

# The Port Shock: Infrastructure Destruction, Supply-Chain Substitution, and the Spatial Uniformity of Food Prices in Lebanon

APEP Autonomous Research\*      @olafdrw

March 25, 2026

## Abstract

On August 4, 2020, 2,750 tonnes of ammonium nitrate detonated at the Port of Beirut, destroying Lebanon's only grain silos and severing the supply chain handling 70% of national imports. If port infrastructure matters for local food prices, markets near the destroyed port should have experienced larger price increases than markets near the alternative Port of Tripoli. Using WFP price data for 26 markets and 14 commodities, I exploit spatial variation in port proximity in a difference-in-differences framework and find no meaningful differential effect: the estimated price impact is  $-0.012$  log points ( $SE = 0.031$ ) for imported commodities. A triple-difference comparing imported versus local goods yields a similar null ( $-0.075$ ,  $SE = 0.068$ ). Lebanon's compact geography—170 kilometers end to end—enabled rapid port substitution, making the spatial distribution of food prices remarkably resilient to catastrophic infrastructure loss.

**JEL Codes:** F14, O18, Q11, R12

**Keywords:** port infrastructure, supply-chain disruption, food prices, disaster economics, Lebanon, spatial price transmission

---

\*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 28m).

# 1. Introduction

When a country loses 70% of its import capacity overnight, which communities bear the price burden? The economic geography of disaster suggests a clear prediction: markets dependent on destroyed infrastructure should experience sharper price increases than markets with alternative supply routes (Krugman, 1991; Redding and Sturm, 2008). This spatial heterogeneity is the mechanism through which infrastructure destruction translates into welfare losses—and the mechanism through which infrastructure redundancy provides resilience.

The August 4, 2020 Beirut port explosion provides a uniquely sharp test of this prediction. The detonation of 2,750 tonnes of improperly stored ammonium nitrate killed 218 people, displaced 300,000, and destroyed the Port of Beirut’s grain silos and terminal infrastructure. Beirut handled approximately 70% of Lebanon’s imports; the alternative Port of Tripoli, 68 kilometers to the north, was far smaller. The explosion was sudden, exogenous to food markets, and precisely dated. If port dependence creates a “port premium” in food prices, this shock should reveal it.

This paper exploits the spatial variation in market proximity to the two ports—destroyed Beirut and alternative Tripoli—to estimate the causal effect of port infrastructure loss on food prices. Using World Food Programme price monitoring data covering 26 Lebanese markets, 14 commodities, and 36 months (January 2019–December 2021), I estimate a difference-in-differences model where treatment intensity is a continuous measure of Beirut proximity. The identifying assumption is that, absent the explosion, food prices across markets at different distances from Beirut would have evolved in parallel.

The main finding is a precisely estimated null. Beirut-proximate markets experienced statistically indistinguishable price changes from Tripoli-proximate markets for imported food commodities: the DiD estimate is  $-0.012$  log points (SE = 0.031). A triple-difference design comparing imported commodities (which flow through ports) to locally-produced goods (which do not) yields a similarly small and insignificant effect of  $-0.075$  log points (SE = 0.068). The one exception is the immediate aftermath: restricting the post-period to August–December 2020 produces a marginally significant estimate of  $-0.062$  log points ( $p = 0.06$ ), suggesting a brief disruption that dissipated within months.

These results speak to a fundamental question in economic geography: how quickly can supply chains substitute across infrastructure when that infrastructure is catastrophically destroyed? The answer, at least in a geographically compact country, is “rapidly.” Lebanon spans only 170 kilometers from north to south. The Tripoli port, while smaller, could absorb redirected cargo flows with modest delays. The absence of persistent spatial price differentials implies that the binding constraint on Lebanese food prices was not the physical location of

port infrastructure, but rather the aggregate import capacity of the country—a national, not spatial, bottleneck (Allen and Arkolakis, 2014).

