Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Can Tourism Extend a Peripheral Region's External Market? Evidence from Madeira Wine

Paulo Henrique Barbosa

paulofbarbosa5@gmail.com

University of Lisbon

António Almeida

Universidade da Madeira

Luiz Pinto Machado

Universidade da Madeira

Research Article

Keywords: wine tourism, regional development, tourism–trade linkages, export promotion, island economies, Madeira wine

Posted Date: August 14th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10471858/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Can Tourism Extend a Peripheral Region's External Market? Evidence from Madeira Wine

July 2026

Abstract

A recurring wine-tourism claim holds that visitors who taste a place-based product later buy it at home, so tourism doubles as export promotion. We test whether it leaves a footprint in administrative data, matching overnight stays by residence country to Madeira wine shipments by destination across twelve European markets, 1993–2024. The question matters for tourism-specialised island regions, whose development strategies often credit the visitor economy with creating export demand. Shipments show no stable association with tourism in the visit year or afterwards, with the estimated association bounded between -0.50 and 0.37 by its 95 percent confidence set, far below the large effects in cross-country tourism–trade studies. One positive association appears in a single specification and window that no other sample reproduces. A calibration shows why a null is not proof of absence: buying by fewer than about 46 percent of guests leaves no detectable trace, and visit-year purchases are effectively invisible. The data record shipments, not what visitors buy, and distribution channels can divert the difference.

Keywords: wine tourism, regional development, tourism–trade linkages, export promotion, island economies, Madeira wine.

JEL classification: R11, Z32, F14, L66, C23.

1 Introduction

Tourism-specialised peripheral regions face a durability question. Visitors generate expenditure while present, but the benefit ends at departure unless the visit creates something that outlives it, and one candidate is external demand for the region's place-based products. This paper asks whether that channel leaves a measurable trace for a region built around both, an island whose visitor economy dominates activity and whose flagship export carries the island's name.

Much of the wine and food tourism literature rests on an appealing idea: a visitor who encounters a distinctive local product on holiday forms an attachment to it and buys it after returning home, so that tourism generates a lasting stream of demand in the visitor's home market. If true, this makes destination tourism a form of export promotion for place-based products, and it is a common justification for cellar-door programmes, wine routes, and gastronomic branding (Hall et al., 2003; Mitchell and Hall, 2006; Croce

and Perri, 2010; Gómez et al., 2019). The idea sits within a longer tradition on the local economic linkages of tourism (Telfer and Wall, 1996; Torres, 2003).

The evidential base for the claim is lopsided in two directions at once, which is what makes a direct test worth having. On one side, the wine tourism literature is dominated by surveys of visitor motivations, satisfaction, and stated purchase intentions, which capture what visitors say they will do rather than what later appears in trade, a reliance the field's own systematic review documents (Gómez et al., 2019). The behavioural work is nonetheless clear about the on-site step: the primary motivations of winery visitors are to taste and to buy wine, the visit is hedonic, and purchase at the point of experience is widespread (Charters and Ali-Knight, 2002; Alant and Bruwer, 2004; Bruwer and Alant, 2009; Getz and Brown, 2006). What surveys cannot show is the second step, whether the attachment survives the flight home and re-emerges as purchases in the visitor's own market, and the channels through which it could operate, information, preference formation, and familiarity with the origin, are exactly those the trade literature credits for tourism's export effects (Kulendran and Wilson, 2000; Santana-Gallego et al., 2016; Kumar et al., 2019).

On the other side, the trade literature does find effects that we review in Section 2, with elasticities close to one and lagged responses in the broad product classes (El-Sahli, 2018), and the study closest to ours in product and design finds a long-run relationship between German tourism to Spain and German imports of Spanish wine (Fischer and Gil-Alana, 2009). A null result against this backdrop is informative in a way a null against brochures would not be: it says that a mechanism which shows up in broad product classes leaves no visible footprint for a flagship heritage product in the very data where promoters would look for it.

Madeira offers an unusually clean confrontation. The island is a mature destination and the home of Madeira wine, a fortified wine with a protected designation of origin. The regional statistical office records tourist overnight stays by country of residence, and the wine institute records wine shipments by destination country, so the two administrative series can be matched country by country and year by year over three decades. If tasting-to-purchase operates at a scale visible in aggregate data, countries whose residents visit Madeira more should, contemporaneously or with a lag, receive more Madeira wine.

We match the two series over 1993–2024 for the 12 covered European origin countries observed in both, 7 of them in every year, and ask three descriptive questions. Do shipments and overnights co-move within a country over time? Does the association persist into later calendar years, the pattern a durable returning-customer channel would produce, rather than on leads? And does a pooled cross-correlation profile peak at any positive horizon? We then do what a null result obliges: we state the detection limits of the design, invert its own inference procedure into a confidence set for the cumulative association within the mechanism's admissible window, and state its detection limits in visitor-conversion units by injecting the hypothesised channel into the outcome and asking when the test would fire, so that the finding reads as a measured footprint with a stated detection limit rather than a claim of absence. Because the sample ends in 2024, the pandemic years receive explicit treatment throughout.

The answer comes in three parts, each carrying its caveat. First, the footprint is absent where the claim predicts it. The cumulative association over the same-year and lagged window is 0.107, bounded between -0.50 and 0.37, no persistent delayed component appears in any sample or outcome, and the cross-correlation profile peaks at no positive horizon.

Second, the one positive cumulative estimate, 0.271 in the pre-2018 window with leads included, is confined to that single specification cell, carried by the same-year term that annual data cannot separate from common demand shocks and importer order timing, and reproduced by no no-lead window. We fold it into the cumulative bound rather than dismissing it on timing grounds.

Third, the design's resolution is measured rather than assumed. Injecting the hypothesised channel into the outcome, detection first occurs above roughly 46 percent of guests under the most detectable purchase-timing calibration, and conversion concentrated in the visit year escapes the design entirely, so the informative object is the confidence set on the recorded shipment association, not a conversion rate. Throughout, the outcome we observe is shipments to importers and distributors, not purchases by returning visitors. Inventory and contracting decouple the two at short horizons, which is why we read cumulative windows rather than individual coefficients, and channels neither outcome sees, direct-to-consumer shipping, travel retail, and third-country routing, are priced by the framework of Section 2.2 as leakage rather than ignored.

The paper makes three contributions. To regional economics, it tests whether tourism specialisation extends a peripheral region's external market for a place-based tradable, matching the origin geography of visitors to the destination geography of the region's flagship export, and it shows that the elasticities near one estimated on cross-country aggregates do not reappear at the product level for this region, so development arithmetic cannot borrow the aggregate number. To empirical practice with small regional panels, it offers a template for informative null results: enumerated wild-cluster inference, confidence sets obtained by test inversion, and a realised-sample detectability calculation that states what the design cannot see before a referee asks. To the wine and gastronomic tourism literature, it provides, to our knowledge, the first country-matched administrative test of the post-visit purchase claim, replacing stated intentions with recorded flows and a four-link framework whose hypotheses the data can reject.

