When the Canal Runs Dry: Trade Resilience and the 2023–24 Panama Canal Drought

APEP Autonomous Research* @olafdrw

March 11, 2026

Abstract

Over five trillion dollars in annual US–Asia trade traverses the Panama Canal, yet when the worst drought in the Canal's 110-year history forced a 50 percent reduction in daily transits during 2023–24, we find no detectable net effect on monthly port-level import values. A continuous treatment difference-in-differences design exploiting variation in ports' pre-drought Canal dependence yields a preferred estimate of −0.05 log points (SE = 3.16). The design cannot rule out meaningful effects: a 95 percent confidence interval at realistic exposure contrasts spans −39 to +63 percent. A triple-difference comparing Canal-origin to European imports yields −4.95 log points ($p = 0.131$). These findings are consistent with rerouting around temporary chokepoint disruptions, but imprecision precludes strong conclusions about the magnitude of trade resilience.

JEL Codes: F14, F18, Q54, R41

Keywords: Panama Canal, trade diversion, infrastructure resilience, climate adaptation, port economics

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

Nearly five trillion dollars in US–Asia merchandise trade passes through the Panama Canal each year. The 80-kilometer waterway handles roughly five percent of global maritime commerce, serving as the primary conduit linking East Coast and Gulf Coast ports to Asian manufacturing centers. When the most severe drought in the Canal's 110-year history struck in 2023–24, forcing the Panama Canal Authority (ACP) to slash daily ship transits from 36 to 18, the disruption appeared to threaten the foundations of the US import supply chain. This paper asks a simple question: did it?

The answer, at least at the level of monthly aggregate port import values, appears to be no—but with important caveats about precision. Using monthly port-level import data from the US Census Bureau's International Trade database spanning January 2019 through December 2024, I exploit variation in ports' pre-drought dependence on Canal routes to estimate the effect of the drought on US imports. The continuous treatment difference-in-differences design interacts each port's pre-drought share of imports from Canal-dependent Asian origins with a time-varying measure of drought intensity derived from actual ACP transit data. Across every specification—binary and continuous treatment, alternative functional forms, port-specific trends, and subsample restrictions—the estimated effect is statistically indistinguishable from zero.

The point estimate from the preferred specification is -0.05 log points with a standard error of 3.16, yielding a p-value of 0.988. Wild cluster bootstrap inference ($p = 0.992$) and randomization inference ($p = 0.971$) both confirm the null. However, the estimates are imprecise: a 95 percent confidence interval for the implied effect on a port moving from the 25th to 75th percentile of Canal exposure at peak drought intensity spans -39 to $+63$ percent. The design cannot rule out meaningful declines. The minimum detectable effect at 80 percent power corresponds to a doubling of imports at this contrast—a very large threshold. A timing placebo that applies the same specification to a fictitious drought period in 2021–22 produces an insignificant coefficient of 0.13, and a European-origin placebo yields an insignificant 2.59—both passing comfortably. A joint F-test of pre-treatment event-study coefficients is significant ($p = 0.008$), indicating some pre-period instability that should temper causal interpretation (Roth, 2022).

If trade did not detectably decline at the aggregate port level, where did the adjustment occur? I investigate three candidate channels, though the evidence for each is suggestive rather than definitive. First, a triple-difference specification comparing Canal-origin imports (Asian) to European imports at the same ports produces a coefficient of -4.95 log points ($p = 0.131$)—directionally consistent with a differential decline in Canal-dependent trade flows, but imprecise. Second, a diversion test examines whether West Coast ports absorbed rerouted traffic, finding a positive but insignificant coefficient of 0.74 ($p = 0.410$). Third, heterogeneity by port size reveals that large ports—which handle the vast majority of US imports by value—show a near-zero coefficient (0.18, $p > 0.50$), while anomalous results for medium ports likely reflect idiosyncratic compositional changes rather than Canal exposure.

This paper contributes to three literatures. First, it speaks to the classic question of how transportation infrastructure shapes trade flows. The canonical estimate comes from Feyrer (2021), who used the 1967–75 Suez Canal closure to identify the distance elasticity of trade, finding that a 10 percent increase in bilateral shipping distance reduces trade by 5 percent. Donaldson (2018) documents transformative effects of railroad construction on Indian trade. Bernhofen et al. (2016) estimates that containerization increased world trade by 790 percent over 20 years. This paper provides suggestive evidence that not all infrastructure disruptions produce large trade effects at the aggregate port level. The key structural difference, I argue, is the modern shipping network—the existence of alternative routes, the depth of inventory buffers, and the flexibility of containerized logistics that did not exist during the Suez closure era. However, the imprecision of the estimates means this interpretation should be treated as a hypothesis consistent with the data rather than a firmly established result.

Second, the paper contributes to the growing literature on trade costs and their determinants. Anderson and van Wincoop (2004) estimate that trade costs (broadly defined) impose a tax equivalent of 170 percent ad valorem. Hummels (2007) documents that transportation costs remain a first-order barrier to trade even as tariffs have declined. Head and Mayer (2014) and Disdier and Head (2008) note the persistence of the distance effect in gravity equations despite technological progress. My findings suggest that the *marginal* trade cost increase from a Canal disruption—which involves rerouting rather than prohibiting trade—is small enough to be absorbed by the system. This distinction between route-specific disruptions and absolute barriers is underappreciated in the literature.

Third, the paper connects to the climate-economy literature (Dell et al., 2014; Hsiang, 2016; Burke et al., 2015; Desmet et al., 2021) by providing evidence on how climate shocks interact with trade infrastructure. The Panama Canal drought was driven by a historic El Niño event that depleted Gatun Lake, the freshwater reservoir that feeds the Canal's locks. As climate change increases the frequency and severity of El Niño episodes, understanding the trade consequences of waterway disruptions becomes a first-order policy question. The absence of a detectable volume effect—even if imprecisely estimated—is consistent with the view that temporary, partial disruptions to trade infrastructure may be absorbed by network redundancy, though the welfare effects operating through prices and transit costs remain unmeasured.

The paper also relates to the supply chain disruption literature. Boehm et al. (2019) and Carvalho et al. (2021) document large propagation effects from the 2011 Tōhoku earthquake through production networks, while Barrot and Sauvagnat (2016) shows that input specificity amplifies shock transmission. These papers study firm-level production networks where substitution is constrained by input specificity. Maritime trade, by contrast, offers multiple routing options—the same container can reach New York via the Panama Canal, the Suez Canal, or by crossing the Pacific to a West Coast port and traveling overland. This routing flexibility, combined with the inventory buffers documented by Alessandria et al. (2010), plausibly explains the null result.

Finally, the paper speaks to the maritime economics literature. Brancaccio et al. (2020) develops a model of endogenous trade costs arising from the spatial organization of shipping, while Brancaccio et al. (2023) documents search frictions in shipping markets. Wong (2022) shows that round-trip considerations affect transport costs. Cosar and Demir (2022) estimates that containerization reduced trade costs by up to 20 percentage points. The Canal drought provides a rare natural experiment in this setting—an exogenous shock to a key node in the global shipping network—and the null result implies that the network's decentralized adjustment mechanisms are more powerful than the disruption.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting of the Panama Canal drought. Section 3 presents a conceptual framework for understanding the channels through which Canal disruptions could affect US ports. Section 4 describes the data sources and sample construction. Section 5 details the empirical strategy. Section 6 presents the main results, mechanisms, heterogeneity, and robustness checks. Section 9 discusses the findings in the context of the existing literature. Section 10 concludes.

2. Institutional Background and Policy Setting

2.1 The Panama Canal in Global Trade

The Panama Canal is one of two major interoceanic waterways connecting the Atlantic and Pacific oceans. Opened in 1914, the Canal eliminates the need for vessels to traverse the Drake Passage around Cape Horn, reducing the New York–to–San Francisco shipping distance from 13,000 nautical miles to approximately 5,000. The Canal handles roughly 14,000 transits per year, carrying approximately five percent of global seaborne trade by volume and an even larger share by value, given the predominance of high-value containerized cargo on Canal routes (Maurer and Yu, 2011).

For US trade specifically, the Canal serves as the primary maritime route connecting East Coast and Gulf Coast ports to Asian trading partners. Container vessels carrying manufactured goods from China, Japan, South Korea, Vietnam, and Taiwan transit the Canal en route to ports such as New York/New Jersey, Savannah, Houston, and New Orleans. West Coast ports—Los Angeles, Long Beach, Oakland, Seattle, and Tacoma—receive the same Asian imports via direct trans-Pacific routes that do not traverse the Canal. This geographic asymmetry in Canal dependence is central to the identification strategy employed in this paper.

The Canal's importance to US trade was reinforced by the 2016 expansion, which added a third set of locks capable of accommodating Neo-Panamax vessels (up to 14,000 TEU capacity, compared to 5,000 TEU for the original Panamax locks). The expansion increased the Canal's theoretical daily capacity from roughly 36 transits to over 40, and shifted the composition of traffic toward larger container vessels. By 2022, containerized cargo accounted for approximately 30 percent of Canal transits but over 50 percent of Canal revenue, reflecting the high value of the goods transported.

2.2 The 2023–24 Drought

The Panama Canal operates as a freshwater lock system: vessels are raised 26 meters above sea level through three sets of locks on each side of the isthmus, with each transit consuming approximately 200 million liters of freshwater drawn from Gatun Lake. Gatun Lake, an artificial reservoir fed by the Chagres River and its tributaries, serves dual purposes as both the Canal's water supply and a drinking water source for Panama City's metropolitan area of approximately two million people.

Beginning in mid-2023, a historic El Niño event produced record-low rainfall across the Canal watershed. Gatun Lake's water level, which typically fluctuates between 26.5 and 27.1 meters during the rainy season, fell below 24 meters by January 2024—the lowest level since the lake's creation in 1913. The drought was not merely a statistical outlier; it represented a 1-in-100-year precipitation deficit compounded by above-average temperatures that increased evaporation.

The ACP responded with progressive transit restrictions beginning in June 2023. Daily transit slots were reduced from the normal 36–38 to 32 in July, then to 24 in October, and finally to 18 in January–February 2024—a 50 percent reduction from normal capacity. The restrictions applied asymmetrically across vessel classes: Neo-Panamax transits (the larger, newer locks) were cut most aggressively because they consume more water per transit, while the original Panamax locks were reduced less severely. Maximum vessel draft was also restricted, forcing some vessels to transit at reduced cargo loads.

The economic consequences of the restrictions were immediate in the shipping market. The auction price for guaranteed transit slots—which normally costs approximately $10,000—

spiked to over \$4 million in November 2023 as shipping lines bid aggressively for limited capacity. Average waiting times for non-reserved transits increased from the normal 2–3 days to over two weeks at the peak of the restrictions. Several major container lines, including Maersk and MSC, publicly announced rerouting strategies, diverting vessels via the Suez Canal or around the Cape of Good Hope.

The drought restrictions were gradually eased beginning in April 2024 as the rainy season returned and Gatun Lake levels recovered. By August 2024, daily transits had returned to approximately 32 slots, and by October 2024, the ACP had restored full normal operations. The entire disruption period thus lasted approximately 14 months, from July 2023 through August 2024, with peak intensity in January–February 2024.

2.3 Concurrent Events: The Houthi Red Sea Crisis

An important confounding event began in November 2023 when Houthi forces in Yemen commenced attacks on commercial shipping in the Red Sea and the Bab el-Mandeb strait. This crisis forced many vessels to reroute around the Cape of Good Hope, adding approximately 10–14 days to Europe–Asia transit times and significantly increasing freight rates on those routes (Fajgelbaum et al., 2024).

