

The Blunt Weapon: Russia's Wine Embargo and Subnational Economic Activity in Moldova

APEP Autonomous Research* @SocialCatalystLab

Abstract

Russia banned Moldovan wine imports in September 2013, collapsing bilateral wine trade by 75%. I exploit this shock using a Bartik shift-share design interacting pre-determined vineyard intensity across 37 districts with embargo timing, measuring local economic activity through monthly satellite nighttime lights. I find no significant decline in nightlight radiance in wine-dependent districts. Point estimates are small and statistically indistinguishable from zero across specifications including pairs cluster bootstrap and randomization inference. Pre-existing differential trends, confirmed by a failed placebo test, limit causal interpretation; the design rules out large negative effects but not moderate ones. The findings are consistent with rapid trade diversion toward EU markets, though the null could also reflect nightlights' insensitivity to agricultural shocks.

JEL: F14, F51, O18, Q17

Keywords: trade embargo, trade weaponization, nighttime lights, Moldova, wine exports, Russia, shift-share

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 27m).

1. Introduction

Russia embargoed Moldovan wine in September 2013, and \$30 million in annual export revenue vanished. For Cahul, a district where one in nine hectares is planted with vines, the shock should have been devastating. Russia had absorbed 30% of Moldova’s wine exports; the ban was sudden, comprehensive, and explicitly punitive—retaliation for Moldova’s decision to negotiate an Association Agreement with the European Union (Popescu, 2014). Yet three years later, Moldova’s total wine exports had recovered, rerouted almost entirely to Romania, Poland, and Germany. This paper asks whether the embargo’s costs were borne locally—whether the districts that grew Moldova’s wine paid a price measurable in economic activity—or whether the shock was absorbed so quickly that it left no subnational trace.

The question matters beyond Moldova. Russia has repeatedly weaponized trade access against smaller neighbors: wine and mineral water bans against Georgia (2006), dairy import restrictions against Lithuania (2009, 2013), and the broader counter-sanctions against EU agricultural imports after 2014 (Drezner, 2003; Felbermayr et al., 2020). The theoretical literature on economic coercion predicts that such measures work when the target economy is concentrated, the sanctioning country controls a large market share, and adjustment costs are high (Hufbauer et al., 2007). Moldova in 2013—small, wine-dependent, with Russia as its dominant trade partner—appeared to satisfy all three conditions.

I test this prediction using a Bartik shift-share design. The “shift” is the September 2013 embargo, a binary event affecting all wine exports to Russia simultaneously. The “share” is pre-embargo vineyard hectares per capita at the district (*raion*) level, drawn from Moldova’s 2011 General Agricultural Census. This cross-sectional variation in wine dependence ranges from near zero in the capital Chisinau to 14% of population-equivalent hectares in the southern wine heartland districts of Leova and Cantemir. I measure local economic activity using VIIRS monthly nighttime light radiance for all 37 Moldovan administrative units from January 2012 to May 2024, providing 20 months of pre-embargo baseline and over a decade of post-embargo adjustment.

The main finding is a null. Across all specifications—continuous treatment intensity, binary above-median splits, log and level outcomes—I find no statistically significant decline in nightlight radiance in wine-dependent districts following the embargo. The point estimate from the preferred continuous Bartik specification is small (0.45 log points per unit vineyard share, $p = 0.40$) and confirmed by pairs cluster bootstrap ($p = 0.41$) and randomization inference ($p = 0.43$). When I add district-specific linear trends to absorb pre-existing differential trajectories, the coefficient turns weakly negative (-0.33 , $p = 0.22$) but remains far from conventional significance. Period-specific estimates show a suggestive short-run

decline in the nine months immediately following the embargo (-0.38 , $p = 0.16$), but the effect dissipates completely in the medium and long run. These estimates correspond to a standardized effect of 0.030 standard deviations—classified as “small positive” in magnitude, well within the null range.

This null is informative for two reasons. First, the trade shock itself is unambiguously large. Moldova’s wine exports to Russia fell from \$40.5 million in 2012 to \$10.2 million in 2014—a 75% collapse confirmed in UN Comtrade data. Second, the null occurs despite wine-dependent districts having genuinely different economic structures: the top-10 wine districts have vineyard concentrations 5–45 times higher than the capital region. If trade weaponization imposes local costs, this setting should reveal them.