The question is also a regional one. Madeira is a tourism-specialised outermost island region whose development strategy pairs a dominant visitor economy with a small portfolio of place-based exports, and a large literature asks whether tourism specialisation feeds broader regional growth (Brida et al., 2016). We ask the narrower, measurable version for the region's flagship export, whether the visitor economy feeds the export base, with the twelve origin markets forming the region's export geography.

The stakes extend beyond the wine literature, because creating foreign demand is a costly public activity. Governments fund export promotion precisely because new markets are expensive to reach, and the evaluation record on paid instruments is instrument-specific. Trade missions often show little effect (Head and Ries, 2010), the benefits of export support can fade (Cadot et al., 2015), and well-targeted services helped firms weather the crisis (Van Biesebroeck et al., 2016). On the other hand, matchmaking and grants raise Portuguese firms' exports, particularly for the smallest firms (citation anonymised for review).

Against that record, post-trip purchasing by visitors is often treated as an incidental export spillover of tourism expenditure, an extra return on promotion budgets that no separate market-access intervention has to earn. The claim's weak link is the last mile. A visitor who tastes the wine in Funchal and looks for it at home months later completes the chain only if an importer lists the product, a retailer stocks it, or a direct or travel-retail channel exists, and if none does, the preference expires or travels in a suitcase the statistics never see. Destination bodies, product boards, exporters, distributors, and

trade-promotion organisations each control different and overlapping parts of that chain, so the policy problem is coordination and measurement rather than a tidy division of labour.

Our results discipline the claim at the point where these parts meet. Whatever tourism does for later demand, it leaves no stable footprint in recorded shipments, and the post-return part of the chain, availability in the origin market, listings, direct channels, and travel retail, remains unmeasured and potentially binding. For regional governments and product boards the immediate implication is about indicators rather than budgets: visitor arrivals are exposure measures, not export-conversion outcomes, so growth in arrivals is not by itself evidence that the home-market link has operated, and programmes that seek that link should provide and evaluate market access separately rather than assume it.

The remainder of the paper proceeds as follows. Section 2 reviews what the wine-tourism and tourism–trade literatures establish and sets out the framework and hypotheses. Section 3 describes the two administrative sources, the residence-to-destination match, and its coverage. Section 4 presents the dynamic estimates, the cumulative association across windows, samples, and controls, and the detectability calibration. Section 5 discusses the results against the trade literature, the measurement gap, and destination management, and Section 6 concludes with what would settle the questions this design cannot.

2 Related literature

2.1 What the literatures establish

Three literatures bear on the question, and the paper sits at a point where the first two make a mechanism plausible while the third shows that its central object is rarely measured.

The first is the econometrics of tourism and trade. Instrumented cross-country estimates report that inbound tourism raises exports of differentiated consumer goods and processed foods, with elasticities close to one and with lagged responses (El-Sahli, 2018), disaggregated European evidence points the same way (Brau and Pinna, 2013; Santana-Gallego et al., 2016; Kulendran and Wilson, 2000), and the closest product-level study finds a long-run link between German tourism to Spain and German imports of Spanish wine (Fischer and Gil-Alana, 2009). The mechanisms these papers invoke are the ones a returning-visitor story needs, information, preference formation, and familiarity with the origin, and they are the same mechanisms the trade-and-mobility literature credits for the effect of migration and networks on bilateral trade in differentiated goods. Against this backdrop our contribution is a disciplined contrast: we look for the footprint the mechanism predicts, cumulative within the admissible window and persistent if the channel is durable, in administrative data for a single product, and find neither a stable full-sample footprint nor a persistent component. The elasticities near one in this literature lie far outside the confidence sets we report.

The second is the behavioural evidence on wine tourism. Surveys of winery visitors consistently find that tasting and buying are the primary motivations, that the visit is hedonic, and that on-site purchase is widespread (Charters and Ali-Knight, 2002; Alant and Bruwer, 2004; Bruwer and Alant, 2009; Getz and Brown, 2006). A longitudinal Australian study that followed cellar-door visitors after their visit reports that roughly two-thirds bought wine at the cellar door and that a substantial minority repurchased the

brand within six months, mostly through domestic retail and direct-to-consumer channels rather than conventional distribution (Bruwer et al., 2015). This is exactly the evidence that makes the mechanism plausible, and exactly the evidence our outcome cannot capture: the repurchase it documents runs through channels that do not appear in shipments to the visitor's country, so the behavioural finding and our null are complementary rather than contradictory.

The third is the measurement problem that separates the two. What the promotional claim asserts is post-visit purchase of a specific place-based product in the visitor's home market. What administrative shipment data record is flows to importers and distributors. We are unaware of a published study that quantifies the share of a region's international visitors who later buy its wine at home, and none documents a lagged tourism-to-shipments effect peaking at a one-year or longer horizon, so there is no direct benchmark for the object the claim is about. Our contribution is to test the closest observable administrative footprint and to bound what that footprint can reveal.

2.2 A framework and testable implications

The claim under test travels a chain with four links, and each link is a distinct economic object: a destination experience creates brand knowledge and preference, preference produces a purchase in the visitor's home market, the purchase draws down an importer's stock, and replenishment appears as a recorded shipment to that market. Figure 1 draws the chain. Every link leaks. Memory and preference decay between the visit and the next purchase occasion. A formed preference can be satisfied through channels that never touch the destination-country import statistics, direct-to-consumer shipping, travel retail, bottles in luggage, or a retailer supplied through a third-country hub. And an importer meets a small demand shift from inventory, so replenishment smooths and delays the footprint that does arrive.

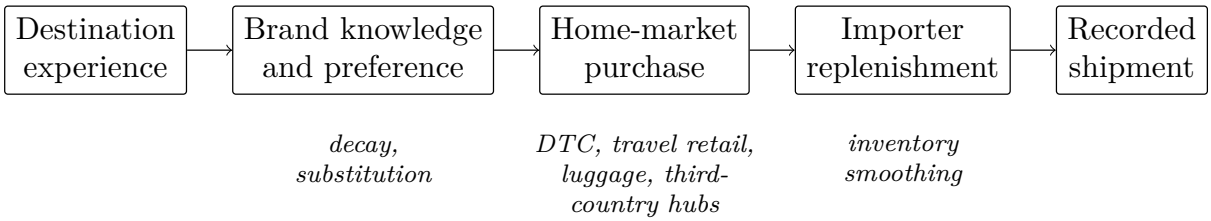


Figure 1: From destination experience to recorded shipments.

Notes: Each link can leak, so the recorded footprint can be attenuated, delayed, or redirected relative to consumer purchases, and the paper reports detectability alongside the estimates.