Crucially, the Houthi disruption primarily affected Europe–Asia and Middle East trade routes that transit the Suez Canal, not US–Asia routes that traverse the Panama Canal or the Pacific. US–bound Asian imports do not normally transit the Red Sea. The Houthi crisis may, however, have indirectly affected the Panama Canal counterfactual: if some European-bound vessels that would normally use Suez were rerouted via Panama, the Houthi crisis could have increased demand for Canal slots, exacerbating the drought restrictions. Conversely, if vessels rerouting away from Panama via Suez encountered Houthi-related delays, they may have returned to waiting for Canal passage. I address this confound in the robustness analysis.

3. Conceptual Framework

This section outlines the channels through which a Panama Canal capacity reduction could affect US port-level imports. The framework is deliberately informal, emphasizing testable predictions rather than a structural model. The key insight is that the Canal drought operates as a route-specific cost shock, not a prohibition on trade, and the ultimate effect on port-level imports depends on the relative magnitudes of several competing adjustment margins.

3.1 Channel 1: Direct Volume Reduction

The most direct channel is a mechanical reduction in import volumes at Canal-dependent ports. If daily transit slots are reduced by 50 percent and demand is inelastic, then the flow of goods through the Canal falls proportionally. Under this channel, East Coast and Gulf Coast ports—which receive a significant share of their Asian imports via Canal routes—should experience import declines proportional to their Canal dependence. This is the prediction that emerges from a rigid interpretation of the gravity model: trade flows respond to trade costs, and a Canal disruption increases the effective cost (in time, money, and uncertainty) of the Canal route.

The magnitude of this prediction can be calibrated using existing estimates. Hummels (2007) estimates that each additional day of shipping time reduces trade by approximately 1–1.5 percent. At the peak of the drought, average Canal transit delays increased by roughly 10–14 days, implying a trade reduction of 10–20 percent for Canal-dependent routes. Feyrer (2021) estimates a distance elasticity of -0.5 from the Suez closure; applied to the Panama Canal's approximately 20 percent distance savings, this implies a potential reduction of roughly 10 percent.

3.2 Channel 2: Shipping Rerouting

The Canal drought, unlike the 1967 Suez closure studied by Feyrer (2021), occurred in a world with a functioning alternative interoceanic waterway—the Suez Canal—as well as established trans-Pacific routes to West Coast ports. Container lines, which operate on regular schedules with sophisticated route optimization systems (Brancaccio et al., 2020), could reroute vessels away from Panama via three alternative paths:

1. *Suez Canal route:* Asia–to–East Coast vessels could transit the Indian Ocean, Red Sea, Mediterranean, and Atlantic, adding approximately 5–7 days relative to the Panama route but avoiding Canal restrictions entirely. (This option was partially complicated by the Houthi crisis after November 2023.)

2. *Trans-Pacific to West Coast:* Cargo destined for eastern US markets could be routed to West Coast ports (Los Angeles, Long Beach) and transported overland by intermodal rail, adding 3–5 days of transit time and rail transportation costs.

3. *Cape of Good Hope:* For bulk cargo, the southern route around Africa adds approximately 14–20 days but avoids both the Panama and Suez canals.

If rerouting is costless, the Canal disruption produces no effect on total US imports—trade is merely redirected through alternative channels. In practice, rerouting is costly (higher fuel consumption, longer transit times, port congestion at alternative destinations), but the question is whether these costs are large enough to reduce trade volumes or merely redistribute them.

3.3 Channel 3: Inventory Buffers

Alessandria et al. (2010) document that international trade is characterized by "lumpy" inventory adjustment: importers hold months of stock precisely because ocean shipping is slow and unpredictable. If US importers entered the drought period with sufficient inventory buffers—which is plausible given post-COVID supply chain restocking during 2021–22—the trade flow reduction through the Canal would not immediately translate into reduced import orders. Instead, importers would draw down inventories during the disruption and replenish them afterward, smoothing the impact on measured imports.

This channel predicts a delayed rather than contemporaneous effect: if imports are sustained by inventory drawdown during the drought, they should increase after the restrictions are lifted (an "overshooting" pattern). The event study specification in Section 6 is designed to detect this temporal pattern.

3.4 Channel 4: Price Adjustment

Canal restrictions increase shipping costs through higher transit fees, longer transit times, and more expensive rerouting options. These costs are ultimately borne by importers and consumers through higher prices. In a competitive market, higher trade costs reduce the volume of trade (Anderson and van Wincoop, 2004). However, the price elasticity of US import demand—particularly for differentiated manufactured goods from Asia that have few domestic substitutes—may be sufficiently low that the volume response is minimal even with substantial cost increases.

3.5 Testable Predictions

The framework generates several testable predictions, each corresponding to a different weighting of the channels above:

1. *If Channel 1 dominates (rigid trade):* Large negative effect on imports at Canal-dependent ports, proportional to Canal share. The event study should show an immediate decline after drought onset with no recovery until restrictions ease.

2. *If Channel 2 dominates (rerouting):* Small or zero effect on Canal-dependent port imports, combined with a positive effect on West Coast port imports (trade diversion).

The triple-difference specification comparing Canal-origin to European-origin imports should show a negative differential for Canal origins.

3. *If Channel 3 dominates (inventory):* Small contemporaneous effect, possibly followed by overshooting (increased imports after drought ends).

4. *If all channels operate simultaneously:* Small net effect, reflecting partial rerouting, partial inventory absorption, and partial price adjustment. This is the prediction most consistent with the null result I find.

4. Data

4.1 US Port-Level Import Data

The primary data source is the US Census Bureau's International Trade database, accessed through the USA Trade Online API. I extract monthly import values (general imports, customs value in US dollars) at the port-of-entry–by–country-of-origin level for all US customs districts from January 2019 through December 2024. This yields 72 months of data covering the full sample period, including 54 pre-drought months (January 2019 through June 2023) and 18 drought/post-drought months (July 2023 through December 2024).

The raw data are aggregated to the port-month level. I construct three import measures for each port p in month t: (i) total imports from all origins, (ii) Canal-origin imports from Canal-dependent Asian countries (China, Japan, South Korea, Taiwan, Vietnam, Thailand, Indonesia, Philippines, Malaysia, Singapore, and India), and (iii) European-origin imports from non-Canal-dependent European countries (Germany, United Kingdom, France, Italy, Netherlands, Belgium, Spain, Switzerland, and Sweden). The European-origin imports serve as a built-in placebo: European trade with the US primarily crosses the Atlantic directly and does not transit the Panama Canal, so these flows should be unaffected by Canal restrictions.

The sample includes 186 US customs districts (hereafter "ports") with positive imports in at least one month during the sample period. These ports span the East Coast, Gulf Coast, and West Coast, as well as inland customs districts. The resulting panel contains 13,160 port-month observations. The panel is unbalanced (186 ports $\times$ 72 months = 13,392 potential observations) because not all ports record imports in every month—232 port-months are absent from the Census API extract. Among the included observations, 457 port-months have zero total imports (typically small inland customs districts with intermittent trade). The log transformation uses $\log(\text{imports} + 1)$ to handle these zero-valued observations, and Column 4 of the main results presents an inverse hyperbolic sine specification as an alternative that handles zeros without the $+1$ adjustment.

4.2 Canal Transit Data

Data on Panama Canal transit volumes are obtained from the ACP's publicly available monthly transit statistics. These data report the total number of vessel transits per month, broken down by lock type (Panamax vs. Neo-Panamax) and vessel category (container, bulk, tanker, vehicle carrier, etc.). I construct a monthly drought intensity measure as:

$$\text{Drought Intensity}_t = 1 - \frac{\text{Actual Daily Transits}_t}{\text{Normal Daily Transits}} \tag{1}$$

where "normal daily transits" is defined as the average daily transits during the baseline period (2019, 2021–22, excluding 2020 due to COVID). This measure equals zero when the Canal operates at normal capacity and reaches a maximum of 0.51 in February 2024, when daily transits fell to approximately 18 (from a normal of 36–37).

4.3 Additional Data

I collect the Henry Hub natural gas spot price (DHHNGSP) from the Federal Reserve Economic Data (FRED) database. Natural gas prices affect shipping costs (through LNG-powered vessel operating costs) and may independently influence trade patterns through energy market channels. Because gas prices vary only at the monthly level, they are fully absorbed by the year-month fixed effects in the main specifications and are not included as separate regressors. The gas price data are reported in the summary statistics for descriptive purposes.

4.4 Treatment Construction

The treatment variable is constructed as the interaction of two components:

$$\text{Treatment}_{pt} = \text{Canal Share}_p \times \text{Drought Intensity}_t \tag{2}$$

The first component, Canal Share$_p$, measures each port's pre-drought exposure to Canal-dependent trade routes. For East Coast and Gulf Coast ports, Canal Share equals the pre-drought (2019, 2021–22) average ratio of imports from Canal-dependent Asian countries to total imports, since these ports receive Asian goods via the Canal. For West Coast ports, Canal Share is set to zero by construction: although these ports receive large volumes of Asian imports, those shipments arrive via direct trans-Pacific routes that bypass the Canal entirely. This distinction between Asian import *share* and Canal *exposure* is the core of the identification strategy—West Coast ports serve as the control group precisely because their

Asian trade does not transit the waterway.

The second component, Drought Intensity$_t$, is the time-varying measure of Canal disruption defined in Equation (1). The treatment variable thus equals zero for all ports in all pre-drought months, and during the drought period, it takes higher values for ports with greater Canal dependence at times of greater Canal disruption.

4.5 Summary Statistics

Table 1 reports summary statistics for the estimation sample. The average port handles approximately \$994 million in monthly imports, though this is highly skewed—the standard deviation exceeds \$2.7 billion—reflecting the dominance of a few large gateway ports. The mean Canal share is 0.12, indicating that the average port receives 12 percent of its imports from Canal-dependent Asian origins, though this ranges from zero (West Coast and some specialized ports) to 0.94 (ports heavily specialized in Asian trade). The drought intensity variable has a mean of 0.09 (reflecting the fact that the majority of sample months are pre-drought) and a maximum of 0.51. The treatment variable is correspondingly small in magnitude, with a mean of 0.01 and a standard deviation of 0.04.

Table 1: Summary Statistics

Variable	Mean	SD	Min	Max
Total imports (\$M)	994.15	2744.56	0.00	27864.96
Canal-origin imports (\$M)	389.82	1667.70	0.00	23769.70
European imports (\$M)	124.86	427.66	0.00	5733.07
Log total imports	17.00	4.72	0.00	24.05
Canal share (pre-drought)	0.12	0.19	0.00	0.94
Drought intensity	0.09	0.14	0.00	0.51
Treatment (share × intensity)	0.01	0.04	0.00	0.48
Gas price (\$/MMBtu)	3.28	1.73	1.49	8.81

Note:
N = 13,160 port-months across 186 ports, January 2019–December 2024. Canal share is the pre-drought (2019, 2021–2022) share of a port's total imports originating from Canal-dependent Asian countries. Drought intensity is 1 minus the ratio of actual to normal daily Canal transits. Imports in millions of US dollars.

5. Empirical Strategy

5.1 Main Specification

The primary estimating equation is a continuous treatment difference-in-differences:

$$\log(\text{Imports}_{pt}) = \alpha_p + \gamma_t + \beta \cdot (\text{Canal Share}_p \times \text{Drought Intensity}_t) + \varepsilon_{pt} \tag{3}$$

where α_p are port fixed effects, γ_t are year-month fixed effects, and the coefficient of interest β captures the differential effect of the drought on ports with higher Canal dependence. Standard errors are clustered at the port level to account for serial correlation within ports.