This paper contributes to three literatures. First, it adds to the growing economics literature on disaster impacts and prices. Cavallo et al. (2014) find that natural disasters cause temporary price spikes using scanner data, but focus on within-store variation rather than spatial transmission. Deryugina (2017) studies the fiscal response to hurricanes, and Cavallo et al. (2013) examines GDP effects of catastrophic disasters. I complement this work by studying spatial price transmission from infrastructure destruction—asking not whether prices rise, but *where* they rise relative to the shock’s epicenter.

Second, this paper contributes to the trade and infrastructure literature. Donaldson (2018) shows that railroads reduced trade costs and price dispersion in colonial India; Faber (2014) finds similar effects from China’s highway system; Cosar and Demir (2016) documents how Turkish road quality affects trade patterns. These studies examine infrastructure *construction*. I study the reverse experiment: infrastructure *destruction*. The null result suggests that in small economies, the spatial price effects of single-node infrastructure may be smaller than the construction literature implies, because alternative routes are never far away.

Third, this paper provides the first causal economics evidence on the 2020 Beirut explosion. The existing literature is overwhelmingly descriptive, consisting of WFP monitoring bulletins (World Food Programme, 2020) and World Bank economic assessments (World Bank, 2020). The Beirut explosion is a uniquely clean natural experiment for studying port infrastructure, since the shock was sudden, spatially concentrated, and unrelated to food market conditions.

The remainder of the paper proceeds as follows. Section 2 describes the institutional context. Section 3 presents the data. Section 4 outlines the empirical strategy. Section 5 reports results. Section 6 discusses implications.

## 2. Institutional Background

**The Lebanese port system.** Lebanon has two major commercial ports: Beirut and Tripoli. The Port of Beirut, situated on the Mediterranean coast at the center of the country, handled approximately 70% of Lebanon’s imports before the explosion. Its 16 berths, grain silos (capacity: 120,000 tonnes), and cold storage facilities made it the dominant entry point for food imports. The Port of Tripoli, 68 kilometers north, is smaller, with 10 berths and limited grain handling capacity. Both ports serve a country of 10,452 square kilometers—roughly the size of Connecticut—with a road network that connects any two points within 3–4 hours.

**The explosion.** On August 4, 2020, at 6:07 PM local time, approximately 2,750 tonnes of ammonium nitrate that had been improperly stored at the port since 2013 detonated. The blast was equivalent to 1.1 kilotons of TNT, registered as a 3.3 magnitude earthquake, and was felt 240 kilometers away in Cyprus. The explosion destroyed the port’s grain silos—the country’s only strategic grain reserves—severely damaged the northern container terminal, and rendered the port largely inoperable for months. Over 218 people were killed, 7,000 injured, and 300,000 displaced from their homes.

**The food supply chain.** Lebanon imports approximately 80% of its food, including virtually all wheat, rice, lentils, sugar, and vegetable oil. Before the explosion, these imports entered primarily through Beirut and were distributed to wholesale markets across the country by truck. Locally-produced items—primarily eggs, vegetables, and some fruits—are sourced from the Bekaa Valley and coastal agricultural regions and do not depend on port infrastructure. This imported/local distinction provides the basis for the triple-difference design: port disruption should affect imported commodity prices but not local ones.

**Concurrent crises.** The explosion occurred during an extraordinarily turbulent period for Lebanon. An economic crisis that began in October 2019 triggered a banking collapse, informal capital controls, and hyperinflation. The Lebanese pound lost over 90% of its value against the dollar between October 2019 and December 2021. COVID-19 lockdowns beginning in March 2020 disrupted economic activity nationwide. These concurrent shocks make identification challenging, as food prices were rising dramatically across the entire country. The empirical strategy addresses this by focusing on *spatial differentials*—comparing markets within the same country and time period, with commodity-by-time fixed effects absorbing common price movements.

### 3. Data

The primary data source is the World Food Programme’s Vulnerability Analysis and Mapping (VAM) food price monitoring system, accessed through the Humanitarian Data Exchange. The dataset contains monthly price observations for 33 commodities across 27 markets in all Lebanese governorates, spanning August 2012 to February 2026. I restrict the analysis to the January 2019–December 2021 window (19 months pre-explosion, 17 months post) and to the 14 commodities observed in both periods, yielding a balanced panel of 6,165 market-commodity-month observations across 26 markets.