The chain delivers four implications that annual administrative data can confront. H1, the footprint: if the chain operates at scale, tourism enters positively through the current-plus-two-lag coefficient sum $\beta_0 + \beta_1 + \beta_2$ in equation (1), because annual aggregation places a purchase that follows a visit within the same calendar year or the next ones. H2, the timing placebo: under a forward-operating channel, future tourism should carry no additional information about current shipments once the past is conditioned on, so lead coefficients near zero are a consistency check rather than identification. H3, persistence, a conditional prediction: mass on the lagged terms appears only when purchasing or replenishment crosses calendar years, so a zero lag sum speaks to the timing and durability of the channel, not to its existence. H4, attenuation: routing, inventory, and alternative

channels can attenuate, delay, or redirect the recorded country-level footprint, so the mapping from purchases to shipments is not one for one, a null footprint does not imply a null consumer channel, and the design's resolution in conversion units must be measured rather than assumed, which Section 4.4 does by injection.

The framework fixes the reading order of the results: the cumulative object of H1 is primary, the leads of H2 are a timing placebo, the lag sum of H3 is evidence on persistence, and H4 obliges the paper to report detectability alongside the estimates.

3 Data

3.1 Sources and matched coverage

We use two administrative sources. Wine shipments come from the commercialisation census of the regional wine institute, which records for every certified operator the volume and value of Madeira wine commercialised in each destination market and year. The census covers the certified sector in full, and its destination breakdown separates shipments to each partner country from sales on the domestic market. Values are recorded in euro and quantities in litres, and we deflate values to constant 2024 euro with the price deflator carried in the source database (2024 = 100), constructing real shipment value, litres, and real unit value for each partner and year. The recorded flows are shipments by producers and shippers to importers and distributors rather than purchases by final consumers, a distinction that matters for interpretation and that we return to in Section 3.2. The collection basis is uniform across partner countries and stable over the sample, so cross-country comparability does not depend on customs or Intrastat reporting rules. The provisional final year is retained, and Table S1 in Online Resource 1 prints the estimates with it removed. Origin-country GDP per capita and population, used as controls in one specification, come from the macro covariate database assembled for this research programme from international statistical sources.

Table 1: Matched coverage: origin countries in both administrative sources

Country	Years	Span	Overnights (000s, mean)	Litres (000s, mean)	Value share (%)
Belgium	32	1993–2024	138	183.5	8.1
Denmark	32	1993–2024	154	64.7	3.8
France	32	1993–2024	421	1122.8	38.9
Germany	32	1993–2024	1,416	325.1	12.9
Netherlands	32	1993–2024	222	74.0	3.6
Sweden	32	1993–2024	169	109.5	5.8
United Kingdom	32	1993–2024	1,451	303.0	24.1
Austria	15	2010–2024	107	30.8	0.7
Finland	15	2010–2024	151	14.3	0.6
Italy	15	2010–2024	90	6.1	0.3
Norway	15	2010–2024	96	15.2	0.8
Spain	15	2010–2024	189	9.2	0.4

Notes. A country is classified as continuously observed only if present in all thirty-two years. Total matched observations 299. Means are taken over each country's matched years, and value shares are of total matched real shipment value. Source: wine commercialisation census of the regional wine institute (shipments by destination) and DREM retrospective tourism statistics (overnights, residence breakdown), values deflated to constant 2024 euro.

Tourism is measured by overnight stays and guests in registered accommodation by country of residence, from the retrospective tourism statistics of the regional statistical office DREM. The residence basis matches the partner-country basis of the shipment data conceptually: a resident of a given country generates overnights recorded against that country and, if the promotional mechanism operates, wine shipped to that same country. We match the two series by country and year over 1993–2024, keeping origin countries present in both sources. Table 1 reports the result, and Figure 2 draws it as the region's export geography: 12 countries and 299 country-year observations, with 7 present in all thirty-two years and five entering in 2010, when the DREM residence breakdown widens to report them separately. The twelve matched markets carry 67 percent of all recorded foreign shipment value and 80 percent of foreign volume over the sample, so the match covers most of the sector's exports while omitting some destinations, and the origin list is fixed by which countries the residence breakdown reports, a criterion independent of shipment outcomes. The matched panel contains no zero-shipment market-years, so the logarithmic specifications lose nothing to zeros, and Table S1 in Online Resource 1 nonetheless reports a Poisson pseudo-likelihood variant because the estimator also guards against log-transformation bias under heteroskedasticity.

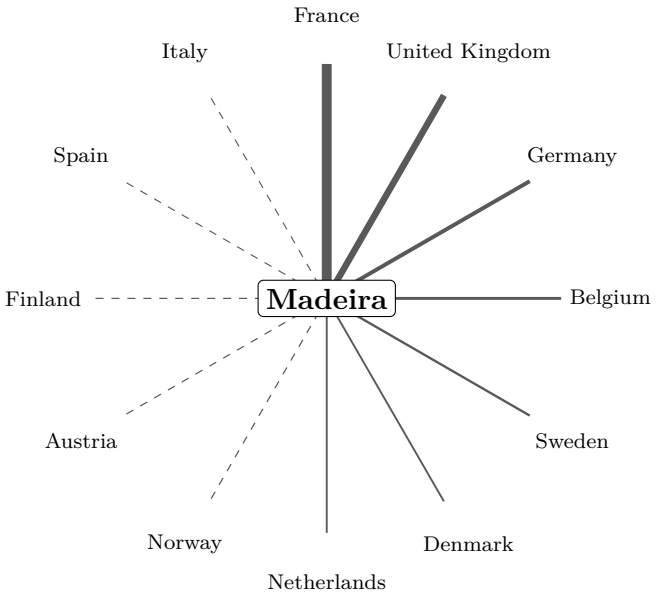


Figure 2: The region's export geography.

Notes: The twelve matched origin markets around the exporting region, with spoke width proportional to each market's share of matched real shipment value (Table 1). Solid spokes mark the seven continuously observed markets and dashed spokes the five entering the residence breakdown in 2010. Visitors travel inward along these spokes and the hypothesised demand travels outward along the same geography, which is the spatial structure the estimation exploits. Source: wine commercialisation census and DREM.

The matched sample is concentrated. France, the United Kingdom, and Germany together carry roughly three-quarters of matched shipment value, and the five countries entering in 2010 are individually small, none reaching one percent of value. Their observations are not negligible in an unweighted panel, however, since the dynamic sample falls from 251 to 196 rows when we restrict to the balanced seven, so we print that specification rather than assert that late entrants cannot matter, and Table S2 in Online Resource 1 removes each country in turn. The concentration is not a defect: these are the destinations with the largest visitor flows and the deepest trade, so they are where a home-market conversion channel, if it existed at the scale the promotional claim implies, would be most visible among the covered European markets.