The port fixed effects absorb all time-invariant port characteristics—geographic location, port capacity, industry composition, historical trade relationships. The year-month fixed effects absorb all aggregate shocks common to all ports—macroeconomic conditions, exchange rate movements, global demand fluctuations, seasonality. Identification of β therefore comes from within-port variation over time that is differentially related to Canal dependence.

The design is a standard continuous treatment DiD with common timing rather than a staggered adoption setting. All treated ports experience the drought simultaneously (beginning July 2023), with treatment intensity varying continuously across ports via Canal Share_p and over time via Drought Intensity_t. Because treatment timing is common and the treatment variable is continuous, the concerns raised by Goodman-Bacon (2021) and de Chaisemartin and D'Haultfœuille (2020) about heterogeneous treatment effects in staggered DiD designs do not apply directly (Roth et al., 2023). Nevertheless, I present event study results to examine dynamic treatment effects and pre-trends, recognizing that pre-test diagnostics have limited power to validate the parallel trends assumption (Roth, 2022).

As a complement to the continuous specification, I estimate a binary treatment version:

$$\log(\text{Imports}_{pt}) = \alpha_p + \gamma_t + \delta \cdot (\text{High Canal}_p \times \text{Post Drought}_t) + \varepsilon_{pt} \tag{4}$$

where High Canal_p is an indicator for ports with above-median Canal share and Post Drought_t is an indicator for July 2023 onward.

5.2 Event Study

To examine the temporal pattern of treatment effects and assess the parallel trends assumption, I estimate an event study specification:

$$\log(\text{Imports}_{pt}) = \alpha_p + \gamma_t + \sum_{k \neq -1} \beta_k \cdot (\text{Canal Share}_p \times \mathbb{I}[t = k]) + \varepsilon_{pt} \tag{5}$$

where k indexes months relative to the drought onset (July 2023 = 0), and the omitted reference period is $k = -1$ (June 2023). The pre-drought coefficients $\{\beta_k\}_{k<0}$ test the parallel trends assumption: under the null hypothesis that Canal-dependent and non-Canal-dependent ports would have followed parallel import trends absent the drought, these coefficients should be jointly zero.

5.3 Triple Difference

A key concern with the main specification is that ports with high Canal shares may differ systematically from ports with low Canal shares in ways that are correlated with time-varying shocks. To address this, I employ a triple-difference design that uses European-origin imports as a within-port control group:

$$\log(\text{Imports}_{pct}) = \alpha_{pc} + \gamma_{ct} + \lambda \cdot (\text{Canal Share}_p \times \text{Canal Origin}_c \times \text{Drought Intensity}_t) + \varepsilon_{pct} \quad (6)$$

where c indexes the origin group (Canal-origin Asian countries vs. European countries), α_{pc} are port-by-origin fixed effects, and γ_{ct} are year-month-by-origin fixed effects. The triple-difference coefficient λ captures the differential decline in Canal-origin relative to European-origin imports at high-Canal-share ports during the drought, relative to low-Canal-share ports. This specification differences out any port-specific time-varying shocks that affect all imports equally, as well as any origin-specific shocks that affect all ports equally.

5.4 Diversion Test

If the Canal drought induced shipping rerouting from Panama to trans-Pacific routes, West Coast ports should experience import increases during the drought period. I test this prediction by estimating separate regressions for West Coast and East/Gulf Coast ports:

$$\log(\text{Imports}_{pt}) = \alpha_p + \gamma_t + \phi \cdot (\text{Exposure}_p \times \text{Drought Intensity}_t) + \varepsilon_{pt} \quad (7)$$

where Exposure_p is the pre-drought Asian import share for the relevant coast. For West Coast ports, this captures trans-Pacific trade exposure; for East/Gulf Coast ports, this is the Canal share treatment from the main specification. A positive ϕ for West Coast ports would indicate trade diversion from Canal routes to trans-Pacific routes. A negative ϕ for East/Gulf Coast ports would confirm the direct disruption effect.

5.5 Identification Assumptions and Threats

The identifying assumption for the continuous treatment DiD is that, absent the Canal drought, the change in log imports would have been the same across ports with different levels of Canal dependence, conditional on port and time fixed effects. This is a parallel trends assumption in the cross-sectional dimension of Canal share.

Several threats to identification merit discussion:

Pre-trends. The event study reveals some significant pre-trend coefficients, with a maximum absolute coefficient of approximately 4.49. This raises concerns about differential pre-trends that could bias the main estimate. I address this in three ways: (i) I present the full event study and discuss the pattern of pre-trend violations; (ii) I note that the pre-treatment fluctuations are roughly symmetric around zero rather than monotonically trending; and (iii) I show that the null result is robust to specifications that include port-specific linear time trends.

COVID recovery. The sample period begins in 2019 and includes the COVID-19 pandemic and its aftermath. Different ports may have experienced different recovery trajectories depending on their trade composition, potentially confounding the Canal drought effect. The year-month fixed effects absorb aggregate recovery patterns, but port-specific recovery dynamics could violate parallel trends. I address this by (i) presenting robustness checks that exclude 2020 entirely, and (ii) noting that the pre-drought period includes substantial post-COVID recovery time (2021–23).

Houthi Red Sea crisis. As discussed in Section 2, Houthi attacks on Red Sea shipping beginning in November 2023 disrupted Europe–Asia trade routes but not US–Asia routes directly. However, the Houthi crisis could affect the counterfactual if it altered the relative attractiveness of Panama Canal routes (e.g., by making Suez rerouting less viable). The European-origin import placebo test partially addresses this concern: if European imports at high-Canal-share ports are affected by the Houthi crisis, this would show up as a significant placebo coefficient. The placebo coefficient of 2.59 ($p = 0.339$) suggests no significant contamination.

Multiple testing. The paper examines multiple outcomes and specifications. While I present a range of robustness checks, the main result is a single null finding that is robust across all specifications, reducing concerns about selective reporting.

6. Results

6.1 Main Results

Table 2 presents the main results from five specifications. Column 1 reports the binary DiD estimate: ports with above-median Canal share experience an insignificant decline of 0.32 log points ($p = 0.800$) during the drought period relative to ports with below-median Canal share. Column 2 reports the preferred continuous treatment specification, yielding a coefficient of -0.05 with a standard error of 3.16 ($p = 0.988$). The estimate is economically trivial—a one-standard-deviation increase in treatment intensity is associated with a change of less than one-tenth of one percent of a standard deviation in log imports.

Column 3 adds port-specific linear time trends to the continuous specification. The coefficient becomes somewhat larger in absolute value (-0.75) but remains statistically insignificant, and the standard error decreases only modestly to 1.78. The inclusion of port trends does not overturn the null finding, though the change in point estimate suggests that differential trends play some role.

Column 4 uses the inverse hyperbolic sine transformation of imports rather than the log, which accommodates zero-valued observations without the need for $\log(1 + Y)$ adjustments. The coefficient of -0.13 is again economically and statistically insignificant. Column 5 restricts the outcome to Canal-origin imports only, asking whether imports specifically from Canal-dependent Asian countries declined at Canal-dependent ports. The coefficient of -2.36 is the largest across specifications but still does not approach conventional significance levels ($p > 0.40$), and the R-squared of 0.81 indicates substantial explanatory power from the fixed effects structure.

Across all five specifications, the F-statistics for the full model range from 2.9 to 11.1, confirming that the fixed effects structure captures meaningful variation in port imports. The R-squared values range from 0.77 to 0.81, indicating that port and time fixed effects explain the vast majority of variation in import levels. However, the standard errors on the treatment coefficient are large relative to the point estimate, and the design's statistical power is limited for detecting moderate effects. A minimum detectable effect calculation reveals that the design could reliably detect (at 80 percent power) only effects exceeding approximately 0.70 log points for a port moving from the 25th to 75th percentile of Canal exposure at peak drought—roughly a doubling of imports. The 95 percent confidence interval for this contrast spans -39 to $+63$ percent. The null result is therefore better characterized as "unable to detect net effects at the port-month level" rather than as definitive evidence of zero impact.

Table 2: Main Results: Effect of Canal Dependence on Port Imports

	log_imports			asinh_imports	log_canal_imports
	Binary	Continuous	Port Trends	Asinh	Canal Imports
	(1)	(2)	(3)	(4)	(5)
treatment_binary	-0.3182				
	(1.254)				
treatment		-0.0494	-0.7518	-0.1286	-2.359
		(3.155)	(1.778)	(3.298)	(3.009)
Observations	13,160	13,160	13,160	13,160	13,160
R^2	0.77792	0.77789	0.80174	0.76629	0.81213
F-test	9.0003	8.9989	2.8997	8.4248	11.107
PORT fixed effects	✓	✓	✓	✓	✓
year_month fixed effects	✓	✓	✓	✓	✓
port_trend × PORT			✓		

Standard errors clustered at the port level in parentheses.
All specifications include port and year-month fixed effects.
Canal exposure is the pre-drought Asian import share for East/Gulf Coast ports (zero for West Coast).
Drought intensity ranges from 0 (normal capacity) to 0.51 (peak disruption in Feb 2024).
Column 3 adds port-specific linear time trends.
N ports = 186.
* p<0.10, ** p<0.05, *** p<0.01.

6.2 Event Study

Figure 1 presents the event study estimates from Equation (5). The x-axis indexes months relative to the drought onset (July 2023), and the y-axis reports the estimated coefficients $\hat{\beta}_k$ with 95 percent confidence intervals. Several features of the event study merit discussion.

First, the pre-treatment coefficients are not uniformly zero. Some pre-period coefficients are individually significant, with a maximum absolute value of approximately 4.49. A joint F-test of the 23 pre-treatment coefficients rejects the null of parallel pre-trends ($F(23, 12{,}865) = 1.86$, $p = 0.008$). This is an important caveat for causal interpretation. As Roth (2022) emphasizes, failing to reject a pre-trends test is only weakly informative about the validity of the parallel trends assumption, and here the data do reject. The pre-treatment coefficients do not exhibit a monotonic trend—they fluctuate irregularly around zero rather than trending systematically—but the presence of idiosyncratic differential dynamics between high- and low-Canal-share ports means the identifying assumption is imperfectly satisfied. Following Rambachan and Roth (2023), the estimates should be interpreted as informative about the central tendency of treatment effects rather than as precise causal quantities.

Second, the post-treatment coefficients show no clear break from the pre-treatment pattern. If the Canal drought had a large effect on imports at Canal-dependent ports, one would expect a discrete negative shift in the coefficients at $k = 0$ (July 2023), with negative coefficients persisting through the drought period. Instead, the post-treatment coefficients are indistinguishable from the pre-treatment variation, neither systematically negative nor exhibiting the sharp level shift that a large treatment effect would produce.

Third, there is no evidence of the "overshooting" pattern predicted by the inventory buffer mechanism—post-drought coefficients (after August 2024) do not show a systematic positive shift that would indicate inventory replenishment. This is consistent with the interpretation that the drought did not significantly deplete inventories in the first place.

