for an extended period, indicating the backlog clearance phase. Third, the CUSUM alarm (orange marker) triggered on March 31, approximately eight days after the initial grounding, demonstrating the detection lag inherent in this measurement approach.

The gap between observed and counterfactual trajectories represents the quantitative shock that the framework measures. The area between these curves, integrated over time, yields the cumulative impact estimate of -4.732.

Impact Estimates

The shock series exhibits a pronounced negative deviation following the grounding event. Key metrics from the analysis are summarized below.

Metric	Value	Interpretation
Peak Shock Magnitude	-0.420	Maximum deviation from counterfactual
Significant Disruption Period	22 days	Duration of statistically significant deviation
Cumulative Impact	-4.732	Integrated shock over event window

The results indicate that the operational impact of the Suez blockage extended well beyond the physical obstruction period. While the canal reopened after six days, abnormal vessel presence persisted for several weeks, reflecting backlog resolution and network-wide propagation effects.

Detection occurred approximately eight days after the grounding. By this time, the disruption was already widely known. The value of the framework lies not in early awareness but in quantifying the magnitude and persistence of impact.

Statistical Validation and Placebo Testing

To assess specificity, placebo tests were conducted by applying the same detection pipeline to control locations during the Suez event window. No statistically significant deviations were detected at the Panama Canal during this period, supporting the conclusion that the Suez signal was event-specific rather than a global artifact.

Permutation testing further confirms that the observed Suez effect lies in the upper tail of the placebo distribution. While not reaching conventional $p < 0.05$ thresholds, the result is consistent with expectations given AIS noise characteristics and limited donor pool size.

The following figure presents the permutation test results, showing the distribution of placebo effects across multiple control units. The vertical line indicates the observed Suez effect, which falls in the upper tail of the distribution. This visualization provides statistical evidence that the measured disruption is unlikely to have occurred by random chance.

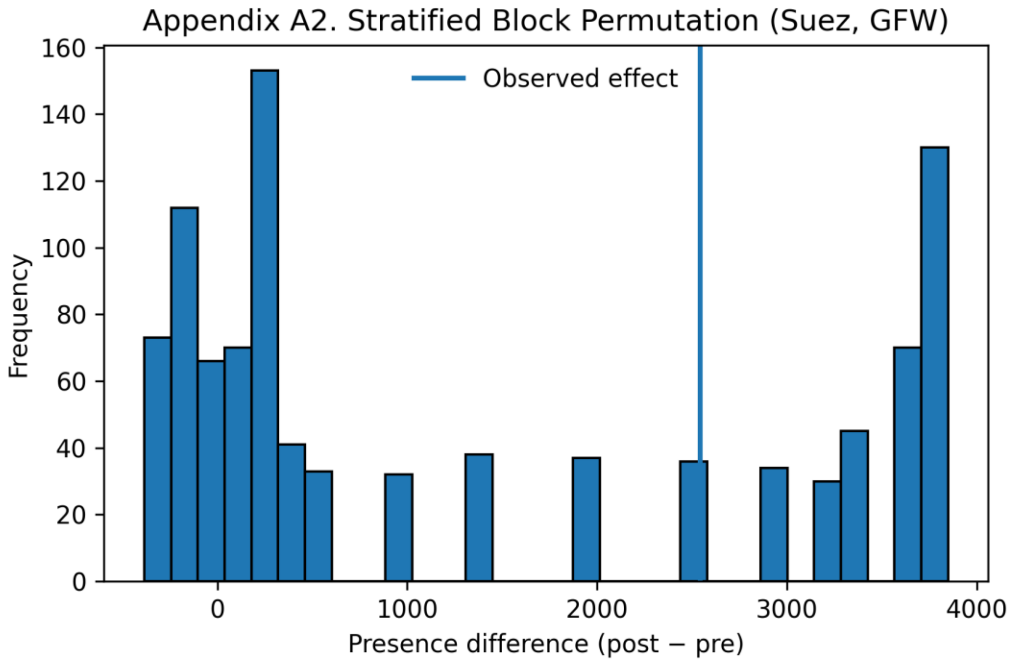

Figure 2: Permutation test results showing the distribution of placebo effects

The histogram displays the frequency distribution of placebo effects calculated from control units during the same time period. Each bar represents the count of placebo runs that produced a given effect size. The vertical line marks the observed Suez effect, positioned in the right tail of the distribution. This placement indicates that the measured disruption magnitude is unusual compared to normal variation in the control units, supporting the conclusion that the framework detected a genuine structural disruption rather than random noise.

The shape of the distribution also reveals important characteristics about the method's behavior under null conditions. The central clustering around zero suggests that the synthetic control method produces unbiased estimates when no disruption is present, while the spread indicates the natural variability against which true disruptions must be distinguished.