I interpret the null as evidence of rapid trade diversion. Moldova signed its EU Association Agreement in June 2014, eight months after the embargo. The Deep and Comprehensive Free Trade Area (DCFTA) eliminated EU tariffs on Moldovan wine, and exports to Romania, Poland, and the Czech Republic surged. By 2019, Moldova’s total wine exports to Russia had partially recovered to \$18 million, but the composition of trade partners had fundamentally shifted. The speed of this adjustment—visible in both aggregate trade data and the absence of local nighttime decline—suggests that credible alternative markets dramatically reduce the coercive power of trade embargoes.

This paper contributes to three literatures. First, it provides the first subnational satellite-based evidence on the local economic costs of geopolitical trade embargoes, complementing aggregate analyses of Russian sanctions effects (Crozet and Hinz, 2020; Ahn et al., 2011). Second, it contributes to the trade adjustment literature by showing that adjustment to a severe sectoral shock can occur rapidly when institutional pathways (the EU accession process) channel reallocation (Autor et al., 2013; Dix-Carneiro and Kovak, 2017). Third, it adds to the growing literature using nighttime lights as proxies for economic activity in data-scarce environments (Henderson et al., 2012; Gibson et al., 2021), documenting both the promise and limitations of this approach for identifying localized trade shocks in small, open economies.

2. Background and Data

The Wine Embargo. Moldova is among the world’s most wine-dependent economies. Vineyards covered approximately 137,000 hectares in 2011—roughly 4% of the country’s total area—and wine accounted for 7–10% of GDP and a quarter of agricultural exports. The wine industry is geographically concentrated in the southern and central regions, with the Cahul, Cantemir, Leova, Stefan Voda, and Gagauzia districts forming the core of commercial

production.

Russia imposed its first wine embargo on Moldova (and Georgia) in March 2006, officially citing quality concerns but widely understood as political retaliation. That embargo was partially lifted in 2007 for select producers. The second embargo, in September 2013, was comprehensive: all Moldovan wine imports were banned. The timing coincided precisely with the intensification of Moldova’s EU Association Agreement negotiations, and Russian officials made the political motivation explicit (Popescu, 2014).

The trade collapse was immediate and severe. Moldova’s wine exports to Russia fell from \$40.5 million in 2012 to \$10.2 million in 2014 (UN Comtrade, HS 2204). The partial recovery after 2015 (\$11–18 million annually) reflects Russia’s selective reopening of the market for certain Moldovan wines, never returning to pre-embargo levels. Crucially, aggregate Moldovan wine exports recovered faster than the Russia-specific channel, as producers redirected to Romania, Poland, Germany, and the Czech Republic. By 2016, EU-destined wine exports had nearly doubled relative to 2012 levels. This trade diversion is central to interpreting the subnational results.

Moldova’s Administrative Geography. Moldova is divided into 37 administrative units (*raions*), plus the autonomous territorial unit of Gagauzia and the contested Transnistria region. These districts range from the capital Chisinau (population 821,000) to small rural raions like Soldanesti (38,000). The wine-producing heartland lies in the southern and central regions, where temperate continental climate, chernozem soils, and long growing seasons create favorable viticultural conditions. Northern raions—Briceni, Soroca, Edinet—are oriented toward grain, sugar beet, and dairy production, providing a natural comparison group for the shift-share design.

Nighttime Light Data. I use monthly VIIRS nighttime light composites from the EOAtlas project (Henderson et al., 2012), which provides pre-extracted zonal statistics for all administrative units worldwide. The dataset covers January 2012 to May 2024, yielding 149 monthly observations per district. The primary outcome is mean radiance (nW/cm²/sr) within each administrative boundary. I also examine total radiance (sum across pixels) as an alternative outcome that weights larger districts proportionally.