3.2 Matched sample and measurement constraints

Three features bound what the exercise can establish, and the second is a scope condition the reader needs before the results. First, the residence breakdown is annual, so we work at annual frequency and cannot exploit within-year timing, a coarse grid for a mechanism that operates over months, and one consequence developed in Section 4.2 is that the same-year coefficient belongs to the mechanism's admissible window. Second, the origin list covers the main European markets and omits, among others, the United States and Switzerland, which have no tourism-residence coverage in the DREM series. These are two of the highest-unit-value destinations for this export, precisely the affluent markets where post-visit conversion is most plausible, so the matched sample excludes part of where the mechanism would bite hardest, and every result should be read as pertaining to the covered European markets. Third, and most fundamentally, the outcome is shipments to importers and distributors, not purchases by returning visitors. Inventory

holding and contracting decouple the two at short horizons, so the design tests the closest administrative footprint of the mechanism, not the mechanism's final object. Guests count arrivals at registered establishments rather than unique visitors. Repeat registrations inflate the count while unregistered stays, day visitors, and visits to friends and relatives escape it, so the mapping from arrivals to unique visitors is not signed, a definition that matters for the conversion calibration and that Section 4.4 carries explicitly. The covered countries account for 86 percent of foreign guests in registered accommodation over the matched years. We do not observe air capacity or direct-route availability, which move with tourism and could shape both series, and no specification controls for them. The sample also spans the pandemic, whose collapse in travel moved overnights and shipments through different channels. The 2020–2021 observations are shaded in the time-series exhibits, and Section 4.5 re-estimates the headline specifications with an exclusion at the regressor level, removing every observation whose own year or whose lag and lead terms reference 2020 or 2021.

Large language model assistance (Claude, Anthropic) was used in code development and manuscript drafting under the authors' direction and verification, and all analytical decisions and the final text are the authors' own.

4 The shipment footprint: estimates, windows, and bounds

4.1 Co-movement

We first ask whether shipments and overnights move together within a country over time, removing country-specific levels. Figure 3 shows the pooled within-country scatter, whose fitted slope is essentially flat. Value behaves differently at the same-year horizon, and Section 4.2 reports that association in full rather than letting the litres scatter speak for both outcomes.

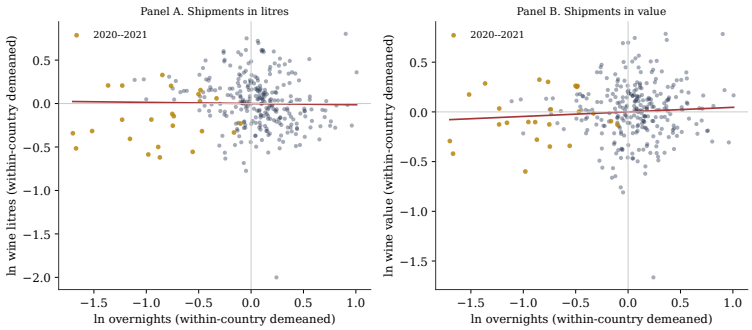


Figure 3: Within-country scatter of log wine shipments against log overnight stays, both demeaned by country.

Notes: Panel A is litres, Panel B is real value. The litres slope is essentially flat (pooled within-country correlation −0.02), while the value panel shows the weak positive same-year tilt that Table 2 quantifies and Section 4.2 shows to be fragile. Pandemic observations are highlighted in gold and sit off the main cloud, previewing Section 4.5. Source: wine commercialisation census and DREM.

Figure 4 plots the two series for the 7 continuously observed countries. There is no consistent co-movement: tourism trends upward in most origins while shipments drift or

decline, and the turning points do not line up. The pooled within-country correlation of the two logs is −0.02 for litres and 0.07 for value, the value figure mattering because the promotional claim centres on premium purchases, so value deserves the same scrutiny as volume. Country by country, the contemporaneous correlation is negative in six of the seven long-sample markets, reaching −0.56 for France, with Sweden the lone positive at 0.56, a dispersion that itself argues against a common mechanism and that Figure S1 in Online Resource 1 unpacks.

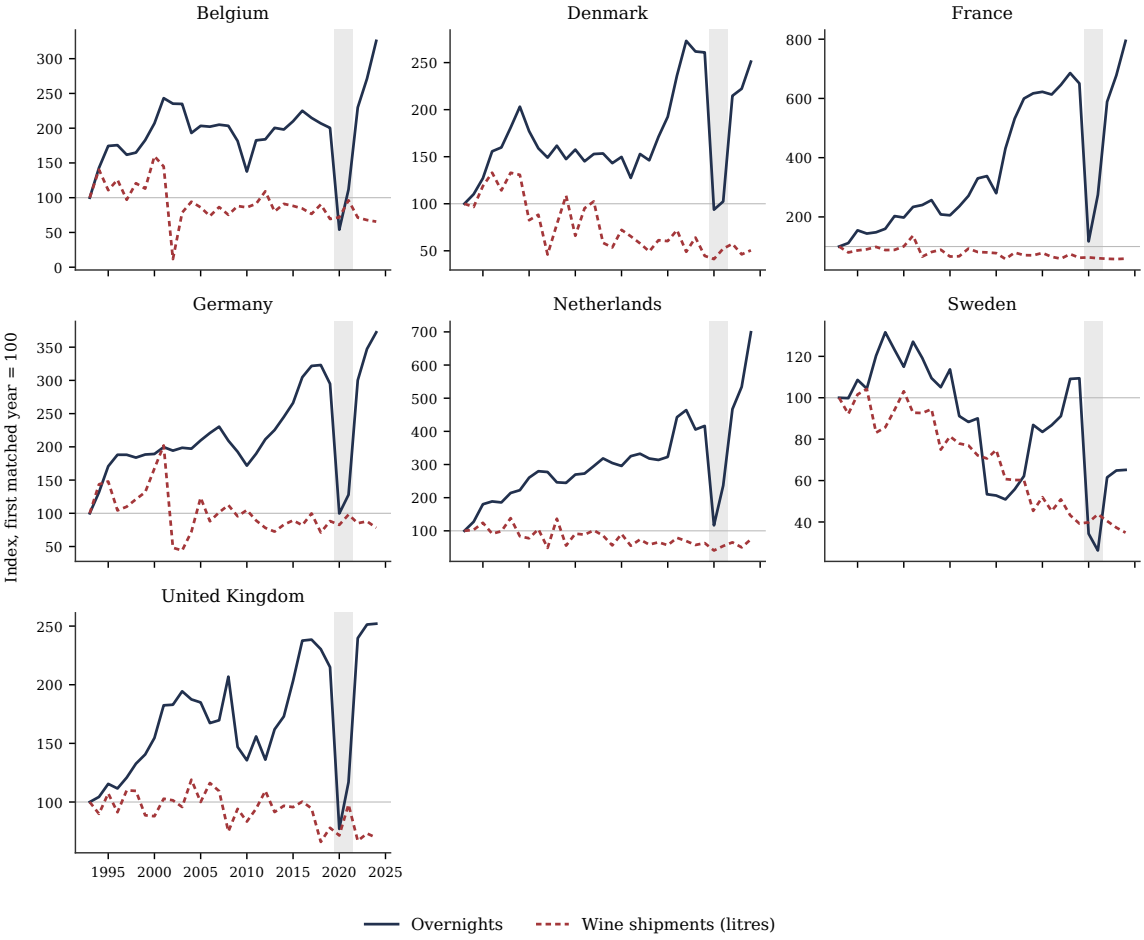


Figure 4: Tourist overnight stays and wine shipments in litres for the 7 continuously observed origin countries, 1993–2024, each series indexed to 100 in the country's first matched year and plotted on a single axis.