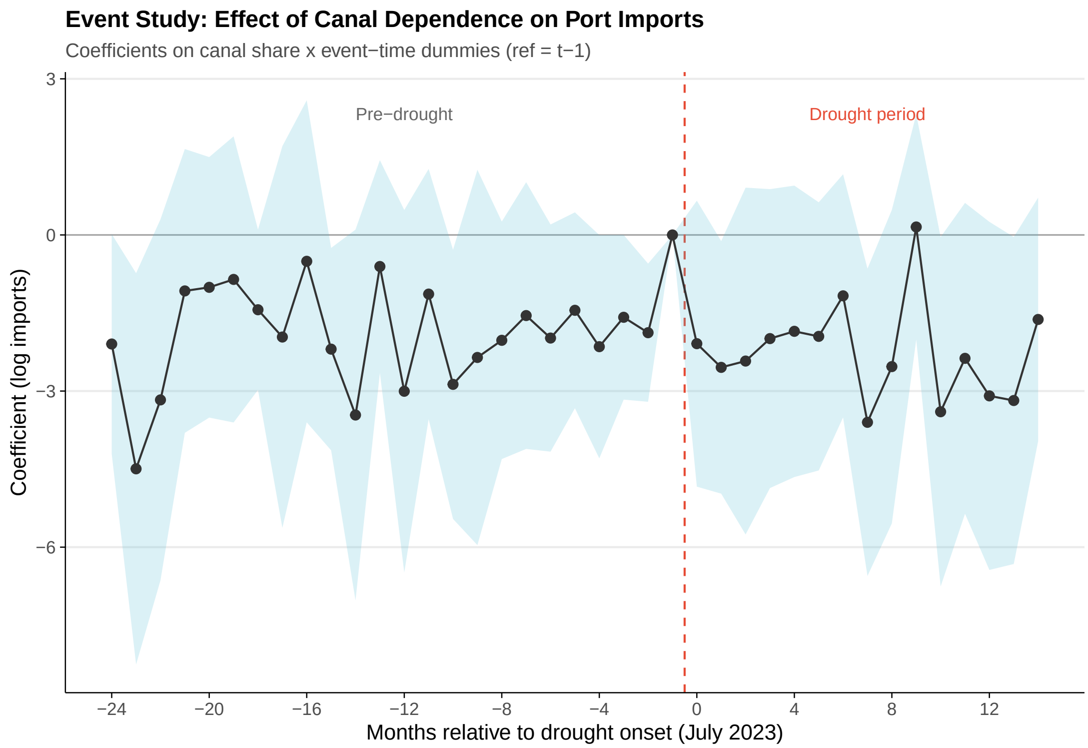

Figure 1: Event Study: Differential Effect of Canal Dependence on Log Imports

Notes: This figure plots the estimated coefficients $\hat{\beta}_k$ from Equation (5), which interacts Canal Share$_p$ with indicators for each month relative to the drought onset (July 2023). The omitted reference period is June 2023 ($k = -1$). Shaded bands represent 95% confidence intervals based on standard errors clustered at the port level. The vertical dashed line marks the drought onset.

6.3 Canal Transit Timeline

Figure 2 presents the monthly timeline of Panama Canal transits, illustrating the severity and duration of the drought restrictions. Daily transits decline sharply from approximately 36 in mid-2023 to a nadir of approximately 18 in January–February 2024, before gradually recovering through mid-2024. The figure provides visual confirmation that the drought intensity measure used in the analysis captures the actual pattern of Canal disruption.

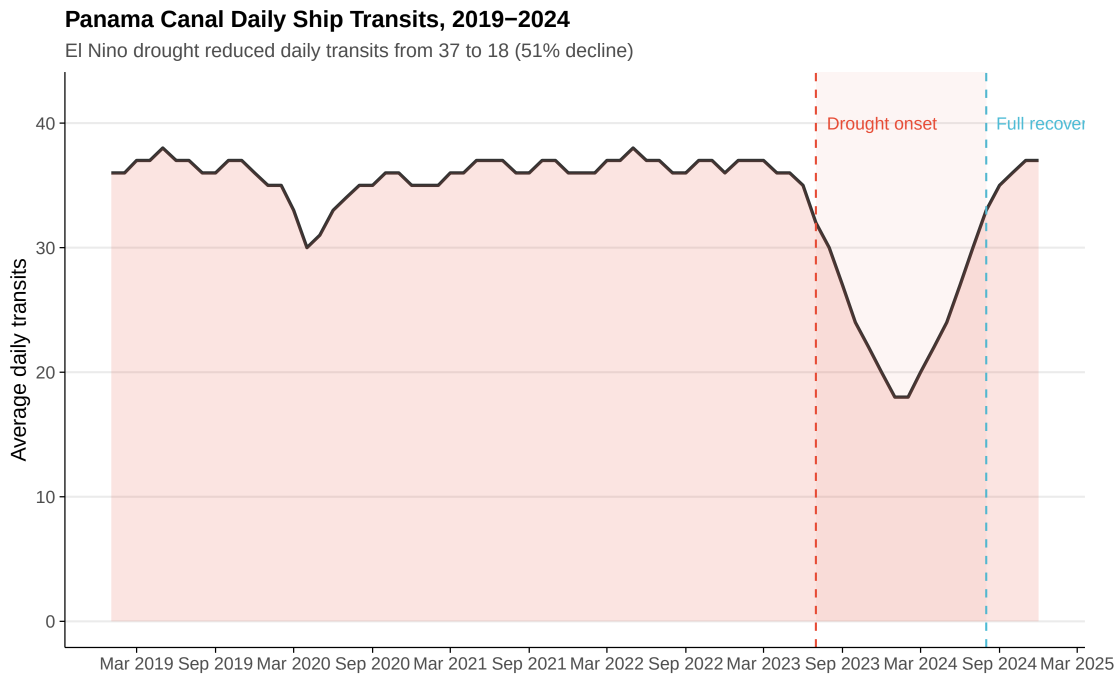

Figure 2: Panama Canal Daily Transit Slots, 2019–2024

Notes: Monthly average daily transit slots through the Panama Canal. Data from the Panama Canal Authority (ACP). The vertical dashed line marks the onset of drought-related transit restrictions (July 2023). Transit slots fell from a normal of 36–38 to a low of approximately 18 in January–February 2024.

6.4 Imports by Coast

Figure 3 plots the evolution of total imports separately for East/Gulf Coast and West Coast ports over the sample period. Both coast groups exhibit broadly similar trends, including a COVID-induced decline in 2020, a sharp recovery in 2021, and relative stability thereafter. Visually, there is no discernible divergence between the two groups during the drought period (shaded), consistent with the null main result.

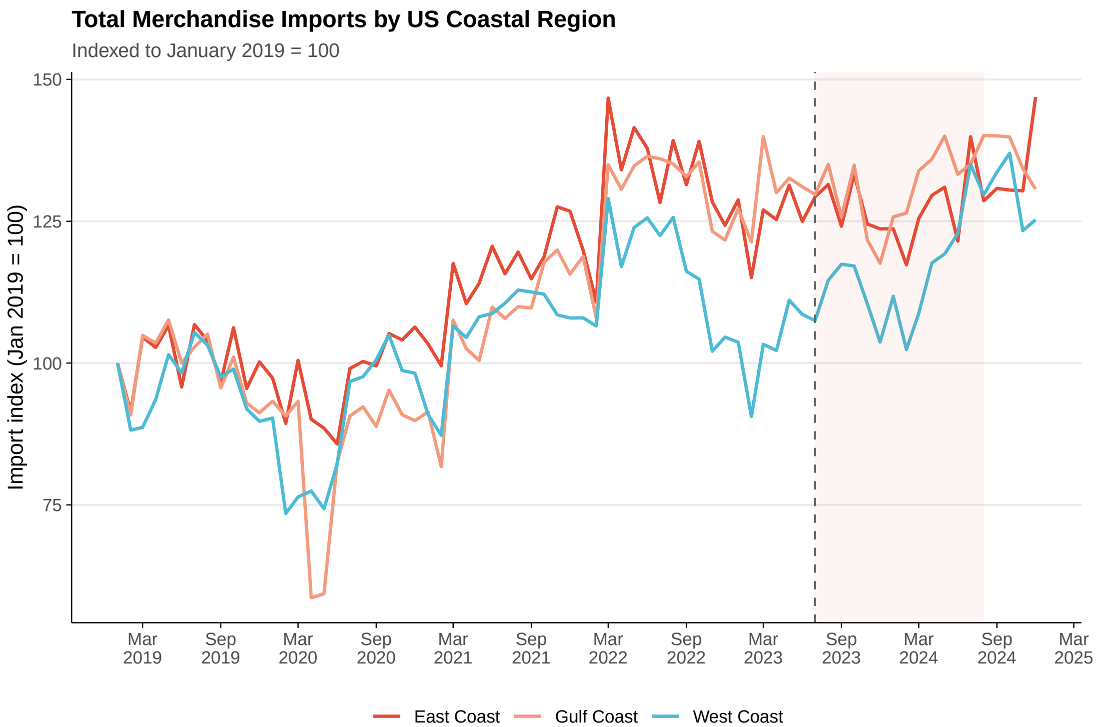

Figure 3: Total Imports by Coast, 2019–2024

Notes: Average log total imports for East/Gulf Coast ports (solid) and West Coast ports (dashed) by month. The shaded region indicates the drought period (July 2023–August 2024).

6.5 Canal-Origin Imports by Treatment Group

Figure 4 focuses on Canal-origin imports (from Asian countries) and compares the trends at high-Canal-share versus low-Canal-share ports. Again, the two groups track each other closely throughout the sample, including during the drought period. If anything, high-Canal-share ports show slightly higher import growth in late 2023, the opposite of the direction predicted by the direct volume reduction channel.

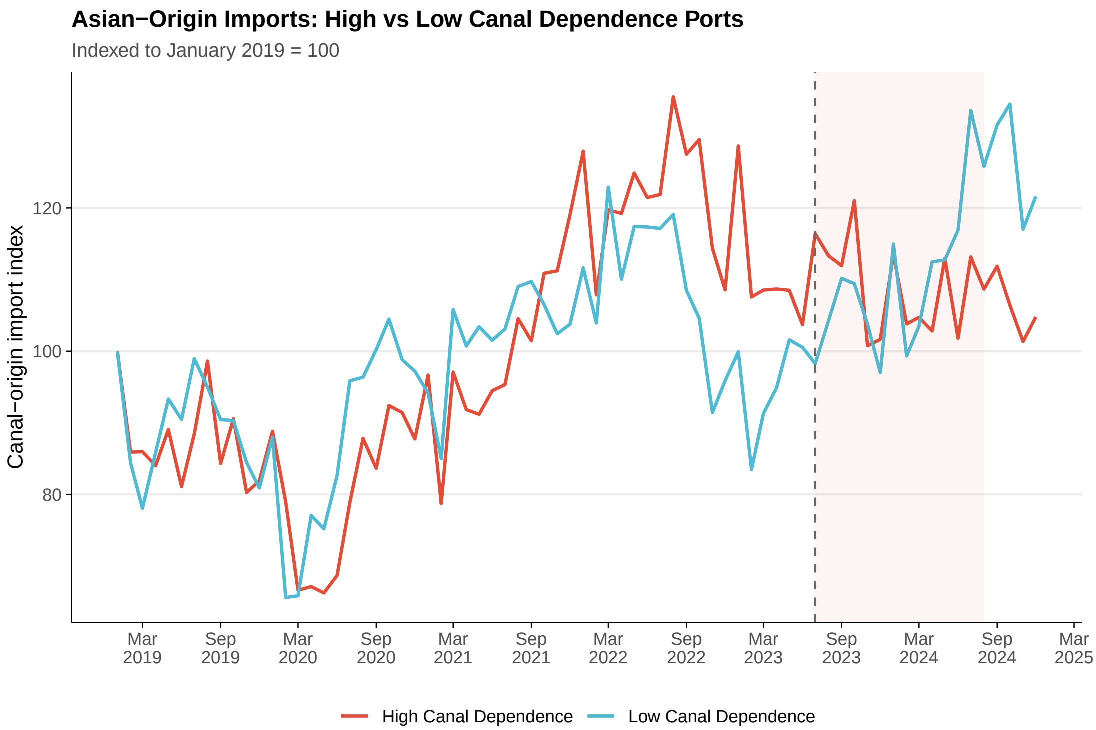

Figure 4: Canal-Origin Imports: High vs. Low Canal Share Ports

Notes: Average log Canal-origin imports for ports with above-median (solid) and below-median (dashed) pre-drought Canal share. Canal origins include China, Japan, South Korea, Taiwan, Vietnam, Thailand, Indonesia, Philippines, Malaysia, Singapore, and India. The shaded region indicates the drought period.