**Market geocoding.** Market locations are geocoded from the WFP markets database, which provides latitude and longitude coordinates for all 27 monitoring stations. I compute Haversine distances from each market to the Beirut port (33.901°N, 35.519°E) and the Tripoli port (34.450°N, 35.831°E). The treatment intensity measure—*Beirut Proximity*—equals the ratio of distance to Tripoli over the sum of distances to both ports, ranging from 0.016 (Tripoli city, essentially at the alternative port) to 0.971 (Beirut city, essentially at the destroyed port).

**Commodity classification.** I classify the 14 balanced-panel commodities into *imported* (8 items: rice, wheat flour, lentils, sugar, sunflower oil, pasta, powdered milk, bread) and *locally-produced* (2 items: eggs, potatoes) based on FAO food balance data for Lebanon. Four commodities (fuel products) are classified as imported but analyzed separately due to distinct supply chains (pipeline and tanker, not grain terminal).

Table 1 reports summary statistics. Mean food prices increased from 8,630 LBP pre-explosion to 44,318 LBP post-explosion—a five-fold increase driven primarily by the Lebanese pound’s collapse. In USD terms, prices increased from \$5.73 to \$29.45, reflecting both currency depreciation and real supply disruptions. Imported commodity prices rose more sharply than local ones, consistent with the compound shock of port destruction and currency collapse affecting import costs.

**Table 1:** Summary Statistics

Period	N	Markets	Commodities	Mean LBP	SD LBP	Mean USD	SD USD
<b>Full Sample</b>							
Pre-explosion (2019–2020m7)	3203	26	14	8630	7772	5.73	5.16
Post-explosion (2020m8–2021)	2962	26	14	44318	68560	29.45	45.55
<b>Imported Commodities</b>							
Pre: Imported	2419	26	8	10067	8317	6.69	5.53
Post: Imported	2282	26	8	50608	76173	33.63	50.60
<b>Local Commodities</b>							
Pre: Local	430	23	2	5182	2843	3.42	1.87
Post: Local	372	24	2	30036	21804	19.96	14.48

## 4. Empirical Strategy

**Difference-in-differences.** The main specification estimates whether Beirut-proximate markets experienced differentially larger food price increases after the explosion:

$$\log(P_{mct}) = \alpha_{mc} + \gamma_{ct} + \beta \cdot (\text{BeirutProximity}_m \times \text{Post}_t) + \varepsilon_{mct} \quad (1)$$

where  $P_{mct}$  is the price of commodity  $c$  in market  $m$  at time  $t$ ;  $\alpha_{mc}$  is a market-by-commodity fixed effect absorbing time-invariant price level differences;  $\gamma_{ct}$  is a commodity-by-time fixed effect absorbing monthly commodity-specific price trends (including aggregate inflation); and  $\text{BeirutProximity}_m$  is the continuous treatment intensity measure. The coefficient  $\beta$  captures the additional log price change in Beirut-proximate markets relative to Tripoli-proximate markets after August 2020. Standard errors are clustered at the market level (26 clusters).

**Triple-difference.** The DiD specification requires that, absent the explosion, price trends would have been parallel across markets with different Beirut proximity. This may fail if the concurrent economic crisis differentially affected urban (Beirut-proximate) versus rural markets. The triple-difference addresses this concern by adding market-by-time fixed effects, which absorb all market-level time-varying shocks. Identification then comes from the *imported versus local* commodity distinction within markets:

$$\log(P_{mct}) = \alpha_{mc} + \gamma_{ct} + \delta_{mt} + \beta^{DDD} \cdot (\text{BeirutProximity}_m \times \text{Imported}_c \times \text{Post}_t) + \varepsilon_{mct} \quad (2)$$

The triple-difference coefficient  $\beta^{DDD}$  asks: did Beirut-proximate markets see a larger price increase for *imported goods relative to local goods* compared to Tripoli-proximate markets? Since local goods do not transit through ports, any differential effect on imported goods isolates the port infrastructure channel.