Notes: Co-movement would appear as the two lines rising and falling together. Instead tourism climbs well above its base in most countries while shipments stay near or fall below theirs. Pandemic years 2020–2021 are shaded. The divergence is the descriptive motivation for the dynamic estimates. Source: wine commercialisation census and DREM.

4.2 Dynamics

A returning-visitor mechanism predicts that visitor flows should precede shipments, so the association should load on current and lagged tourism, while future tourism should not predict current shipments. We estimate

$$\ln \text{Ship}_{it} = \alpha_i + \tau_t + \beta_0 \ln \text{Over}_{it} + \sum_{k=1}^{2} \beta_k \ln \text{Over}_{i,t-k} + \sum_{k=1}^{2} \delta_k \ln \text{Over}_{i,t+k} + \varepsilon_{it}, \qquad (1)$$

with country and year fixed effects and standard errors clustered by country, for shipments in litres and in real value, and with guests replacing overnights as the exposure measure in a variant, since exposure is people while overnights conflate arrivals with length of stay.

Annual data blur the mechanism's timing at the same-year horizon. A visitor who travels early in year t can purchase at home later in the same calendar year, so the admissible window for a returning-visitor response is the contemporaneous term together with the lags, not the lags alone. The same-year term is also where common origin-country demand shocks and importer order timing operate, so it cannot separate rapid post-visit purchasing from confounding on its own.

We therefore read the table through three objects, with the cumulative association defined as the current-plus-two-lag coefficient sum $\beta_0 + \beta_1 + \beta_2$ rather than coefficient by coefficient, which matters doubly because overnights are highly persistent and the individual terms of equation (1) are collinear by construction. The sum of the lag coefficients isolates the delayed component the promotional literature emphasises. The cumulative sum $\beta_0 + \beta_1 + \beta_2$ covers the full admissible window, and its confidence set bounds the recorded shipment association over that window.

Because tourism is persistent, conditioning on future tourism can absorb part of the very variation a genuine channel would use, so we estimate the cumulative window both with the leads included and in a no-lead model of current and lagged tourism, and we treat the no-lead interval as the headline bound. The leads, which the mechanism requires to be zero, are then tested as a falsification rather than serving only as controls. First-difference and five-year long-difference variants remove the persistence directly. Table 2 reports the results, with a one-lag value variant, in which the contemporaneous value coefficient is significant, printed in the auxiliary table (Online Resource 1, Table S1) so that the specification dependence is visible rather than buried.

Table 2 and the cumulative object of Table 3 carry the main reading, with the contemporaneous-coefficient battery in Table S1 in Online Resource 1, and the estimates read in two movements. In the full sample, every object in the admissible window is statistically indistinguishable from zero. Read against the framework, H2 passes in every sample, the persistent component of H3 appears nowhere, and H1's positive footprint emerges only in the pre-2018 window of Table 3.

In litres, the contemporaneous coefficient is 0.077 (wild $p = 0.50$), the summed lags are -0.030 (wild $p = 0.83$), the cumulative association over the admissible window is 0.107 ([-0.50, 0.37]) without leads and 0.046 ([-0.25, 0.20]) with them, and the joint test gives $p = 0.78$ analytically and 0.95 under the wild bootstrap. The estimate stays small and indistinguishable from zero excluding Sweden (0.095), excluding the United Kingdom (0.062), restricting to the balanced seven (-0.039), adding origin GDP and population controls (0.073), and allowing country-specific trends (0.140), and Table S2 in Online Resource 1 shows no single country drives the null, with the leave-one-out cumulative sum ranging only from -0.045 to 0.109. The unweighted coefficient is an average market-level projection, and the value-weighted variant printed in Table S1 in Online Resource 1 keeps the sector-wide reading visible.

Table 2: Dynamic association of shipments and tourism, symmetric model

	Litres	Value
ln overnights, t	0.077 (0.095)	0.095 (0.088)
wild bootstrap p	0.50	0.264
ln overnights, $t-1$	0.004 (0.121)	−0.022 (0.093)
ln overnights, $t-2$	−0.035 (0.092)	0.028 (0.107)
ln overnights, $t+1$	−0.095 (0.131)	−0.151 (0.109)
ln overnights, $t+2$	0.125 (0.123)	0.192 (0.103)
Sum of lags	−0.030	0.006
(p, wild p)	(0.83, 0.83)	(0.97, 0.97)
95% wild CI	[−0.42, 0.25]	[−0.38, 0.34]
Cumulative sum, t to $t-2$	0.046 (0.096)	0.101 (0.150)
wild p, 95% wild CI	0.62, [−0.25, 0.20]	0.62, [−0.35, 0.43]
Cumulative, no-lead model	0.107 (0.128)	0.141 (0.164)
wild p, 95% wild CI	0.61, [−0.50, 0.37]	0.63, [−0.62, 0.47]
Sum of leads	0.029	0.041
wild p	0.77	0.63
Joint test, all terms (p, wild p)	0.78, 0.95	0.38, 0.75
Observations	251	251
Country and year fixed effects	Yes	

Notes. Standard errors clustered by country in parentheses. Wild-cluster bootstrap p-values impose the null and enumerate all 2^{12} Rademacher sign patterns, which removes simulation error from the bootstrap distribution but does not make the test exact in finite samples. Sum, cumulative, and lead-sum wild p-values and confidence sets are obtained through exact reparametrisations, and the confidence sets invert the same enumerated test on a dense grid, each acceptance region verified to be a single interval. Joint tests report the analytic F p-value beside a wild bootstrap of the Wald statistic under the joint null. The no-lead cumulative rows use the current-plus-lags model on 275 observations, with leads reserved for the falsification rows. Under six-point Webb weights the headline wild p-values are 0.027 (one-lag value contemporaneous, Table S1 in Online Resource 1), 0.82 and 0.97 (lag sums), and 0.63 and 0.61 (cumulative sums). A returning-visitor mechanism predicts a positive cumulative association within the admissible same-year and lagged window (H1), zero leads (H2), and lagged mass only when purchasing or replenishment crosses calendar years (H3). Source: author's calculations on DREM tourism statistics and the wine commercialisation census.

In value the lag structure is equally dead, with summed lags of 0.006 (wild $p = 0.97$) and a cumulative window of 0.101 bounded by [−0.35, 0.43]. The one association that does appear is fragile and sits at the horizon confounding shares. The contemporaneous value coefficient is 0.170 with a wild p of 0.023 in the one-lag specification. Single-country exclusions move that wild p between 0.007 and 0.083 (Table S2 in Online Resource 1), the estimate attenuates to 0.095 (wild $p = 0.264$) once symmetric second-order terms enter, and it is 0.426 with a standard error of 0.320 outside the pandemic years, a larger point estimate that the shorter sample cannot pin down. Because the data are annual, same-year co-movement in value is compatible both with rapid post-visit purchasing and with common demand shocks in the origin country, importer order timing aligned with the travel season, or marketing cycles, and the aggregate data cannot separate these, which is why the cumulative bound above, not the timing of this coefficient, carries the

paper's conclusion.