Nighttime lights are an established proxy for economic activity at the subnational level, with an elasticity of approximately 0.3 with respect to GDP in developing countries (Henderson et al., 2012; Gibson et al., 2021). However, the relationship is noisier for small administrative units and agricultural economies where much activity occurs during daytime. For Moldova’s wine regions, where the primary economic activity is vineyard cultivation and wine production during daylight hours, nightlights may understate the true economic impact of the embargo.

This measurement concern works against finding an effect, making any null result harder to interpret as evidence of no economic damage.

Treatment Intensity. Vineyard hectares per capita by district come from Moldova’s 2011 General Agricultural Census, measured two years before the embargo and thus predetermined. This predetermined nature satisfies the exclusion restriction for the Bartik shift-share design (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022): conditional on district and time fixed effects, the 2011 vineyard share should affect nightlights only through exposure to the 2013 trade shock.

The distribution of vineyard intensity exhibits substantial variation across Moldova’s districts. Southern wine districts like Leova (0.14 ha/cap) and Cantemir (0.14 ha/cap) are 45 times more vineyard-intensive than Chisinau (0.003 ha/cap). The top 10 wine-producing districts average 0.10 vineyard hectares per capita, compared to 0.01 for the bottom 10. I use this continuous measure as the Bartik “share” and also construct a binary indicator for above-median wine dependence (18 high-wine vs. 19 low-wine districts).

Table 1 presents summary statistics by wine dependence group. High-wine districts are substantially dimmer in nightlight terms (mean radiance 0.48 vs. 1.15 nW/cm²/sr pre-embargo), consistent with their more rural, agricultural character. This level difference is absorbed by district fixed effects and does not threaten identification. The two groups are observed for 20 pre-embargo months and 129 post-embargo months, providing a panel of 5,513 district-month observations (5,365 after excluding months with missing radiance values).

3. Empirical Strategy

The identification strategy exploits the interaction of a common aggregate shock (the embargo) with cross-sectional variation in exposure (vineyard intensity):

$$Y_{rt} = \alpha_r + \gamma_t + \beta \cdot (V_r \times \text{Post}_t) + \varepsilon_{rt} \quad (1)$$

where Y_{rt} is log mean nighttime radiance in district r at month t , α_r are district fixed effects, γ_t are year-month fixed effects, V_r is vineyard hectares per capita (predetermined), and Post_t indicates months from September 2013 onward. Standard errors are clustered at the district level (37 clusters). Given the small number of clusters, I supplement analytic standard errors with pairs cluster bootstrap (999 replications) and randomization inference (999 permutations of vineyard shares across districts).

The coefficient β captures whether districts with higher pre-embargo wine dependence experienced differential changes in nightlight radiance after the embargo, relative to less wine-

dependent districts, controlling for common time shocks and permanent district differences.

Threats to Identification. Three concerns merit discussion. First, the EU Association Agreement (signed June 2014) provides a concurrent positive shock to districts positioned to capture EU trade. If wine-producing districts benefited disproportionately from EU market access, this would attenuate any embargo-induced decline and bias β toward zero. I address this by estimating period-specific effects that separate the “pure embargo” window (September 2013–May 2014) from the EU integration period.

Second, Russia’s 2006 wine embargo may have already triggered adaptation in wine-dependent districts, reducing their vulnerability to the 2013 shock. This would make the 2013 null less surprising but more informative: it would suggest that learning from prior coercion erodes the weapon’s effectiveness over time.

Third, with 37 clusters, inference requires care. I report three approaches: analytic cluster-robust standard errors, pairs cluster bootstrap, and randomization inference. All three yield consistent conclusions.

4. Results

Main Results. [Table 2](#) presents the main estimates. Column (2), the preferred specification with log radiance and continuous treatment, yields $\hat{\beta} = 0.447$ (SE = 0.530, $p = 0.40$). The point estimate is positive—the opposite of the expected sign—and statistically indistinguishable from zero. Pairs cluster bootstrap confirms: SE = 0.573, 95% CI [−0.42, 1.64], $p = 0.41$. The binary specification in Column (4) shows a marginally significant positive effect ($\hat{\beta} = 0.099$, $p = 0.037$), suggesting that if anything, high-wine districts grew slightly faster in nightlight terms after the embargo. However, this result should be interpreted cautiously given the pre-trend evidence below.