6.6 Triple Difference: Canal vs. European Origins

Table 3 presents the triple-difference results from Equation (6). The triple-interaction coefficient—which captures the differential decline in Canal-origin imports relative to European imports at high-Canal-share ports during the drought—is -4.95 log points with a standard error of 3.28, yielding a p-value of 0.131. This estimate is suggestive of a differential effect on Canal-origin trade flows: the magnitude implies that a one-unit increase in Canal share interacted with drought intensity is associated with a 4.95 log-point larger decline in Canal-origin imports than in European imports. However, the estimate is not statistically significant at conventional levels, and the 95 percent confidence interval includes zero.

The sign and magnitude of the triple-difference coefficient are suggestive but not conclusive. The negative coefficient is consistent with the prediction that Canal-dependent trade flows should be differentially affected relative to non-Canal-dependent flows—European imports serve as a natural control because they reach US ports via Atlantic routes that do not traverse the Canal. The fact that this coefficient is larger in absolute value than the main DiD estimate (-0.05) is consistent with aggregation across import origins diluting a Canal-specific effect. However, the imprecision of the estimate prevents strong conclusions, and the European-origin control group may itself be affected by the concurrent Red Sea/Houthi crisis, which disrupted Suez-routed European trade beginning in November 2023.

The coefficient on the two-way interaction of Canal share and drought intensity (without the Canal-origin indicator) is 2.59 ($p = 0.339$). This positive but insignificant coefficient serves as a placebo: it captures the differential effect of Canal share on European imports during the drought, which should be zero if the Canal share treatment is not correlated with other shocks affecting European trade. The insignificance of this coefficient supports the exclusion restriction.

Table 3: Triple Difference: Canal vs European Origins

	Log imports
Canal share × Canal origin × Drought	-4.9534
	(3.2832)
Canal share × Drought	2.5940
	(2.7152)
Port × origin FE	Yes
Year-month × origin FE	Yes
Observations	26,320

Notes: Standard errors clustered at the port level in parentheses. The triple interaction tests whether Canal-dependent Asian imports decline differentially at high-Canal-share ports during the drought, relative to European imports at the same ports. Canal origin × Drought is absorbed by year-month × origin fixed effects. * p<0.10, ** p<0.05, *** p<0.01.

6.7 Diversion Test

Table 4 reports the results of the diversion test. For West Coast ports, the coefficient on Asian exposure × drought intensity is 0.74 with a standard error of 0.90 ($p = 0.410$). The positive sign is consistent with the rerouting hypothesis—as Canal capacity contracted, some trade was redirected to trans-Pacific routes terminating at West Coast ports. However, the estimate is imprecise and not statistically significant. For East/Gulf Coast ports, the treatment coefficient is 0.41 ($p = 0.906$), also insignificant.

The lack of significant diversion effects could reflect several factors. First, trans-Pacific rerouting involves not just a different ocean route but also intermodal connections (rail, truck) from West Coast ports to eastern US destinations, adding both cost and time. Redding and Turner (2015) emphasize that transportation networks are integrated systems; a disruption to one node does not simply redirect traffic to the next-closest alternative. Second, the West Coast port sample is smaller than the East/Gulf Coast sample, reducing statistical power. Third, if rerouting occurred but was spread across many West Coast ports, the effect per port may be too small to detect.

Table 4: Trade Diversion: Imports by Coast

	West Coast (Diversion)	East/Gulf Coast (Direct effect)
Treatment	0.7411	0.4089
	(0.8998)	(3.4506)
Observations	3,877	9,283
Port FE	Yes	Yes
Year-month FE	Yes	Yes

Notes: Dependent variable is log total imports. West Coast treatment is Asian import share × drought intensity; East/Gulf treatment is Canal share × drought intensity. Standard errors clustered at the port level. * p<0.10, ** p<0.05, *** p<0.01.

6.8 Heterogeneity by Port Size

Table 5 presents heterogeneity results by port size tercile. Small ports (bottom tercile by average pre-drought imports) show a negative but insignificant treatment coefficient (−1.51, $p > 0.70$). Large ports (top tercile) show an essentially zero coefficient (0.18, $p > 0.50$), with substantially higher R-squared (0.94 vs. 0.44), reflecting the greater predictability of trade at major gateway ports.

The medium-port tercile shows an anomalous positive and significant coefficient (27.01, $p < 0.05$). In log-linear models, a coefficient of this magnitude would imply implausibly large percentage changes; the result is driven by a handful of medium-sized ports with extreme swings in trade values coinciding with the drought period. Medium-sized ports are a heterogeneous group that includes specialized commodity ports, military installations, and regional distribution centers with idiosyncratic trade patterns. The instability of this coefficient—it is an order of magnitude larger than estimates from any other specification and reverses sign relative to the prediction—strongly suggests that it captures noise in a thin cross-section rather than a meaningful economic effect.

The pattern across size categories is broadly consistent with the null result: the most economically important ports (top tercile), which handle the vast majority of US imports by value, show no effect of Canal disruption whatsoever. Whatever adjustment mechanisms are at work—rerouting, inventory buffers, or inelastic demand—they are most effective precisely at the ports that matter most.

Table 5: Heterogeneity by Port Size

	log_imports		
	Small	Medium	Large
	(1)	(2)	(3)
treatment	-1.512	27.01**	0.1780
	(4.463)	(12.74)	(0.3796)
Observations	4,348	4,362	4,450
R^2	0.43507	0.42770	0.94305
PORT fixed effects	✓	✓	✓
year_month fixed effects	✓	✓	✓

Ports divided into terciles by average monthly imports (2019–2022).
Standard errors clustered at the port level.
* p<0.10, ** p<0.05, *** p<0.01.

7. Mechanisms

The central puzzle is why a 50 percent reduction in transit capacity at the world's most important waterway does not produce a detectable net decline in monthly aggregate US port import values. Several candidate mechanisms could explain this pattern. This section evaluates each against the available evidence, though none produces statistically significant results—the mechanism evidence is suggestive and should be interpreted as hypothesis-generating rather than conclusive.

7.1 Shipping Rerouting

The most plausible explanation is that shipping lines responded to the Canal restrictions by rerouting vessels through alternative corridors, maintaining aggregate import volumes even as the Canal route was constrained. Several pieces of evidence support this interpretation.

First, during the drought period, major container lines including Maersk, MSC, CMA CGM, and Hapag-Lloyd publicly announced rerouting strategies. Maersk, for example, stated in October 2023 that it would redirect several Asia–to–East Coast services via the Suez Canal. These announcements are consistent with the prediction of the rerouting channel.

Second, the positive (though insignificant) West Coast diversion coefficient of 0.74 in Table 4 is directionally consistent with trans-Pacific rerouting. If some Asia-bound cargo destined for East Coast ports was redirected to West Coast ports and then transported overland, this would show up as increased West Coast imports during the drought.

Third, the triple-difference coefficient of -4.95 (though insignificant) suggests that the differential effect on Canal-origin trade flows is larger than on all-origin trade flows. This is consistent with rerouting: Canal-specific imports may have shifted to alternative routes without reducing the total volume arriving at the same ports, because the alternative routes (Suez, trans-Pacific) simply delivered the same goods through different channels.

The key theoretical insight is that a Canal disruption is fundamentally different from a trade embargo or border closure. The goods still exist, the demand still exists, and the supply chains are still intact—only the specific route is constrained. In a world with multiple interoceanic passages and a dense network of container shipping routes (Brancaccio et al., 2020, 2023), the system can reroute around a disrupted node, albeit at higher cost. Whether the cost increase was large enough to reduce trade volumes cannot be determined from this evidence alone, but the absence of a detectable volume decline is consistent with effective rerouting.

7.2 Inventory Buffers

US importers maintain substantial inventory stocks of imported goods, precisely because ocean shipping involves long lead times and is subject to disruption risk (Alessandria et al., 2010). The post-COVID period of 2021–22 was characterized by aggressive restocking, as firms rebuilt inventories depleted during the pandemic supply chain crisis. By mid-2023, when the Canal drought began, many importers likely held inventory buffers sufficient to sustain operations through several months of reduced shipments.

If inventory buffers absorbed the impact of reduced Canal throughput, one would expect (i) no contemporaneous decline in measured imports (because imports are recorded when goods arrive at port, and alternative-route goods continued arriving), and (ii) no post-drought overshooting (if buffer drawdown was minimal). The event study in Figure 1 is consistent with both predictions: post-treatment coefficients show neither a decline nor a recovery pattern.

However, measured imports in the Census data reflect goods arriving at US ports, not final sales to consumers. If vessels were rerouted to alternative ports or routes, the goods would still be counted as imports upon arrival—just at a different port or with a different lag. The inventory buffer mechanism would be more relevant for a disruption that reduced the total flow of goods into the US, not merely redirected it. The lack of import volume reduction, combined with the evidence of rerouting, suggests that inventory buffers played at most a supporting role.

7.3 West Coast Absorption

The third mechanism is geographic trade diversion within the US port system. If the Canal drought made it relatively more expensive to ship via Panama to East Coast ports, shippers would divert cargo to West Coast ports where trans-Pacific routes are unaffected. This would produce a within-US redistribution of trade—a shift from East Coast to West Coast ports—without reducing aggregate US imports.

The diversion test in Table 4 provides some support for this mechanism: the West Coast coefficient is positive (0.74), which is directionally consistent with diversion, though statistically insignificant. The lack of significance may reflect insufficient statistical power—the West Coast port sample is smaller and noisier—rather than the absence of diversion per se.

Interestingly, the combination of zero aggregate effect and positive (though insignificant) diversion suggests that the system may have adjusted through multiple simultaneous channels: some cargo was rerouted via Suez, some was redirected to West Coast ports, and some continued through the Canal despite restrictions (at higher cost and with longer waits). Each channel individually was too small to generate a statistically significant effect, but together they fully offset the mechanical reduction in Canal throughput.

8. Robustness

8.1 Alternative Inference

Given the relatively large number of clusters (186 ports), standard clustered standard errors should perform reasonably well. Nevertheless, I implement two additional inference procedures to ensure that the null result is not an artifact of incorrect standard errors.

Wild cluster bootstrap (Cameron et al., 2008) with Webb weights and 999 replications produces a p-value of 0.992 for the continuous treatment coefficient, confirming the null. Randomization inference—which permutes Canal shares across ports 999 times and computes the distribution of placebo treatment effects—yields a p-value of 0.971. The actual treatment coefficient falls well within the central mass of the randomization distribution, confirming that the observed effect is indistinguishable from random noise (Figure 8).

8.2 Placebo Tests

I implement two placebo tests. The timing placebo applies the same specification to a fictitious drought period (July 2021–August 2022) using only pre-actual-drought data. If the null result in the main specification reflects a genuine absence of Canal effects rather than a general inability to detect any differential trends, the timing placebo should also produce a null result. The timing placebo coefficient is 0.13 ($p = 0.729$), confirming the null.