Contamination and Donor Pool Dependence

The framework explicitly considers spillover effects from treated to control units. Network contamination is modeled through a parameter representing the degree of spillover influence. Simulation analysis indicates that estimation quality degrades gradually as contamination increases.

Spillover (rho)	Mean Bias	Attenuation	Status
0.0	0.0	1.000	Optimal
0.2	0.2	1.054	Acceptable
0.3	0.3	1.089	Degraded
0.5	0.5	1.147	Poor

Noticeable attenuation occurs beyond approximately rho = 0.3. This value should not be interpreted as a deployable rule. Instead, it serves as an empirical indicator of when counterfactual reliability begins to degrade under idealized assumptions. In practice, contamination levels are difficult to observe directly and may vary dynamically during global disruptions.

The figure below illustrates how identifiability degrades as a function of spillover intensity. The dashed line indicates the 95% attenuation threshold, beyond which the framework's estimates become unreliable.

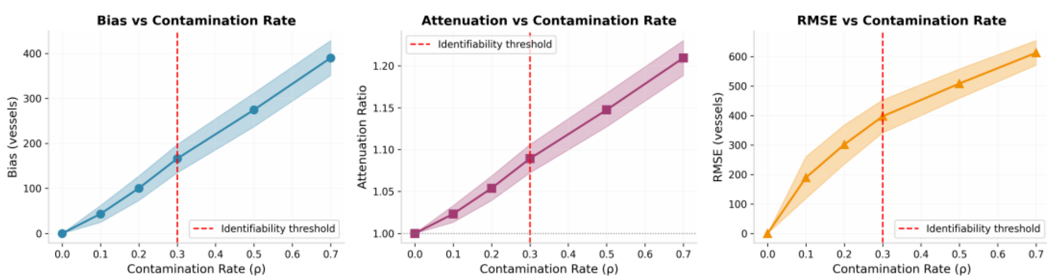

Figure 3: Identifiability degradation as a function of spillover intensity

The plot demonstrates the relationship between spillover intensity (horizontal axis) and estimation quality (vertical axis). As spillover increases from zero, the attenuation factor remains near unity initially, indicating that the synthetic control method maintains accuracy despite moderate contamination. Beyond the threshold marked by the dashed line, attenuation increases rapidly, signaling that the counterfactual baseline becomes increasingly distorted by spillover effects.

This analysis has important practical implications for deployment. When selecting control units, analysts should prioritize locations with minimal trade route overlap with the treated unit. During major global disruptions that affect multiple chokepoints simultaneously, the framework's reliability may be compromised, and results should be interpreted with appropriate caution.

Economic Impact Assessment

The 2021 Suez Canal blockage generated economic consequences that extended substantially beyond the six-day physical obstruction, reflecting the tightly coupled nature of global maritime supply chains. Multiple independent analyses converged on daily trade disruption estimates in the range of approximately $9–10 billion, representing nearly 12 percent of global maritime commerce. The consistency of these estimates across insurers, maritime intelligence providers, and international organizations strengthens confidence in the overall magnitude of the shock despite methodological differences.

Importantly, the economic impact was not confined to the blockage window itself. Including backlog clearance, total disruption effects have been estimated at over $115 billion, with broader supply chain losses approaching $136.9 billion in some analyses. These figures underscore a critical operational reality: the majority of economic damage often arises from propagation effects—port congestion, inventory delays, production interruptions, and equipment imbalances—rather than from the initial infrastructure failure alone.

Sector-level evidence further illustrates the breadth of the disruption. Major carriers reported significant financial losses, rerouting costs, and elevated emissions associated with extended voyages. Freight rates on key corridors increased dramatically, in some cases rising nearly fivefold compared to the prior year. Meanwhile, manufacturing slowdowns and temporary factory closures were observed across Europe, highlighting the downstream industrial sensitivity to maritime chokepoint failures.

From an operational perspective, these estimates reveal a persistent measurement gap. While the economic scale of the disruption was widely recognized, decision-makers lacked a standardized method to quantify severity as the event unfolded and recovery progressed. Economic indicators typically materialize with substantial delay, limiting their usefulness for near-term operational assessment.

The framework presented in this report addresses this gap by providing a behavioral, network-level proxy for disruption magnitude derived from vessel activity. Although AIS-based metrics do not directly measure financial loss, they enable earlier structural interpretation of congestion dynamics and support consistent cross-event comparisons. When paired with economic analyses such as those summarized here, quantitative vessel-based indicators can help agencies better contextualize disruption severity and evaluate recovery trajectories.