Event Study and Pre-Trends. The event study ([Table 5](#)) reveals substantial noise in pre-treatment coefficients. Several pre-embargo bins are statistically significant, with magnitudes comparable to post-embargo coefficients. This pattern suggests that wine-dependent and non-wine districts were already on different trajectories before September 2013, likely reflecting urbanization trends, the 2006 embargo’s lingering effects, or structural divergence between agricultural and service-oriented districts. The pre-trend violations complicate causal interpretation of the post-embargo coefficients.

Period-Specific Effects. [Table 3](#) decomposes the post-embargo period into three windows. The short-run estimate (September 2013–May 2014) is negative ($\hat{\beta} = -0.382$, $p = 0.16$),

consistent with a transient local cost during the pure embargo window before the EU agreement. The medium-run (June 2014–December 2016) estimate is essentially zero (0.048), and the long-run (2017–2024) estimate is positive (0.669) but insignificant. This pattern is consistent with an initial shock followed by rapid adjustment and potential catch-up.

Robustness. Table 4 presents six robustness checks. The most informative is the specification with district-specific linear time trends (Panel C), which absorbs the pre-existing differential trajectories documented in the event study. With this control, the coefficient turns negative (-0.334 , $SE = 0.267$) but remains insignificant ($p = 0.22$). The placebo test at September 2012 (Panel B) yields a significant coefficient (-0.306 , $p = 0.002$), confirming that wine-dependent districts were already declining relative to others before the embargo. Excluding Transnistria or Chisinau does not materially change the results. Randomization inference yields $p = 0.43$.

The leave-one-out analysis reveals that the estimate is sensitive to the inclusion of Bender (a small urban district with extreme nightlight values): dropping Bender shifts the coefficient from 0.45 to 0.01. This sensitivity underscores the fragility of the positive point estimate.

Statistical Power. Given the noisy outcome and 37 clusters, it is important to ask what effect size the design could detect. Using the analytic standard error from the preferred specification ($SE = 0.53$) and a one-sided test at $\alpha = 0.05$, the minimum detectable effect (MDE) at 80% power is approximately $|\beta| = 0.53 \times (1.64 + 0.84) = 1.31$ in the continuous Bartik specification, or 0.087 in standardized units ($SDE = \beta \times SD(X) / SD(Y) = 1.31 \times 0.043 / 0.653$). This means the design can rule out large negative effects ($SDE > 0.087$, approximately a 9% of a standard deviation decline) but cannot detect moderate effects in the 0.02–0.08 SDE range. The trend-controlled specification, with its smaller standard error (0.267), yields a tighter MDE of 0.66 ($SDE = 0.044$), capable of ruling out effects above 4.4% of a standard deviation. The design is thus informative for ruling out catastrophic local decline but may miss economically meaningful moderate costs.

5. Discussion

Why No Local Decline?. Three mechanisms could explain the null. First, rapid trade diversion to the EU limited the aggregate trade loss. Moldova’s wine exports to Romania alone grew from \$3 million in 2012 to \$12 million by 2016, while exports to Poland and the Czech Republic also surged. The EU Association Agreement’s DCFTA provisions, which eliminated EU tariffs on Moldovan wine beginning in September 2014, created an institutional channel for this adjustment that was unavailable during the 2006 embargo. The timing is revealing: the short-run coefficient (-0.38 , September 2013–May 2014) is the most negative

period-specific estimate, suggesting a transient cost during the nine months before EU market access expanded. That this short-run effect is statistically insignificant does not rule out meaningful but transient local costs; it may instead reflect nightlights' limited sensitivity to agricultural sector shocks.

Second, the 2006 embargo may have already induced adaptation—reducing Russia-specific capital, encouraging market diversification, and building export capacity for EU standards. If wine-dependent districts had already begun reorienting toward EU markets after the 2006–2007 experience, the 2013 shock would have been less disruptive than a first embargo. This interpretation is consistent with the “learning” hypothesis in the sanctions literature: repeated use of trade weaponization erodes the weapon’s coercive power as targets develop alternative networks (Hufbauer et al., 2007).