Table 3: The cumulative shipment object across windows, samples, and controls

	Estimate (se)	wild p	95% wild CI	N
Full sample	0.107 (0.128)	0.61	[−0.50, 0.37]	275
Clean window, through 2017	0.112 (0.088)	0.39	[−0.66, 0.80]	191
Same rows, leads included	0.271 (0.102)	0.01	[0.11, 0.50]	191
Extended clean window (adds 2018–19, 2024)	0.135 (0.103)	0.45	[−0.55, 0.46]	227
Balanced seven	0.155 (0.129)	0.51	[−1.81, +∞)	210
GDP per capita and population	0.088 (0.114)	0.53	[−0.42, 0.33]	275
Country-specific trends	0.012 (0.104)	0.91	[−0.22, 0.24]	275
Bilateral real exchange rate	−0.050 (0.143)	0.79	[−0.47, 0.66]	168

Notes. The object is the cumulative same-year and lagged sum $\beta_0 + \beta_1 + \beta_2$ from the current-plus-lags model without leads, except the same-rows leads-included row, which reports the symmetric model on the identical observations. The clean window drops every observation whose own year, lags, or leads touch 2020 or 2021, which ends it in 2017, and the extended clean window drops only own-year and lag references, adding 2018, 2019, and 2024. Wild p-values and confidence sets invert the null-imposed enumerated wild-cluster bootstrap, with subsamples enumerating their own cluster counts. The real exchange rate is a CPI-based bilateral index against Portugal, available for a shorter window, so that row runs on fewer observations. The balanced-seven upper endpoint is not bounded by the seven-cluster enumerated bootstrap, a known weak-tail property of null-imposed wild inference with few clusters, and the interval is reported half open. Hypotheses H1 to H3 of Section 2.2 are read against this table together with the lead tests of Table 2. Source: author's calculations on DREM tourism statistics and the wine commercialisation census.

The first-difference row makes the pandemic's role explicit: the differenced contemporaneous coefficient of 0.184 sits at the boundary of significance in the full sample (wild $p = 0.05$) and falls to 0.072 (se 0.157) once every difference touching 2020 or 2021 is removed, so the short-run differenced co-movement is concentrated in the pandemic years and we do not read it as a mechanism.

These magnitudes sit at the low end of the tourism-trade literature and lack its timing. Cross-country studies report positive tourism-export elasticities, in some estimates close to one for differentiated and food products and with lagged responses (El-Sahli, 2018; Brau and Pinna, 2013), and the closest product study finds a long-run tourism-imports link for Spanish wine in Germany (Fischer and Gil-Alana, 2009). Our full-sample confidence sets exclude one at every horizon combination we test.

4.3 Cross-correlation profile

We then let the timing be fully flexible. For each country we correlate within-country demeaned overnights at year t with demeaned shipments at year $t + k$ for k from -3 to $+3$ and pool across countries. Figure 5 plots the profile with a naive reference band of $\pm 1.96/\sqrt{n_k}$, about ± 0.116 at one year, which understates the true uncertainty because within-country observations are dependent, and which we therefore treat as descriptive scaling rather than a test. If tourism led exports, the profile would peak at a positive k. It does not. The values at positive horizons are -0.15, -0.13, and -0.10, mildly negative throughout, and the largest of them in magnitude sits at exactly the one-year horizon where the mechanism predicts a positive peak, carrying the opposite sign. Computed without the pandemic years, the positive-horizon values are -0.21, -0.21, and -0.18, and

the profile again peaks at no positive horizon. Nothing in the profile suggests that visitor flows anticipate shipments at any horizon, and we read it as descriptive corroboration of the regression evidence rather than as inference in its own right.

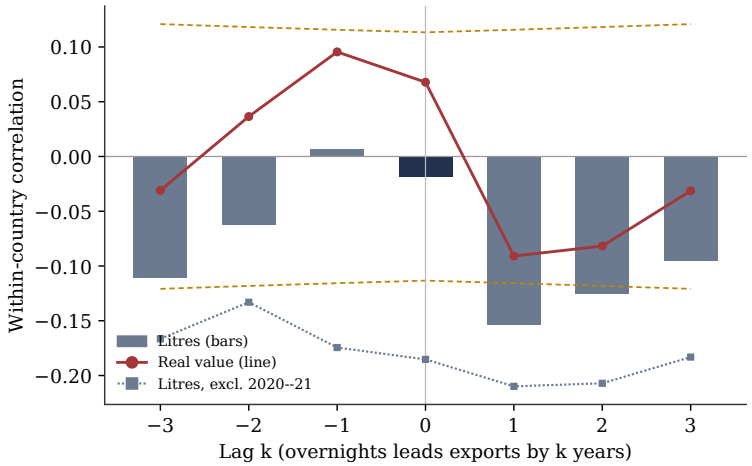


Figure 5: Pooled cross-correlation profile of within-country demeaned overnights and shipments.

Notes: Positive k means overnights lead shipments by k years. The correlation at $k = 0$ is -0.02 and at $k = +1$ is -0.15, against a naive ± 0.116 band at one year, drawn as a visual benchmark only, since within-country dependence widens it. The rust line shows the same profile for real shipment value, which also peaks at no positive horizon, and the dotted line shows the litres profile computed without the 2020–2021 observations.

4.4 Power, detection, and what the null can bound

A null is only as informative as the effects it can detect, and a detection limit is not a bound, so we report the two objects separately and let each do its own work. The minimum detectable elasticity at conventional power is about 0.27 at the contemporaneous horizon, 0.34 at the one-year lag, and 0.33 for the guest-based lag, ex-ante quantities that describe the design's resolution rather than what the data rule out. The bound itself comes from inverting the design's own inference. The 95 percent wild-cluster confidence set for the summed lag response in litres is $[-0.42, 0.25]$, the cumulative same-year and lagged window gives $[-0.50, 0.37]$ without leads and $[-0.25, 0.20]$ with them, and the value counterparts are $[-0.62, 0.47]$ and $[-0.35, 0.43]$, so the elasticities near one that the cross-country literature reports lie well outside what these data admit at any admissible horizon. Each acceptance region was inverted on a dense grid and verified to be a single interval. The intervals bound the recorded shipment response, not consumer conversion. Omitted origin-year demand shocks that move tourism and shipments together would push the association up, while measurement error in exposure, third-country routing, direct channels, and inventory buffering push the observable response down, so we do not convert the intervals into a bound on the underlying consumer channel.

Conversion units are a different matter. Suppose a fraction κ of a country's guests buys two bottles at home in the following year, 1.5 litres per convert. We add that channel to observed shipments country by country and year by year, re-estimate equation (1) exactly as specified, and ask at which κ the enumerated wild test rejects. The answer is sobering for anyone hoping the null pins down conversion. The cumulative test rejects only above roughly 46 percent of guests, the guest-based cumulative test rejects only above about 46

percent of guests, and the lag-sum test does not reject at any rate up to full conversion of guests.