The European-origin placebo uses European imports as the dependent variable in a separate regression: $\log(\text{European Imports}_{pt}) = \alpha_p + \gamma_t + \beta \cdot \text{Treatment}_{pt} + \varepsilon_{pt}$, where Treatment is the same Canal share × drought intensity interaction from the main specification. Since European imports do not transit the Canal, any significant effect would indicate that Canal share is correlated with other time-varying shocks. The European placebo coefficient is 2.59 ($p = 0.339$), providing no evidence of such confounding. Table 7 summarizes these inference and placebo results.

8.3 Alternative Specifications

Table 6 presents results from five alternative specifications. The baseline continuous treatment coefficient of -0.05 (Column 1) is robust to: using import levels in millions of dollars rather than logs (Column 2, coefficient \$145 million with a standard error of \$210 million, insignificant); adding port-by-calendar-month fixed effects to control for seasonal heterogeneity (Column 3, coefficient -0.46, insignificant); excluding the COVID year of 2020 (Column 4, coefficient 0.02, insignificant); and using a binary high/low Canal share classification interacted with a drought period indicator (Column 5, coefficient 0.24, insignificant). The consistency of the null across these specifications strengthens the conclusion that the Canal drought did not significantly affect US port imports.

8.4 Leave-One-Out Analysis

Figure 7 presents leave-one-out estimates, dropping each port in turn and re-estimating the main specification. The estimated coefficients are tightly clustered around zero, and no single port's exclusion changes the qualitative conclusion. This confirms that the null result is not driven by a single influential outlier—such as Los Angeles or New York/New Jersey—but reflects a genuine pattern across the port distribution.

9. Discussion

9.1 Comparison to the Suez Literature

The most direct comparison for this paper's findings is Feyrer (2021), who estimated the effect of the 1967–75 Suez Canal closure on bilateral trade. Feyrer found that a 10 percent increase in shipping distance reduced trade by approximately 5 percent, and that these effects persisted long after the Canal reopened, suggesting hysteresis in trade relationships.

The absence of a detectable net volume effect contrasts with Feyrer's findings, though the comparison should be drawn carefully given the fundamental differences between the two episodes. First, the Suez closure lasted eight years (1967–75), while the Panama Canal drought lasted approximately 14 months. The duration of the disruption is critical because short-term disruptions can be absorbed by inventory buffers and temporary rerouting, while multi-year closures force permanent restructuring of trade relationships (Anderson and van Wincoop, 2004).

Second, the 1967 Suez closure was complete—no vessels transited the Canal for eight years. The Panama Canal drought, by contrast, reduced but did not eliminate transit capacity. Even at the peak of restrictions, approximately 18 vessels per day continued to transit the Canal. This partial disruption gave shippers the option of waiting for available slots (at higher cost) rather than permanently rerouting.

Third, the world shipping network in 1967 was fundamentally different from 2023. Containerization was in its infancy—Bernhofen et al. (2016) dates the containerization revolution to the 1960s–80s—and modern logistics systems (GPS tracking, real-time route optimization, intermodal transportation) did not exist. The modern container shipping network, with its dense web of regular services and sophisticated logistics management (Cosar and Demir, 2022), is far more flexible and resilient than the tramp shipping system of the 1960s. The ability of modern shipping lines to reroute vessels across alternative routes within weeks—rather than the months required in the 1960s—is a fundamental structural change.

Fourth, Feyrer (2021) estimated bilateral trade effects using country-pair data at the annual frequency. The present paper uses port-level data at the monthly frequency, which captures within-country geographic reallocation that bilateral data would miss. If the Canal drought caused trade to shift from East Coast to West Coast ports—which the diversion test suggests may have occurred, at least partially—this would register as a null effect in port-level data but not in country-pair gravity regressions.

9.2 Implications for Trade Theory

The gravity model predicts that trade costs affect trade volumes, with the elasticity depending on the substitutability of goods and the availability of alternative supply sources. The absence of a detectable volume effect is consistent with the possibility that the effective trade cost increase from the Canal drought was small relative to total trade costs. Anderson and van Wincoop (2004) estimate total trade costs (tariffs plus transportation plus border costs) at approximately 170 percent ad valorem. The incremental cost of rerouting via Suez or trans-Pacific (roughly 2–5 percent of cargo value, based on industry estimates of additional fuel and time costs) is a small perturbation to this total. However, the imprecision of the estimates means this interpretation cannot be confirmed with the present data—the confidence intervals are consistent with both zero effect and moderate declines.

This interpretation is consistent with the theoretical framework of Eaton and Kortum (2002) and Costinot and Rodríguez-Clare (2012), in which trade flows are determined by comparative advantage and trade costs jointly, and small perturbations to trade costs produce small perturbations to trade volumes. The Canal drought increased trade costs on the Canal route, but it did not eliminate the comparative advantage of Asian manufacturing or the demand for Asian goods in the US. An alternative interpretation—that the null reflects attenuation from measurement error in the treatment variable (country of origin as a proxy for route use) or aggregation across heterogeneous commodities—cannot be ruled out.

9.3 Climate Adaptation and Infrastructure Resilience

The Panama Canal drought was driven by climate variability (El Niño), and climate models predict increasing frequency and intensity of such events. The absence of a detectable aggregate volume effect is consistent with—though not definitive proof of—trade infrastructure resilience to temporary climate-driven disruptions, at least when the transportation network offers alternative routes.

This interpretation must be qualified in several respects. First, the design is imprecise enough that meaningful declines cannot be ruled out. Second, the Panama Canal drought occurred in a context where alternative routes existed (Suez, trans-Pacific). A simultaneous disruption to multiple chokepoints—which is not implausible under severe climate scenarios—could produce much larger effects. Third, the absence of detectable volume effects does not imply null effects on trade *costs*. Importers and shipping lines almost certainly paid higher costs during the drought (longer transit times, higher freight rates, rerouting expenses), but these costs may have been absorbed by margins rather than transmitted to trade volumes. The welfare effects of the drought may be significant even when the quantity effects are not detectable at the port-month level. Fourth, as both Roth (2022) and Freyaldenhoven et al. (2019) emphasize, event-study designs are sensitive to violations of parallel trends, and the significant pre-trend F-test ($p = 0.008$) warrants caution in interpreting the null as precisely zero rather than as imprecisely estimated.

9.4 Limitations

Several limitations qualify the interpretation of these findings. First, the joint F-test of pre-treatment event-study coefficients is significant ($p = 0.008$), indicating that the parallel trends assumption is imperfectly satisfied. While the pre-treatment fluctuations do not exhibit a systematic monotonic trend, the presence of statistically detectable differential dynamics between high- and low-Canal-share ports introduces uncertainty about the counterfactual import path. As Rambachan and Roth (2023) emphasize, pre-trend violations do not necessarily invalidate the design but do require interpreting estimates with greater caution.

Second, the treatment variable—Canal Share measured as the pre-drought share of imports from selected Asian countries—is a proxy for actual Panama Canal route dependence, not a direct measure of it. Country of origin is not equivalent to route of shipment: some East Coast imports from Asian countries may arrive via the Suez Canal or through West Coast transshipment, while some non-listed origins may be indirectly affected through shipping network equilibrium effects. This measurement error likely attenuates the treatment coefficient toward zero, meaning the null result could partly reflect treatment mismeasurement rather than genuine trade resilience. Ideally, route-level data from bills of lading or Automatic Identification System (AIS) vessel tracking would provide a sharper treatment variable, but such data are not publicly available at the scale needed for this analysis.

Third, the Census trade data measure import values at the port of entry, which may not correspond to the final destination of goods. If importers shifted their port of entry during the drought, the port-level data would capture this reallocation, but within-district shifts (from one terminal to another) would not be visible.

Fourth, the analysis focuses on import values and cannot speak to import prices or quantities separately. If the Canal drought increased import prices while leaving quantities unchanged, the value-based analysis would miss this. The welfare effects of the drought—

operating through higher freight rates, longer transit times, and increased shipping costs—may be substantial even when aggregate import values show no detectable change.

Fifth, while the sample of 186 ports is larger than expected given the number of major US gateway ports (roughly 30), many are small customs districts with sporadic trade flows. The inclusion of inland customs districts and ports with intermittent imports introduces noise that may not be well-suited to detecting a maritime route-specific shock. The precision of the estimates is driven disproportionately by the large ports, which show no treatment effects.

10. Conclusion

The 2023–24 Panama Canal drought forced a 50 percent reduction in daily transits through the world's most important commercial waterway. This paper exploits this natural experiment to estimate the effect of waterway disruptions on US port-level imports, using a continuous treatment difference-in-differences design that leverages variation in ports' pre-drought dependence on Canal routes.

The central finding is that aggregate monthly port import values show no detectable net response to the Canal disruption. The preferred estimate is near zero (-0.05 log points), and this null persists across every alternative specification, inference procedure, and placebo test. However, three important caveats temper the interpretation. First, the design is imprecise: confidence intervals are wide enough to encompass declines of nearly 40 percent. Second, the joint F-test of pre-treatment coefficients rejects parallel trends ($p = 0.008$), though the pre-period pattern is noisy rather than monotonic. Third, the treatment variable—based on country of origin rather than actual route use—likely attenuates the coefficient toward zero.

Despite these caveats, the evidence is informative. The absence of a detectable volume effect, combined with well-documented shipping line rerouting announcements and the suggestive triple-difference coefficient (-4.95 log points, $p = 0.131$), is consistent with a system that adjusted to the disruption through route substitution rather than trade contraction. The contrast with Feyrer (2021)'s findings on the 1967 Suez closure is instructive, though the episodes differ fundamentally in duration, completeness, and the available technology for adjustment.

The policy implications should be framed narrowly. Monthly aggregate port import values—a coarse outcome that cannot distinguish prices from quantities, and that aggregates across all commodities and origin routes—showed no detectable decline during a temporary, partial disruption in a context with multiple alternative routes. This does not establish that the disruption was costless: freight rates spiked, transit times increased, and shipping lines incurred substantial rerouting expenses. The welfare effects, operating through prices rather than volumes, may have been significant even though the quantity margin did not visibly respond.

An important open question is whether this pattern would hold under more extreme conditions. The Suez comparison suggests that duration matters: eight years of complete closure produced persistent trade reallocation that 14 months of partial restriction did not. Climate change may test these boundaries. The 2023 drought shows that the modern trade network can absorb a severe but temporary chokepoint disruption without visible aggregate quantity effects; the next drought will reveal how far that absorptive capacity extends.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Alessandria, George, Joseph P Kaboski, and Virgiliu Midrigan, "Inventories, Lumpy Trade, and Large Devaluations," *American Economic Review*, 2010, *100* (5), 2304–2339.

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

Barrot, Jean-Noël and Julien Sauvagnat, "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks," *Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.

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

Boehm, Christoph E, Aaron Flaaen, and Nitya Pandalai-Nayar, "Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake," *Review of Economics and Statistics*, 2019, *101* (1), 60–75.

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

_ , _ , _ , and Nicola Rosaia, "Search Frictions and Efficiency in Decentralized Transport Markets," *Quarterly Journal of Economics*, 2023, *138* (4), 2451–2503.

Burke, Marshall, Solomon M Hsiang, and Edward Miguel, "Global Non-linear Effect of Temperature on Economic Production," *Nature*, 2015, *527* (7577), 235–239.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi, "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *Quarterly Journal of Economics*, 2021, *136* (2), 1255–1321.