Third, nighttime lights may be insufficiently sensitive to detect a shock concentrated in the agricultural sector. Wine production occurs during daylight hours; processing, storage, and transport contribute modestly to nightlight-generating activity. A 75% trade collapse that reduces winery revenue without closing factories or reducing urban commercial activity could plausibly leave no satellite signature. This is a genuine measurement limitation, and future work with firm-level or household survey data would provide a sharper test.

What the Paper Cannot Identify. The null in nightlights does not mean the embargo imposed no costs. Adjustment to new markets involves real resources: compliance with EU phytosanitary standards, investment in new distribution networks, rebranding for EU consumers, and the deadweight loss of destroyed trade relationships. These costs may be absorbed by wineries and workers in ways that reduce welfare without reducing aggregate district-level nightlight radiance. Similarly, the embargo may have depressed wages or employment in wine-producing districts without reducing the total number of lit structures or commercial establishments that drive nightlight variation.

Firm-level or worker-level data—unavailable for Moldova at this subnational granularity—would be needed to detect such micro-level adjustment costs. Moldova’s Labor Force Survey and enterprise registries could in principle provide this evidence, but access requires partnerships with the National Bureau of Statistics that are beyond the scope of this study.

Limitations. The pre-trend violations documented in the event study and placebo test represent the paper’s most serious limitation. The significant placebo coefficient at September 2012 (-0.31 , $p = 0.002$) indicates that wine-dependent districts were already declining relative to non-wine districts before the embargo. This pre-existing divergence could reflect urbanization trends drawing activity toward Chisinau and Balti, secular agricultural decline in rural raions, or lingering effects of the 2006 embargo. While the trend-controlled specification

($\beta = -0.33$, $p = 0.22$) partially addresses this concern by absorbing district-specific linear trajectories, it cannot rule out nonlinear differential trends that mimic or mask the embargo’s effect.

The leave-one-out analysis reveals that the positive point estimate in the baseline specification is sensitive to the inclusion of Bender, a small urban district with extreme nightlight values. Dropping Bender shifts the continuous Bartik coefficient from 0.45 to 0.01, underscoring the fragility of any conclusion about the sign of the effect. The vineyard hectare data from the 2011 census may contain measurement error, though classical measurement error in the treatment variable would attenuate the estimates toward zero, making the null harder to reject. Finally, the study uses administrative boundaries that do not perfectly correspond to wine-producing zones—some “wine districts” contain substantial non-wine economic activity that dilutes the treatment effect.

6. Conclusion

Russia’s 2013 wine embargo destroyed 75% of Moldova’s wine trade with its largest partner. Yet wine-dependent districts show no detectable decline in satellite-measured economic activity relative to other districts. The absence of local costs—despite a severe and sudden trade shock—is consistent with rapid trade diversion enabled by Moldova’s concurrent EU integration. The finding suggests that geopolitical trade weaponization is a blunt instrument when the target has credible alternative market access. For Moldova, Russia’s embargo may have inadvertently accelerated the very westward economic reorientation it sought to prevent.

References

- Ahn, JaeBin, Amit K. Khandelwal, and Shang-Jin Wei**, “The Role of Intermediaries in Facilitating Trade,” *Journal of International Economics*, 2011, *84* (1), 73–85.
- Autor, David H., David Dorn, and Gordon H. Hanson**, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 2013, *103* (6), 2121–2168.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-Experimental Shift-Share Research Designs,” *Review of Economic Studies*, 2022, *89* (1), 181–213.
- Crozet, Matthieu and Julian Hinz**, “Friendly Fire: The Trade Impact of the Russia Sanctions and Counter-Sanctions,” *Economic Policy*, 2020, *35* (101), 97–146.
- Dix-Carneiro, Rafael and Brian K. Kovak**, “Trade Liberalization and Regional Dynamics,” *American Economic Review*, 2017, *107* (10), 2908–2946.
- Drezner, Daniel W.**, “The Hidden Hand of Economic Coercion,” *International Organization*, 2003, *57* (3), 643–659.
- Felbermayr, Gabriel, Constantinos Syropoulos, Erdal Yalcin, and Yoto V. Yotov**, “The Economic Effects of Sanctions on Russia,” *European Economic Review*, 2020, *129*, 103561.
- Gibson, John, Susan Olivia, Geua Boe-Gibson, and Chao Li**, “Night Lights in Economics: Sources and Uses,” *Journal of Economic Surveys*, 2021, *35* (4), 955–980.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik Instruments: What, When, Why, and How,” *American Economic Review*, 2020, *110* (8), 2586–2624.
- Henderson, J. Vernon, Adam Storeygard, and David N. Weil**, “Measuring Economic Growth from Outer Space,” *American Economic Review*, 2012, *102* (2), 994–1028.
- Hufbauer, Gary Clyde, Jeffrey J. Schott, Kimberly Ann Elliott, and Barbara Oegg**, *Economic Sanctions Reconsidered*, 3rd ed., Peterson Institute for International Economics, 2007.
- Popescu, Nicu**, “Moldova’s Fragile Pluralism,” in “EU Neighbours in Crisis,” European Union Institute for Security Studies, 2014.