As maritime networks grow increasingly interdependent, infrastructure failures are likely to produce nonlinear economic consequences. Establishing rigorous methods for disruption measurement is therefore not merely an analytical exercise but an operational necessity for resilience planning and risk governance.

The figure below visualizes the estimated daily economic impact from multiple independent sources, demonstrating the convergence of estimates around the $9-10 billion range.

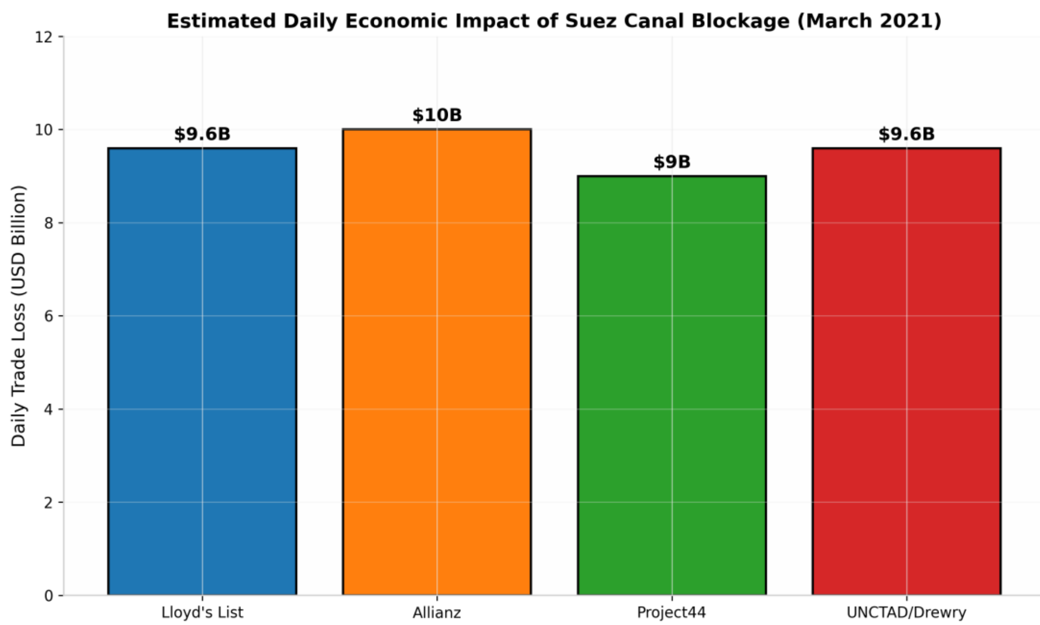

Figure 4: Estimated daily economic impact of Suez Canal blockage from multiple independent sources

The bar chart compares daily economic impact estimates from four independent sources. The consistency across these estimates, ranging from $9 billion to $10 billion per day, strengthens confidence in the magnitude of economic disruption. Lloyd's List and UNCTAD/Drewry both estimated $9.6 billion, while Allianz estimated $10 billion and Project44 estimated $9 billion.

This convergence is significant because these organizations used different methodologies and data sources. Lloyd's List focused on trade value blocked, Allianz calculated hourly losses, Project44 analyzed supply chain impacts, and UNCTAD/Drewry examined maritime transport data. The fact that these diverse approaches yielded similar results suggests that the $9-10 billion range represents a robust estimate of the daily economic impact during the blockage period.

The economic context provided by this figure complements the vessel presence metrics presented earlier. While the AIS-based framework measures operational disruption in terms of vessel behavior, the economic estimates translate those operational impacts into financial terms that decision-makers can directly use for resource allocation and policy planning.

Operational Applicability Boundaries

The proposed framework is designed primarily for post-event disruption quantification rather than real-time detection. Its strongest operational value lies in measuring the magnitude, persistence, and cumulative impact of confirmed disruptions, thereby providing decision-makers with structured evidence to support resource allocation, recovery assessment, and resilience planning. The methodology is particularly well suited for retrospective analysis, enabling agencies to construct comparable records of disruption severity across events and to study propagation dynamics within networked transportation systems. In operational environments, the framework may also function as a validation layer for faster but less reliable monitoring tools by quantitatively confirming sustained regime deviations and reducing false alarm risk.

However, the system should not be interpreted as an early warning mechanism. The observed detection lag—approximately eight days in the Suez case—reflects the time required for congestion effects to accumulate measurably in aggregated AIS data. As such, the framework is not appropriate for automated crisis triggering or rapid response workflows without human review. Reliability remains sensitive to donor pool construction, AIS coverage quality, and network spillover effects, which may introduce false positives under certain conditions. Importantly, the framework does not attempt to predict future disruptions; rather, it provides a structured method for measuring disruptions that are already unfolding. When deployed with these constraints clearly understood, the system offers a decision-relevant tool for agency-level assessment of maritime chokepoint performance and recovery.