By design, the exercise is a detectability calculation on the realised sample rather than a repeated-sample power statement, the $p(\kappa)$ profile is computed on a full grid and crosses the five percent line once, and it maintains one-for-one pass-through of purchases into recorded same-country shipments in the stated year, an assumption the routing and inventory frictions above make favourable to detection. Two features of the data produce this bluntness. Overnights are highly persistent, so an injected signal keyed to last year's guests spreads across adjacent tourism terms rather than concentrating where the test looks. And the guest-to-shipment ratio varies by a factor of twenty across countries, with an implied share at a κ of five percent running from 1.2 percent in France to 26.4 percent in Spain, so a uniform conversion rate produces cluster-heterogeneous effects that twelve-cluster inference prices in. Ratio shortcuts miss both features: dividing the lag detection limit by the median implied elasticity suggests a threshold of about 23 percent, and mapping the cumulative confidence limit linearly suggests about 11 percent, both far below the simulated threshold. The simulation is the design's best resolution in conversion units, and it says the data cannot bound modest conversion rates.

Table 4 reports the implied shipment shares behind these thresholds, computed on 2015–2019 means, with the cross-country dispersion printed because a median alone conceals it. The conversion rate κ is therefore a net shipment-equivalent rate among registered-accommodation arrivals rather than a repurchase rate among unique visitors, and because repeat registrations and uncovered stays pull the arrival count in opposite directions, we do not adjust the thresholds in either direction.

The conclusion is therefore asymmetric across units, and stating the asymmetry is the point. In elasticity units the bound is informative in every sample: the confidence sets place the cumulative association no higher than 0.37 in litres in the full sample and at 0.50 in the pre-2018 window, and each of these confidence sets excludes the literature's benchmark of one. In conversion units the design is blunt: rejection first occurs in the realised-sample injection above roughly 46 percent of guests.

Table 4 repeats the exercise with the purchase allocated to the same year, split evenly across the same and following years, and spread over the admissible three-year window, so the detection limit is read against the framework's own timing rather than a single channel. The abstract quotes the smallest, which the following-year allocation produces, while the same-year and three-year allocations generate no rejection at any conversion rate, so conversion concentrated in the visit year is effectively invisible to the design, so smaller, commercially meaningful conversion rates remain open, not because the data hint at them but because annual shipment data cannot see them.

We are unaware of a published study that reports a post-visit home-market repurchase rate for a region's international visitors that would pin down where in this range the truth lies. The nearest evidence is domestic and on-channel: a longitudinal Australian study finds that about two-thirds of cellar-door visitors buy on site and a minority repurchase within six months, but through domestic retail and direct-to-consumer channels (Bruwer et al., 2015), which are precisely the channels absent from destination-country shipments. The behavioural evidence and our bound therefore describe different links in the same chain.

Table 4: Calibration: visitor conversion, purchase timing, and detection

Conversion κ	Median share (%)	Mean share (%)	Range (%)	Implied elast.
1%	1.5	1.9		0.015
5%	7.5	9.5	1.2–26.4	0.075
10%	15.0	19.0		0.150
Detection thresholds for the following-year channel, by simulation				
Cumulative test	rejects above about 46 percent of guests			
Guest-based test	rejects only above about 46 percent of guests			
Lag-sum test	does not reject at any rate up to full conversion of guests			
Timing of the injected purchase, two bottles per convert				
Following year (baseline)	first rejects above about 46 percent of guests			
Same year	does not reject up to full conversion of guests			
Even split, same and following	first rejects above about 88 percent of guests			
Spread over three years	does not reject up to full conversion of guests			

Notes. A fraction κ of each country's guests is assumed to buy 1.5 litres at home in the following year. Implied shares divide the resulting addition by observed shipments, on 2015–2019 country means, with the median, mean, and five-percent range across the 12 matched countries reported, and shares scale linearly in κ. The implied elasticity column converts the median share into an elasticity contribution. Injection thresholds add the channel to observed litres, re-estimate equation (1), and apply the enumerated wild-cluster test. One bottle per convert (0.75 litres) moves the cumulative threshold to about 92 percent and four bottles (3.0 litres) to about 23 percent. A linear mapping of the cumulative confidence limit would suggest about 11 percent, far below the simulated threshold, which quantifies how much ratio shortcuts overstate the design's resolution once persistence and cross-country heterogeneity are priced in. Source: author's calculations on DREM tourism statistics and the wine commercialisation census. A naive per-guest ratio for the lag test gives about 23 percent; it appears only in this note because it overstates the design's resolution, as the text explains.

The logic is the point: with both series in hand, the tasting-to-purchase claim stops being unfalsifiable rhetoric and becomes an elasticity with a confidence set and a conversion rate with a measured detection limit.

4.5 The pandemic years

The sample ends in 2024 and therefore spans the collapse and recovery of travel. For the five countries entering in 2010, the 2020–2021 window is a large share of their within-country variance, and it moves overnights through border closures while shipments respond to entirely different constraints, so pandemic observations can dominate correlations in either direction. Dropping the 2020–2021 rows is not enough, because adjacent rows keep pandemic values inside their lag and lead terms, so the exclusion operates at the regressor level: an observation enters only if neither its own year nor any year referenced by its lags and leads falls in 2020–2021. On that basis the pooled within-country correlation is -0.19, the contemporaneous coefficient is 0.242 (0.246) with a one-year lag of -0.047 (0.224) on 227 observations, and the symmetric specification gives a summed lag response of 0.204 (wild $p = 0.31$) on 191 observations. The picture that emerges is not uniform, and we report it as it is. On the symmetric clean sample, which the lead requirement ends in 2017, the cumulative window is 0.271 (wild $p = 0.01$) with a confidence set of [0.11, 0.50], and the guest-based counterpart is 0.277 (wild $p = 0.01$, confidence set [0.08, 0.53]), a modest positive same-year-plus-lags co-movement the full

sample does not show. The no-lead clean sample requires no future years and therefore also includes 2018, 2019, and 2024, and there the cumulative window is 0.135 (wild $p = 0.45$, confidence set $[-0.55, 0.46]$) on 227 observations. Table 3 separates the model from the window exactly, by estimating the no-lead model on the identical pre-2018 rows (0.112, wild $p = 0.39$, confidence set $[-0.66, 0.80]$), alongside the balanced seven and specifications with origin GDP per capita and population, country-specific trends, and a bilateral real exchange-rate control. The decomposition lands on the model rather than the window: the headline no-lead object is null in every window we print, and the symmetric pre-2018 positive appears only once future tourism is partialled out, the conditioning Section 4.2 flags as absorbing mechanism-relevant variation. We therefore read it as a conditional association inside one specification cell, not as evidence of a footprint the headline object misses. Its composition disciplines its reading further. The delayed component stays indistinguishable from zero on the same clean sample, as the summed lag response above shows, so the significance comes from the same-year term together with individually insignificant lags, exactly the horizon Section 4.2 argued cannot separate rapid post-visit purchasing from common origin-year demand shocks and importer order timing. What survives every sample and window is the timing and the absence of any large footprint: the full-sample, window, controls, and trends confidence sets all lie below the elasticities near one in the cross-country literature, and no sample or specification produces the persistent delayed component of H3. Because the symmetric clean-sample cumulative test already rejects zero with no injected channel, the conversion detectability thresholds of Section 4.4 are full-sample objects. The cross-correlation profile recomputed without the pandemic years peaks at no positive horizon, as Section 4.3 reports. The paper's bound and the absence of the delayed component do not rest on any window choice. The behaviour of the cumulative co-movement does, null in the full sample, positive in the pre-2018 symmetric window, absent from the no-lead windows, and we state that dependence rather than bury it.