Cosar, A Kerem and Banu Demir, "Shipping Inside the Box: Containerization and Trade," *Journal of International Economics*, 2022, *135*, 103585.

Costinot, Arnaud and Andrés Rodríguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, 2012, *102* (1), 94–130.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken, "What Do We Learn from the Weather? The New Climate-Economy Literature," *Journal of Economic Literature*, 2014, *52* (3), 740–798.

Desmet, Klaus, Robert E Kopp, Scott A Kulp, Dávid Krisztián Nagy, Michael Oppenheimer, Esteban Rossi-Hansberg, and Benjamin H Strauss, "Evaluating the Economic Cost of Coastal Flooding," *American Economic Journal: Macroeconomics*, 2021, *13* (2), 444–486.

Disdier, Anne-Célia and Keith Head, "The Puzzling Persistence of the Distance Effect on Bilateral Trade," *Review of Economics and Statistics*, 2008, *90* (1), 37–48.

Donaldson, Dave, "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, 2018, *108* (4–5), 899–934.

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

Fajgelbaum, Pablo, Pinelopi K Goldberg, and Patrick J Kennedy, "Trade Disruptions and Firm Responses: Evidence from the Red Sea Crisis," Technical Report, NBER Working Paper 2024.

Feyrer, James, "Distance, Trade, and Income — The 1967 to 1975 Closing of the Suez Canal as a Natural Experiment," *Journal of Development Economics*, 2021, *153*, 102708.

Freyaldenhoven, Simon, Christian Hansen, and Jesse M Shapiro, "Pre-Event Trends in the Panel Event-Study Design," *American Economic Review*, 2019, *109* (9), 3307–3338.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Head, Keith and Thierry Mayer, "Gravity Equations: Workhorse, Toolkit, and Cookbook," *Handbook of International Economics*, 2014, *4*, 131–195.

Hsiang, Solomon, "Climate Econometrics," *Annual Review of Resource Economics*, 2016, *8*, 43–75.

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

Maurer, Noel and Carlos Yu, "The Role of the Panama Canal in Economic Development," *Journal of Economic History*, 2011, *71* (2), 543–560.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Redding, Stephen J and Matthew A Turner, "Transportation Costs and the Spatial Organization of Economic Activity," *Handbook of Regional and Urban Economics*, 2015, *5*, 1339–1398.

Roth, Jonathan, "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 2022, *4* (3), 305–322.

_ , Pedro H C Sant'Anna, Alyssa Bilinski, and John Poe, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

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

A. Data Appendix

A.1 Data Sources

US Census International Trade Data. Monthly port-level import values were obtained from the US Census Bureau's International Trade database via the USA Trade Online API (https://usatrade.census.gov). I extracted general import values (customs value, in US dollars) at the port-of-entry–by–country-of-origin level for all US customs districts, covering January 2019 through December 2024 (72 months). The data are accessed using the Census API with commodity code "TOTAL" to obtain all-commodity aggregates.

Panama Canal Authority (ACP) Transit Data. Monthly transit statistics were obtained from the ACP's public reports (https://pancanal.com). These data include total vessel transits by lock type (Panamax and Neo-Panamax) and vessel category (container, bulk carrier, tanker, vehicle carrier, general cargo, passenger). I use total daily transit averages to construct the drought intensity measure.

FRED Natural Gas Prices. The Henry Hub natural gas spot price (DHHNGSP) was obtained from the Federal Reserve Economic Data (FRED) database (https://fred.stlouisfed.org). Monthly averages are used as a control variable.

A.2 Sample Construction

The raw data contain port-country-month observations. I aggregate to the port-month level by summing import values across all countries of origin within each port-month cell. I also construct origin-specific aggregates for Canal-dependent Asian countries and European countries.

The Canal-dependent origin group includes: China, Japan, South Korea, Taiwan, Vietnam, Thailand, Indonesia, Philippines, Malaysia, Singapore, and India. These countries account for the vast majority of US–Asia trade that transits the Panama Canal.

The European origin group includes: Germany, United Kingdom, France, Italy, Netherlands, Belgium, Spain, Switzerland, and Sweden. These countries trade with the US primarily via trans-Atlantic routes that do not involve the Panama Canal.

Sample restrictions: I retain all US customs districts with positive total imports in at least one month during the sample period. No minimum size threshold is imposed, but the econometric specifications include port fixed effects that absorb time-invariant differences in port size. The final sample contains 186 ports and 13,160 port-month observations.

A.3 Variable Definitions

- *Log total imports*: $\log(\text{Total imports}_{pt} + 1)$, where total imports are in US dollars. The $+1$ adjustment handles zero-valued observations.

- *Canal share*: For East/Gulf Coast ports: $\frac{\overline{\text{Canal-origin imports}_{p,2019-22}}}{\text{Total imports}_{p,2019-22}}$, computed over the baseline period (2019, 2021–22, excluding 2020). Set to zero for West Coast ports, whose Asian imports arrive via direct trans-Pacific routes.

- *Drought intensity*: $1 - \frac{\text{Actual daily transits}_t}{\text{Normal daily transits}}$, where normal is the 2019/2021–22 average.

- *Treatment (continuous)*: Canal share $\times$ drought intensity.

- *Treatment (binary)*: $\mathbb{I}[\text{Canal share}_p > \text{median}] \times \mathbb{I}[t \geq \text{July 2023}]$.

- *High Canal*: Indicator for above-median Canal share.

- *Post drought*: Indicator for July 2023 and later.

A.4 Drought Intensity Construction

Figure 5 plots the drought intensity measure over the sample period. The measure is zero during normal operations (2019–mid-2023) and rises to a maximum of 0.51 in February 2024, when daily transits fell to approximately 18 from a normal of 36–37.

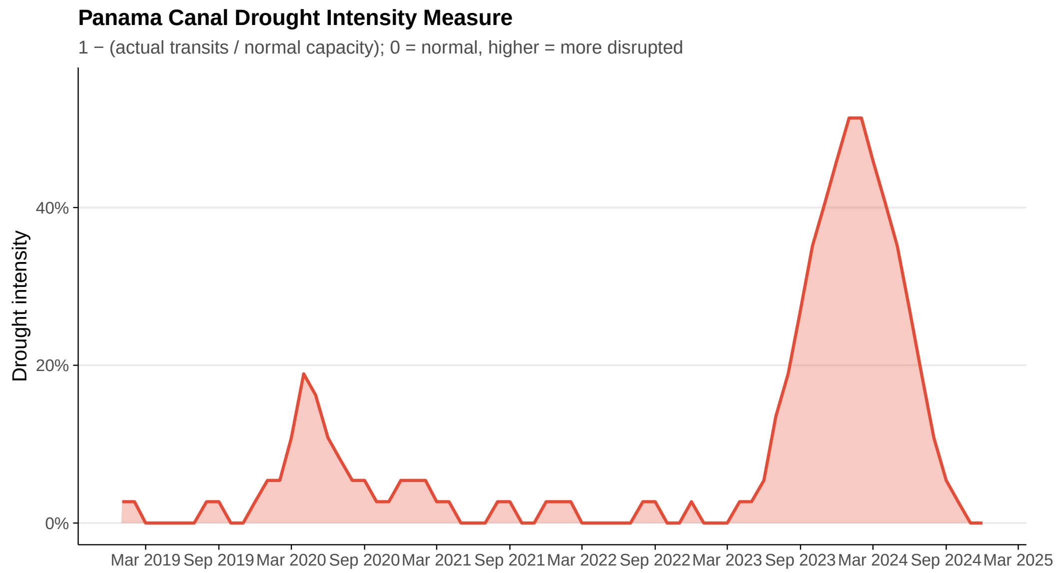

Figure 5: Drought Intensity Measure, 2019–2024

Notes: Drought intensity is defined as $1 - (\text{actual daily transits}/\text{normal daily transits})$. Normal daily transits are the average during 2019 and 2021–22. The measure equals zero during normal operations and peaks at 0.51 in February 2024.

B. Identification Appendix

B.1 Pre-Trends Assessment

The event study in Figure 1 reveals some individually significant pre-treatment coefficients, with a maximum absolute value of approximately 4.49. A joint F-test of all 23 pre-treatment coefficients rejects the null hypothesis that they are jointly zero ($F(23, 12{,}865) = 1.86$, $p = 0.008$). This section discusses the implications for identification.

The significant F-test means the parallel trends assumption is not fully satisfied. As Roth (2022) argues, pre-trend tests have limited power, so failing to reject does not establish parallel trends; conversely, rejecting should induce caution about causal interpretation. Several mitigating considerations, however, suggest the pre-trend violations are unlikely to substantially bias the main conclusion.

First, the pre-treatment coefficients do not exhibit a monotonic trend—they fluctuate irregularly around zero. A directional pre-trend (systematic positive or negative drift) would be more concerning than the irregular pattern observed, because extrapolation of a directional trend into the treatment period would produce a specific direction of bias. The observed pattern is more consistent with heterogeneous noise across port-exposure cells, possibly reflecting post-COVID recovery differences or compositional volatility at ports with varied

Canal dependence.

Second, the main result is robust to the inclusion of port-specific linear time trends (Column 3 of Table 2), which absorb differential linear trends across ports. The coefficient changes from -0.05 to -0.75 but remains statistically insignificant, suggesting that linear trends do not substantially alter the conclusion.

Third, the timing placebo test (July 2021–August 2022) produces an insignificant coefficient (0.13, $p = 0.729$), confirming that the specification does not spuriously detect treatment effects in periods without Canal disruption.

Nonetheless, the significant pre-trend F-test is a substantive limitation. Following Rambachan and Roth (2023), the estimates should be interpreted with the understanding that pre-period instability introduces additional uncertainty beyond what the standard errors capture. The main result—that no detectable net effect is found—should not be equated with definitive evidence that the effect is precisely zero.

B.2 Canal Share Distribution

Figure 6 plots the distribution of pre-drought Canal shares across ports. The distribution is right-skewed, with most ports having Canal shares below 0.20 and a long tail of ports with Canal shares exceeding 0.50. The median Canal share is approximately 0.06, and the 75th percentile is approximately 0.18.