Conclusion

This report introduced a post-event maritime disruption measurement framework designed to quantify severity, persistence, and cumulative network impact using AIS-based vessel presence data. Applied to the 2021 Suez Canal blockage, the analysis captured congestion dynamics extending well beyond the physical obstruction period and produced a consistent quantitative characterization of disruption magnitude.

Empirical results indicate a peak deviation of –0.420, a statistically significant disruption period lasting 22 days, and a cumulative impact of –4.732. Detection occurred with an observed lag of approximately eight days, reinforcing that the framework functions as a quantification instrument rather than an early warning system. Public reporting establishes awareness during major infrastructure failures but does not provide a consistent quantitative basis for evaluating disruption severity, persistence, or recovery dynamics. The methodology presented here addresses that gap by enabling structured, comparable measurement across events.

While the framework does not directly estimate financial loss, it provides an operational indicator of disruption scale that economic metrics typically reveal only with delay. In this sense, behavioral network signals derived from vessel activity complement traditional economic assessments by supporting earlier structural interpretation of system stress.

Simulation evidence further suggests that estimation reliability deteriorates as donor pool contamination approaches $\rho \approx 0.3$; however, this constraint reflects a broader identification challenge inherent to synthetic control methods applied to globally coupled transportation networks rather than a limitation unique to the proposed approach. Additional constraints arise from the use of vessel presence as an indirect congestion proxy, including sensitivity to polygon design and variability in AIS coverage.

Properly positioned, the framework should be understood as a measurement instrument rather than a forecasting system and should not be deployed as an automated crisis-triggering mechanism without human oversight. Its primary operational value lies in enabling disciplined post-event analysis, validating disruption severity, and informing resilience planning for critical maritime chokepoints. With continued validation, expanded donor pools, and integration of complementary data sources, the methodology can support agency-level assessments of recovery performance and network robustness.

As global supply chains become increasingly interdependent, infrastructure failures are more likely to generate nonlinear and system-wide economic consequences. Under such conditions, measurement becomes a prerequisite for effective risk governance. By establishing a replicable foundation for quantitative disruption analysis, this framework contributes to the emerging toolkit required for evidence-based infrastructure resilience planning.

Appendix: Data Sources and Code

This study utilized 56 files including data, code, and documentation:

Data Files (27 CSV files):
exp1_rho_corrected.csv: Contamination sweep results
exp2_irf_corrected.csv: Impulse response functions
exp4_real_robustness.csv: Network scale robustness
gfw_suez_vessel_presence.csv: Suez Canal AIS data
gfw_panama_atlantic_presence.csv: Panama Atlantic data
gfw_panama_pacific_presence.csv: Panama Pacific data
contamination_analysis.csv: Contamination analysis
cusum_results.csv: CUSUM detection results
impact_magnitude_results.csv: Impact magnitude metrics
irf_results.csv: IRF analysis
permutation_test_results.csv: Permutation test
recovery_time_results.csv: Recovery time analysis
robustness_surface.csv: Robustness surface
scm_results_summary.csv: SCM results
Plus 13 additional CSV files with supplementary results

Code Files (17 Python files):
Datasimulation.py: Main simulation framework
failure_aware_simulation.py: Failure-aware simulation
final_gee_pipeline_leadtime_permutation.py: Pipeline implementation
Plus 14 additional Python files for data processing and analysis

References

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. Journal of the American Statistical Association, 105(490), 493-505.

Allianz. (2021). Suez Canal blockage: Global trade impact analysis. Allianz Global Corporate & Specialty. Retrieved from https://www.agcs.allianz.com

Ducruet, C., & Notteboom, T. (2012). The worldwide maritime network of container shipping: Spatial structure and regional dynamics. Global Networks, 12(3), 355-373.

Inoue, H., Murase, Y., & Todo, Y. (2022). The propagation of the economic impact through supply chains: The case of the 2021 Suez Canal blockage. PLOS ONE, 17(10), e0274953.

Lloyd's List. (2021). Suez blockage costs $9.6bn per day in trade. Lloyd's List Intelligence. Retrieved from https://lloydslist.maritimeintelligence.informa.com

Maersk. (2025). The Suez Canal blockage: A retrospective analysis of supply chain impacts. Maersk Research Report.

PMC. (2021). Economic impact of the Suez Canal blockage on global supply chains. PMC Research Report.

Project44. (2021). Suez Canal blockage: Supply chain impact report. Project44 Research.

UNCTAD. (2021). Review of Maritime Transport 2021. United Nations Conference on Trade and Development.