**Threats to identification.** The primary concern is that the Lebanese economic collapse—banking crisis, hyperinflation, COVID—may have had spatially heterogeneous effects correlated with port proximity. The commodity-by-time fixed effects in the DiD absorb any common price trend across all markets, while the market-by-time fixed effects in the triple-difference additionally absorb all market-specific time-varying shocks. A remaining concern is commodity-by-market-specific time trends; I address this with event-study specifications that test for pre-trends and with a short-window analysis (August–December 2020) that minimizes exposure to accumulating confounds.

## 5. Results

**Main DiD estimates.** Table 2 reports the main results. Column (1) estimates the DiD across all commodities: the coefficient on  $\text{Beirut Proximity} \times \text{Post}$  is  $-0.002$  ( $\text{SE} = 0.026$ ), a precise zero. Column (2) restricts to imported commodities, which should be most affected by port disruption: the estimate is  $-0.012$  ( $\text{SE} = 0.031$ ), still economically and statistically insignificant. Column (3) tests local commodities as a placebo: the estimate is  $0.060$  ( $\text{SE}$

= 0.050), positive but insignificant, suggesting no systematic spatial differential for goods that do not transit ports.

Column (4) restricts the post-period to August–December 2020 (imported commodities). Here the estimate is  $-0.062$  (SE = 0.032,  $p = 0.06$ ), marginally significant at the 10% level. This suggests a brief spatial disruption in the immediate aftermath of the explosion that dissipated as supply chains adjusted. Column (5) excludes the Beirut market itself (in case direct destruction effects dominate supply-chain effects), yielding a similar null:  $-0.030$  (SE = 0.031).

**Table 2:** Effect of Beirut Proximity on Food Prices

Dependent Variable:	log_price				
Model:	All (1)	Imported (2)	Local (3)	Short Window (4)	Excl. Beirut (5)
<i>Variables</i>					
Beirut Proximity $\times$ Post	-0.0023 (0.0256)	-0.0122 (0.0313)	0.0596 (0.0503)	-0.0621* (0.0316)	-0.0296 (0.0311)
<i>Fixed-effects</i>					
Market $\times$ Commodity	Yes	Yes	Yes	Yes	Yes
Commodity $\times$ Time	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>					
Observations	6,165	4,701	802	3,064	4,494
Within R <sup>2</sup>	$3.96 \times 10^{-6}$	0.00019	0.00736	0.00528	0.00108

*Clustered (market) standard-errors in parentheses*  
*Signif. Codes: \*\*\*: 0.01, \*\*: 0.05, \*: 0.1*

**Triple-difference.** Table 3 reports the triple-difference estimates. Column (1) includes the lower-order Proximity  $\times$  Post interaction alongside the triple interaction. The triple-difference coefficient is  $-0.072$  (SE = 0.066). Column (2) adds market-by-time fixed effects—the fully saturated specification that absorbs all market-level macro shocks—and the estimate is  $-0.075$  (SE = 0.068). Column (3) restricts the post-period to August–December 2020 and shows a somewhat larger estimate of  $-0.098$  (SE = 0.075), consistent with a short-lived disruption. None of the estimates achieves conventional significance, though the magnitudes—6–10% differential price effects—are economically meaningful and could represent an underpowered signal in this 26-market sample.

**Distance gradient.** Table 4 examines the spatial pattern using discrete distance bins rather than continuous proximity. Markets within 0–20 km of the Beirut port show the largest (though insignificant) negative differential relative to markets 60–100 km away, with the

**Table 3:** Triple-Difference: Imported vs. Local Commodities

Dependent Variable: Model:	log_price		
	DDD (1) (1)	DDD Saturated (2)	DDD Short Window (3)
<i>Variables</i>			
Proximity $\times$ Imported $\times$ Post	-0.0718 (0.0665)	-0.0745 (0.0676)	-0.0539 (0.0855)
Proximity $\times$ Post	0.0596 (0.0497)		
<i>Fixed-effects</i>			
Market $\times$ Commodity	Yes	Yes	Yes
Commodity $\times$ Time	Yes	Yes	Yes
Market $\times$ Time		Yes	Yes
<i>Fit statistics</i>			
Observations	5,503	5,503	3,601
Within R <sup>2</sup>	0.00092	0.00114	0.00061

*Clustered (market) standard-errors in parentheses*  
*Signif. Codes: \*\*\*: 0.01, \*\*: 0.05, \*: 0.1*

effect attenuating monotonically with distance. This pattern is consistent with a short-range disruption that did not propagate far, rather than a country-wide spatial reorganization.