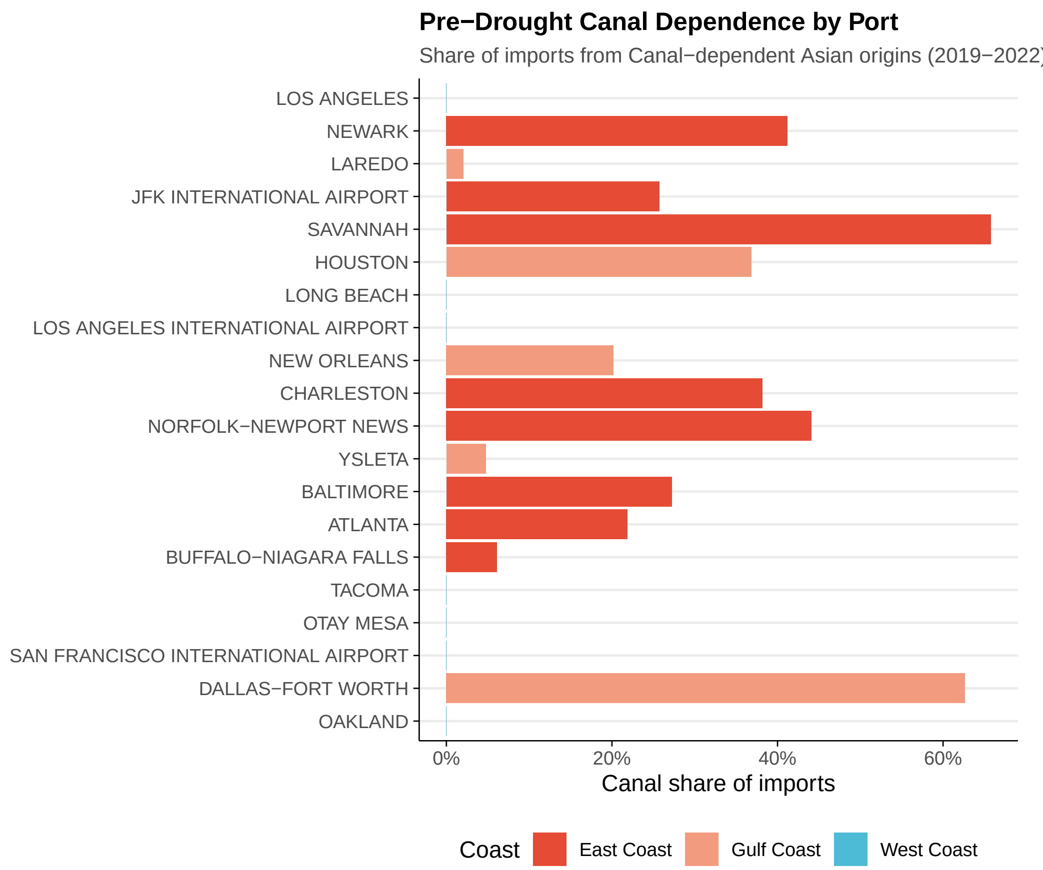

Figure 6: Distribution of Pre-Drought Canal Share Across Ports

Notes: Histogram of Canal share (pre-drought Asian import share) across 186 US customs districts. Canal share is the ratio of imports from Canal-dependent Asian countries to total imports, averaged over the baseline period (2019, 2021–22).

C. Robustness Appendix

C.1 Alternative Specifications

Table 6 presents five alternative specifications testing the robustness of the main result. The baseline coefficient of -0.05 is stable across import levels (Column 2), port-by-calendar-month fixed effects (Column 3), exclusion of 2020 (Column 4), and binary Canal share classification (Column 5). Additionally, winsorizing the dependent variable at the 1st and 99th percentiles—to address the concern that the medium-port anomaly in Table 5 reflects extreme outliers masking a true effect—produces a nearly identical coefficient (-0.049, $p = 0.988$), confirming that outliers are not driving the null.

Table 6: Robustness: Alternative Specifications

	(1) Baseline	(2) Levels ($M)	(3) Port×Month FE	(4) Excl. COVID	(5) Binary
Treatment	-0.0494	145.32	-0.4640	0.0175	0.2419
	(3.1550)	(209.73)	(3.1037)	(3.2343)	(0.5455)
Observations	13,160	13,160	13,160	10,966	13,160
R^2	0.778	0.968	0.818	0.780	0.778
Port FE	Yes	Yes		Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes	Yes
Port × calendar-month FE			Yes		

Notes: Col. 1 is the baseline from Table 2. Col. 2 uses levels ($M). Col. 3 adds port×calendar-month FE. Col. 4 excludes 2020. Col. 5 uses binary treatment (above-median Canal share × drought indicator). SEs clustered at port level. * p<0.10, ** p<0.05, *** p<0.01.

C.2 Robustness Summary

Table 7 summarizes the alternative inference and placebo test results. Both wild cluster bootstrap and randomization inference confirm the null, and both placebos pass.

Table 7: Robustness Summary

Test	Coefficient	*p*-value	Result
Inference			
Wild cluster bootstrap	-0.0494	0.9920	Not significant
Randomization inference	-0.0494	0.9710	Not significant
Placebos			
Timing placebo (2021–2022)	0.1306	0.7291	Null (passes)
European imports placebo	2.5940	0.3394	Null (passes)

Notes: Wild cluster bootstrap uses Webb weights with 999 replications. Randomization inference permutes Canal share across ports 999 times. Timing placebo applies the same specification to a fake drought period (July 2021–August 2022) using only pre-actual-drought data. European imports placebo tests whether non-Canal-dependent imports respond to the Canal share treatment.

C.3 Leave-One-Out Analysis

Figure 7 presents leave-one-out estimates. Each point represents the estimated treatment coefficient when one port is dropped from the sample. The estimates are tightly clustered around zero, and no single port drives the result.

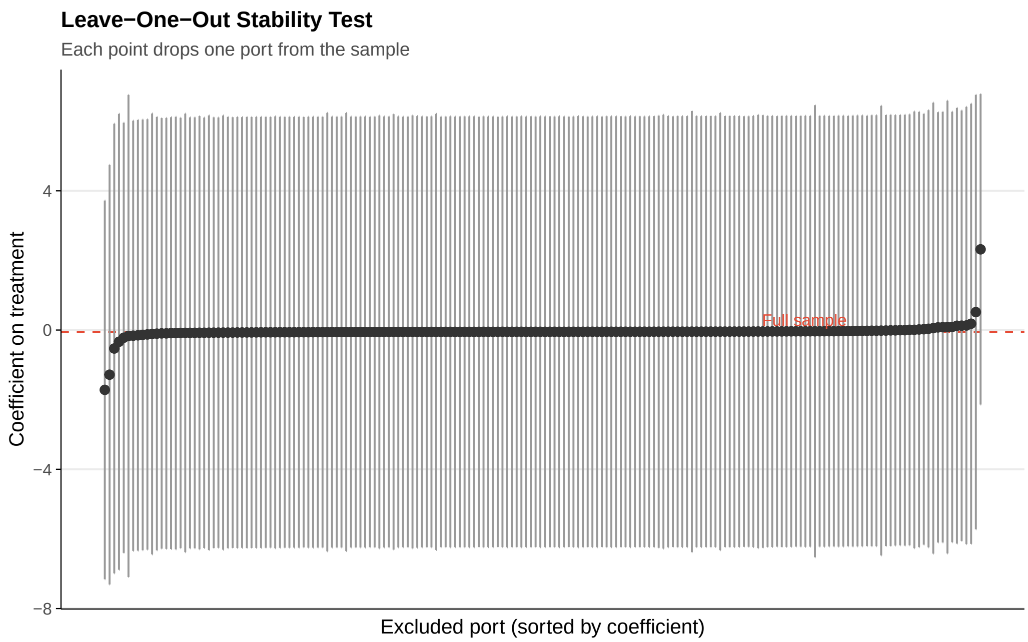

Figure 7: Leave-One-Out Estimates of Treatment Effect

Notes: Each point represents the estimated continuous treatment coefficient from the main specification (Equation (3)) with one port dropped. The horizontal line indicates the full-sample estimate.

C.4 Randomization Inference Distribution

Figure 8 plots the distribution of placebo treatment effects from the randomization inference exercise. The actual treatment coefficient (vertical line) falls well within the central mass of the distribution, confirming that the observed effect is indistinguishable from random assignment of Canal shares.

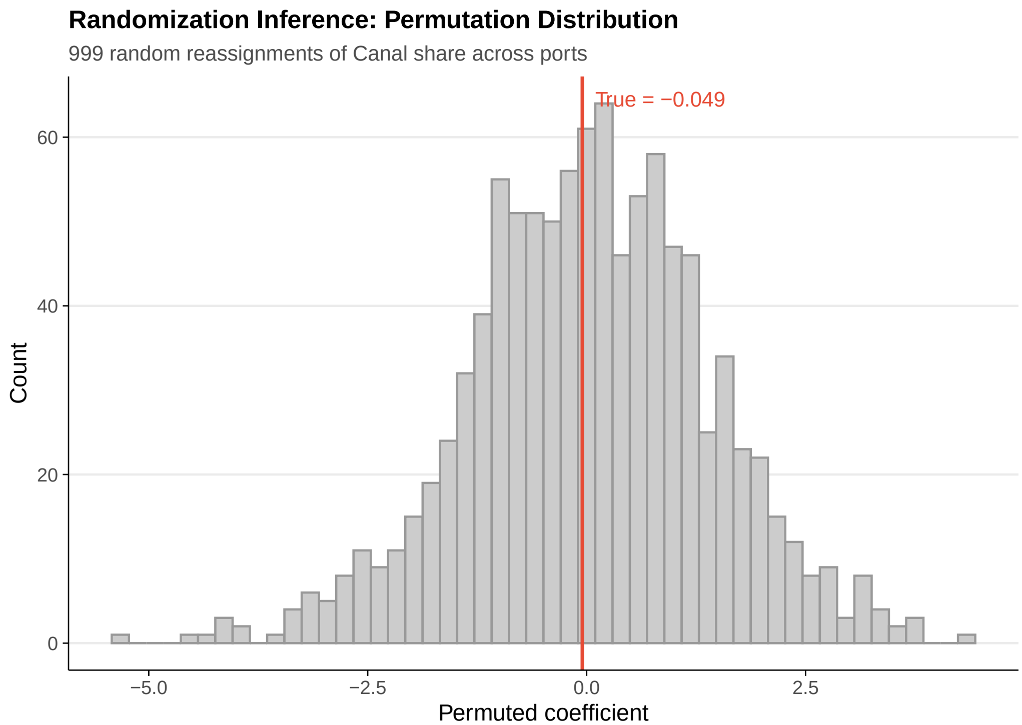

Figure 8: Randomization Inference: Distribution of Placebo Treatment Effects

Notes: Distribution of treatment coefficients from 999 permutations of Canal share across ports. The vertical line indicates the actual treatment coefficient (-0.05). The two-sided RI p-value is 0.971.

D. Heterogeneity Appendix

The main text presents heterogeneity by port size tercile (Table 5). The primary finding is that large ports—which handle the vast majority of US imports by value—show essentially zero treatment effects. Small ports show a negative but insignificant effect. The medium-port tercile shows a positive and significant coefficient that is an order of magnitude larger than estimates from other specifications—this anomalous result likely reflects compositional idiosyncrasies in this heterogeneous group (which includes specialized commodity ports, military installations, and regional distribution centers) rather than a genuine Canal exposure effect.

E. Additional Figures and Tables

All figures and tables referenced in the main text and appendices have been presented above. Supplementary exhibits, including detailed regression output and additional robustness checks, are available in the replication code.

F. Standardized Effect Sizes

Table 8: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
Log total imports	Table 2, Col. 2	-0.0494	0.0354	4.7231	-0.0004	0.0236	Null
Log Canal-origin imports	Table 2, Col. 5	-2.3595	0.0354	7.5106	-0.0111	0.0142	Small negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For continuous treatments, SDE = $\hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$, which gives the effect of a one-standard-deviation change in the treatment variable, measured in standard deviations of the outcome. SD(Y) and SD(X) are unconditional standard deviations from Table 1.

Research question: How does Panama Canal drought-induced transit capacity reduction affect US port-level merchandise imports, exploiting differential Canal dependence across ports? **Treatment:** Continuous interaction of pre-drought Canal share and monthly drought intensity. **Data:** US Census International Trade API, January 2019–December 2024, 186 ports, 13,160 port-months. **Method:** Continuous treatment DiD with port and year-month fixed effects, port-clustered SEs. **Sample:** US customs districts with positive imports in at least one month during the sample period.

Classification thresholds (7 categories): large negative (< -0.15), moderate negative (-0.15 to -0.05), small negative (-0.05 to -0.005), null (-0.005 to 0.005), small positive (0.005 to 0.05), moderate positive (0.05 to 0.15), large positive (> 0.15). Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. "Null" denotes a near-zero effect size ($|\text{SDE}| < 0.005$), not a failure to reject a null hypothesis.