Table 1: Summary Statistics: Nighttime Light Radiance by Wine Dependence

	Pre-Embargo (Jan 2012–Aug 2013)		Post-Embargo (Sep 2013–May 2024)	
	High Wine	Low Wine	High Wine	Low Wine
Mean radiance (nW/cm ² /sr)	0.481	1.148	0.497	1.042
SD	(0.153)	(1.713)	(0.201)	(1.581)
Log mean radiance	-0.752	-0.351	-0.744	-0.442
Vineyard ha per capita	0.079	0.014	—	—
Raions	18	19	18	19
Raion-months	360	380	2,322	2,451

Notes: Nighttime light radiance from VIIRS monthly composites (EOAtlas, 2012–2024). High Wine = above-median vineyard hectares per capita from the 2011 General Agricultural Census. 37 administrative units (raions) observed monthly. Embargo date: September 2013.

Appendix

Table 2: Effect of Wine Embargo on Nighttime Light Radiance

	(1)	(2)	(3)	(4)
	Level	Log	Level	Log
Vineyard Share \times Post	0.979 (2.031)	0.447 (0.530)		
High Wine \times Post			0.122 (0.172)	0.099** (0.046)
Treatment	Continuous	Continuous	Binary	Binary
Raion FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Observations	5,365	5,365	5,365	5,365
R^2	0.798	0.948	0.798	0.948
Within R^2	0.0007	0.0019	0.0015	0.0129
Bootstrap SE		[0.573]		
Bootstrap p		0.408		

Notes: Standard errors clustered at the raion level in parentheses. Pairs cluster bootstrap SE in brackets (999 replications). Vineyard Share = vineyard hectares per capita from 2011 Agricultural Census. High Wine = above-median vineyard share. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix B: Standardized Effect Sizes

Acknowledgements

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

Contributors: @SocialCatalystLab

First Contributor: <https://github.com/SocialCatalystLab>

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

Table 3: Period-Specific Effects of the Wine Embargo

	(1)	(2)
	Log Radiance Pooled Post	Log Radiance By Period
Vineyard Share \times Post	0.447 (0.530)	
Vineyard Share \times Short Run (Sep 2013–May 2014)		-0.382 (0.267)
Vineyard Share \times Medium Run (Jun 2014–Dec 2016)		0.048 (0.352)
Vineyard Share \times Long Run (Jan 2017–May 2024)		0.669 (0.658)
Raion FE	Yes	Yes
Year-month FE	Yes	Yes
Observations	5,365	5,365

Notes: Dependent variable: log mean nighttime radiance. Standard errors clustered at the raion level. Short Run = September 2013 to May 2014 (pre-EU Association Agreement). Medium Run = June 2014 to December 2016 (EU agreement signed, initial adjustment). Long Run = January 2017 to May 2024 (full adjustment). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4: Robustness Checks