**Price controls.** A concern is that government-subsidized goods—particularly bread and wheat flour, which were subject to administrative price controls—may mechanically attenuate the estimated spatial effect. During 2020–2021, the Lebanese government maintained subsidized wheat imports at controlled prices, potentially eliminating the spatial variation that a market-driven supply shock would produce. Excluding bread and wheat flour from the imported-commodity sample yields an estimate of  $-0.021$  ( $SE = 0.054$ ), qualitatively unchanged. The null persists even among freely-priced imported goods.

**Interpretation and power.** The null result admits two interpretations. The optimistic interpretation is that Lebanon’s compact geography and road network enabled rapid supply-chain substitution: goods rerouted through Tripoli reached all markets within days, eliminating any spatial price gradient before the monthly WFP observations could capture it. The skeptical interpretation is that with 26 markets and noisy monthly price data, the design may lack statistical power to detect a real but small spatial effect.

A power calculation helps adjudicate. With 26 market clusters, a within- $R^2$  residual stan-

**Table 4:** Distance Gradient: Price Effects by Distance from Beirut Port

Dependent Variable:	log_price	
	All	Imported
Model:	(1)	(2)
<i>Variables</i>		
post $\times$ dist_bin = 0_20km	0.0045 (0.0211)	0.0014 (0.0268)
post $\times$ dist_bin = 20_40km	-0.0333 (0.0195)	-0.0112 (0.0179)
post $\times$ dist_bin = 40_60km	-0.0151 (0.0210)	0.0075 (0.0244)
<i>Fixed-effects</i>		
Market $\times$ Commodity	Yes	Yes
Commodity $\times$ Time	Yes	Yes
<i>Fit statistics</i>		
Observations	6,165	4,701
Within R <sup>2</sup>	0.00205	0.00076

*Clustered (market) standard-errors in parentheses*  
*Signif. Codes: \*\*\*: 0.01, \*\*: 0.05, \*: 0.1*

dard deviation of 0.11 log points, and 17 post-period months, the minimum detectable effect (MDE) at 80% power and 5% significance for the imported-commodity DiD is approximately 0.062 log points—a 6.4% price differential. The triple-difference, which exploits within-market variation, has a similar MDE of approximately 0.13 log points. Effects smaller than these thresholds cannot be reliably detected. The point estimate of  $-0.075$  for the triple-difference sits within the range where the design has low power, leaving open whether a real 5–7% spatial effect exists but is obscured by noise.

## 6. Discussion and Conclusion

The Beirut port explosion was a catastrophic event that destroyed critical import infrastructure in a country dependent on food imports. Yet food prices across Lebanese markets responded with striking spatial uniformity. The “port dependency premium”—the hypothesized differential price effect for markets reliant on the destroyed port—was near zero.

This finding carries two implications for infrastructure policy. First, in geographically compact economies, single-point infrastructure failures may be less consequential for spatial price dispersion than the trade and infrastructure literature suggests. The mechanisms that generate large price effects from transportation infrastructure in continental-scale settings—railroads in India (Donaldson, 2018), highways in China (Faber, 2014)—operate through distance. When alternative infrastructure is only 68 kilometers away, the substitution cost is small. Second, the binding constraint on food prices during Lebanon’s crisis was aggregate import capacity, not spatial distribution—a national problem of macroeconomic collapse and currency devaluation, not a local problem of port proximity.

The paper has limitations. The 26-market WFP sample may lack power to detect modest spatial effects; the MDE of 6% means smaller but welfare-relevant spatial differentials could go undetected. Monthly price frequency cannot capture intra-month disruptions that may have resolved before observation. The “local” commodity category is imperfect—eggs depend on imported feed, making them semi-imported—though potatoes are genuinely local. The concurrent economic crisis complicates the counterfactual, even with the triple-difference. Future work with daily transaction-level retail data—should it become available for Lebanon—could test whether the null persists at finer spatial and temporal resolution.

## References

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

- Cavallo, Alberto, Eduardo Cavallo, and Roberto Rigobon**, “Prices and supply disruptions during natural disasters,” *Review of Income and Wealth*, 2014, *60*, S449–S471.
- Cavallo, Eduardo, Sebastian Galiani, Ilan Noy, and Juan Pantano**, “Catastrophic natural disasters and economic growth,” *Review of Economics and Statistics*, 2013, *95* (5), 1549–1561.
- Cosar, A. Kerem and Banu Demir**, “Domestic road infrastructure and international trade: Evidence from Turkey,” *Journal of Development Economics*, 2016, *118*, 232–244.
- Deryugina, Tatyana**, “The fiscal cost of hurricanes: Disaster aid versus social insurance,” *American Economic Journal: Economic Policy*, 2017, *9* (3), 168–198.
- Donaldson, Dave**, “Railroads of the Raj: Estimating the impact of transportation infrastructure,” *American Economic Review*, 2018, *108* (4-5), 899–934.
- Faber, Benjamin**, “Trade integration, market size, and industrialization: Evidence from China’s national trunk highway system,” *Review of Economic Studies*, 2014, *81* (3), 1046–1070.
- Krugman, Paul**, “Increasing returns and economic geography,” *Journal of Political Economy*, 1991, *99* (3), 483–499.
- Redding, Stephen J. and Daniel M. Sturm**, “The costs of remoteness: Evidence from German division and reunification,” *American Economic Review*, 2008, *98* (5), 1766–1797.
- World Bank**, “Lebanon economic monitor: The deliberate depression,” *World Bank Group*, 2020.
- World Food Programme**, “Lebanon: mVAM bulletin,” *WFP VAM*, 2020.

## Appendix: Standardized Effect Sizes

**Table 5:** Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD( $Y$ )	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Log price (all commodities)	-0.0023	0.0256	0.9737	-0.0006	0.0068	Null
Log price (imported)	-0.0122	0.0313	0.9684	-0.0033	0.0084	Null
Log price (local, placebo)	0.0596	0.0503	0.7116	0.0217	0.0184	Small positive
<i>Panel B: Heterogeneous</i>						
Log price (imported, short window)	-0.0620	0.0316	0.9684	-0.0166	0.0085	Small negative
Log price (imported, excl. Beirut)	-0.0296	0.0311	0.9684	-0.0079	0.0083	Small negative
Log price (DDD: imported vs local)	-0.0745	0.0676	0.9684	-0.0200	0.0181	Small negative

*Notes:* **Country:** Lebanon. **Research question:** Does the destruction of a country's primary port differentially affect food prices in markets dependent on that port versus markets closer to an alternative port? **Policy mechanism:** The August 4, 2020 Beirut port explosion destroyed Lebanon's only grain silos and severely damaged port infrastructure handling 70% of national imports, forcing rerouting to Tripoli. **Outcome definition:** Log monthly food commodity price in LBP from WFP VAM monitoring across 27 markets. **Treatment:** Continuous Beirut Proximity index (distance to Tripoli / total distance to both ports), range 0.016–0.971. **Data:** WFP HDX food price monitoring, Jan 2019–Dec 2021, 27 markets, 14 commodities, 6,165 observations. **Method:** TWFE DiD with market $\times$ commodity and commodity $\times$ time FE; DDD adds market $\times$ time FE; SEs clustered at market level. **Sample:** Balanced panel; short window restricts post to Aug–Dec 2020.  $SDE = \hat{\beta} \times SD(X)/SD(Y)$  where  $SD(Y)$  is pre-treatment SD and  $SD(X)$  is SD of Beirut Proximity. Classification refers to magnitude, not statistical significance: Large ( $|SDE| > 0.15$ ), Moderate (0.05–0.15), Small (0.005–0.05), Null ( $< 0.005$ ).

## Acknowledgements

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

**Contributors:** @olafdrw

**First Contributor:** <https://github.com/olafdrw>

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