	Coefficient	SE	N
<i>Panel A: Main specification</i>	0.447	(0.530)	5,365
<i>Panel B: Placebo (Sep 2012)</i>	-0.306***	(0.092)	740
<i>Panel C: Raion-specific trends</i>	-0.334	(0.267)	5,365
<i>Panel D: Exclude Transnistria</i>	0.401	(0.555)	5,220
<i>Panel E: Exclude Chisinau</i>	0.634	(0.525)	5,220
<i>Panel F: Log sum radiance</i>	0.453	(0.532)	5,365
Randomization inference p -value		0.432	

Notes: Dependent variable: log mean nighttime radiance in all panels except F. All specifications include raion and year-month fixed effects. Standard errors clustered at the raion level. Panel B uses only pre-embargo observations (Jan 2012–Aug 2013) with a placebo treatment date of September 2012. Panel C adds raion-specific linear time trends. Randomization inference permutes vineyard shares across raions (999 replications). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: Event Study: Bimonthly Coefficients

	Estimate	SE
<i>Pre-embargo</i>		
$t - 20$	-0.561	(0.374)
$t - 18$	-1.156***	(0.415)
$t - 16$	-0.736	(0.695)
$t - 14$	-0.906**	(0.365)
$t - 12$	-0.069	(0.611)
$t - 10$	-1.786***	(0.457)
$t - 8$	-1.556***	(0.478)
$t - 6$	-0.705	(0.464)
$t - 4$	-2.759***	(0.477)
<i>Post-embargo</i>		
$t + 0$	-0.301	(0.463)
$t + 2$	-1.520***	(0.380)
$t + 4$	-2.351***	(0.520)
$t + 6$	-1.206**	(0.483)
$t + 8$	-0.299	(0.494)
$t + 10$	-1.429**	(0.546)
$t + 12$	-0.389	(0.547)
$t + 14$	-0.159	(0.803)
$t + 16$	0.789	(0.656)
$t + 18$	-0.899*	(0.499)
$t + 20$	-2.351***	(0.757)
$t + 22$	0.490	(0.681)
$t + 24$	-1.171**	(0.484)
Reference period	$t - 2$ (Jul–Aug 2013)	
Observations	1,628	

Notes: Coefficients from interacting bimonthly event-time indicators with vineyard hectares per capita. Reference period: two months before embargo (July–August 2013). Includes raion and year-month fixed effects. Standard errors clustered at the raion level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Log radiance (continuous)	0.447	0.530	0.653	0.030	0.035	Small positive
Log radiance (w/ trends)	-0.334	0.267	0.653	-0.022	0.018	Small negative
<i>Panel B: Heterogeneous</i>						
Rural raions only ($N = 34$)	0.371	0.209	0.653	0.025	0.014	Small positive
Short run only (Sep 2013–May 2014)	-0.382	0.267	0.653	-0.025	0.018	Small negative

Notes: **Country:** Moldova. **Research question:** Does Russia’s September 2013 wine import embargo impose subnational economic costs on wine-dependent administrative districts in Moldova? **Policy mechanism:** Russia banned Moldovan wine imports in September 2013 as geopolitical retaliation for Moldova’s EU Association Agreement negotiations, collapsing bilateral wine exports by 75% (\$40M to \$10M annually) and disproportionately affecting raions with high pre-embargo vineyard concentration. **Outcome definition:** Monthly mean VIIRS nighttime light radiance (nW/cm²/sr) at the administrative district (raion) level, a satellite-based proxy for local economic activity. **Treatment:** Continuous — vineyard hectares per capita from the 2011 General Agricultural Census, measuring pre-determined exposure to the wine export shock. **Data:** EOAtlas VIIRS monthly nightlights (2012–2024), 37 raions, 5,365 raion-months after excluding missing observations; UN Comtrade for aggregate trade validation. **Method:** Bartik shift-share DiD with raion and year-month fixed effects, standard errors clustered at raion level, pairs cluster bootstrap (999 replications), and randomization inference (999 permutations). **Sample:** All 37 Moldovan administrative districts (raions) including Gagauzia autonomous region and Transnistria; 20 pre-embargo months and 129 post-embargo months. $SDE = \hat{\beta} \times SD(X)/SD(Y)$ where $SD(X)$ is the cross-raion standard deviation of vineyard hectares per capita (0.043) and $SD(Y)$ is the pre-embargo standard deviation of log mean radiance (0.653). Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).