5 Discussion

Across co-movement, symmetric dynamics, and a flexible cross-correlation profile, the persistent delayed component a durable returning-customer channel would produce is not present in the covered European markets within the stated detection limits, and the associations that do appear, a contemporaneous one in shipment value and a positive cumulative window in the pre-2018 sample, are confined to the horizon most exposed to confounding, with the delayed component null in every sample. Three readings discipline what follows from this.

First, against the trade literature. Instrumented cross-country studies find tourism raising exports of differentiated and processed food products with elasticities near one and with lagged effects (El-Sahli, 2018; Brau and Pinna, 2013), and the closest product study finds a long-run tourism-imports link for Spanish wine in Germany (Fischer and Gil-Alana, 2009). Our bounded null for a flagship protected-origin wine, in an administrative census of the certified sector, says those broad-class findings do not license product-level promotional claims: whatever aggregates to a positive cross-country elasticity near one, nothing that large is visible for this product in these markets at this frequency, the delayed component the promotional story names is absent outright, and marketing arithmetic for a specific heritage good cannot borrow the cross-country number.

Second, against the measurement gap. Our outcome is shipments to distributors, not purchases by returning visitors, and importers hold inventory and reorder on their own schedules. Smoothing of that kind spreads a demand shift across adjacent years, which is why the cumulative window and the five-year long difference carry weight, and the conclusion is a bound at low frequency rather than an acquittal at all frequencies: the conditional annual shipment association is bounded by the confidence sets of Section 4.4, in the covered markets, at annual horizons up to three years, a small country-level footprint can reflect routing through hubs, direct channels, and travel retail rather than absent demand, and bounding consumer conversion itself would require a formal mapping from purchases to recorded shipments that these data cannot support.

Third, the separation of channels. The point-of-experience channel and the home-market shipment channel are different economic objects. This paper measures only the second, and the separation stands on its own framework: strong purchasing where the product is tasted is fully compatible with the bounded footprint here, because the links from formed preference to a recorded destination shipment, availability after return, distribution, direct channels, and travel retail, are exactly where the leakage sits. A destination can succeed at the first channel and still show nothing in the second.

5.1 Policy implications: from visitor exposure to market conversion

What the evidence changes is the status of visitor counts. At the upper endpoint of the primary full-sample confidence set (0.37), a sustained ten percent rise in overnights corresponds to roughly a four percent change in recorded shipments within the estimated conditional association, and the point estimates are small and statistically indistinguishable from zero. Because omitted origin-market shocks, measurement error, routing, and inventory stand between this projection and any causal multiplier, the number describes the recorded association, not a policy elasticity. The operational content is narrower and firmer: arrivals and overnights are exposure measures, and this paper finds no basis for reading their growth as demonstrated export conversion.

What the framework suggests is conditional, and it is complementarity. Destination experience and foreign-market access are distinct inputs into any post-trip sale, controlled by different and overlapping institutions, destination bodies and product boards on one side, exporters, distributors, and trade-promotion organisations on the other, in Portugal the remit of AICEP. The framework motivates a testable policy hypothesis rather than a recommendation: where export spillovers are an explicit objective, exposure and market access should be designed and evaluated jointly, and a pilot in a few high-visitor origin markets with weak current availability, combining importer engagement, listings, and post-visit communication against comparable control markets, would test the complementarity directly.

What should be measured next follows from the chain itself. Each link has its own observable, tasting participation and route use at the exposure stage, availability and listings in origin markets at the access stage, buyer-country retail or direct sales at the conversion stage, and importer orders at the replenishment stage, which is the only stage this paper's outcome captures, and imperfectly. The paper therefore informs programme architecture and accountability, not optimal public expenditure, and it cannot say which link binds or what any intervention returns. The conclusion lists the data that would let each stage be tracked, and collecting them is within the reach of the same institutions.

6 Conclusion

This paper asked whether tourism acts as export promotion for a place-based wine product through a delayed home-market shipment channel, matching tourist overnights by country of residence with Madeira wine shipments to the same country over three decades and testing the footprint the mechanism predicts.

The answer is a bounded null for the delayed channel, and for the full window a bounded co-movement whose significance depends on the sample window, with the bounded object the recorded shipment response. The persistent delayed component a durable returning-customer channel would produce is undetected in both litres and value and in every sample: summed lags are indistinguishable from zero, the cumulative same-year and lagged window is bounded by a 95 percent confidence set of $[-0.50, 0.37]$ in litres in the full sample, with the remaining window and control intervals reported in Table 3, full-sample joint tests of all tourism terms cannot reject zero, and the cross-correlation profile peaks at no positive horizon. Two positive associations appear and both sit at the horizon most exposed to confounding: the contemporaneous value coefficient, fragile across specifications, and the pre-2018 cumulative window, whose delayed component is itself null and which the no-lead windows do not reproduce. At annual frequency that horizon mixes rapid post-visit purchasing with common demand shocks and importer order timing, so both enter the cumulative bound rather than being dismissed or credited on timing alone.

The contribution is threefold. Firstly, a mechanism that shows up in broad product classes leaves no stable full-sample footprint and no persistent delayed component for a flagship heritage product in the very data where promoters would look for it. Secondly, the paper supplies an administrative, country-matched, dynamic test of a claim usually supported by stated intentions, together with a bound built from the design's own inference: the confidence sets place the cumulative shipment association no higher than 0.37 in litres and exclude the elasticities near one that the cross-country literature reports, while an injection exercise shows that conversion below roughly 46 percent of guests escapes the design's own test under its most detectable purchase-timing calibration, so the shipment-association bound rather than a conversion rate is the paper's summary object and smaller, commercially meaningful conversion remains open. Thirdly, the paper leaves a usable framework and an honest map of what annual administrative data can resolve for the tasting-to-purchase claim: a four-link chain whose leaks are priced by a measured detection limit, a cumulative footprint whose significance is window dependent and reported as such, and a persistent delayed component that is undetected everywhere.

The policy reading is deliberately narrow. The paper does not evaluate the on-site, experiential, or destination-branding returns to wine tourism, and nothing here argues against tastings, routes, or cellar-door investment on those grounds. It does withdraw one specific justification: visitor growth alone is not evidence that a home-market export link has operated, at least not in recorded shipments at annual frequency. Programmes that seek that link should provide market access and evaluate it separately, the hypothesis the Discussion states in testable form.

The bound carries its scope conditions honestly. It holds at annual frequency, for the covered European markets, on shipments to distributors rather than final purchases, and it is silent on the United States and Switzerland, which have no tourism-residence coverage and are among the destinations where conversion is most plausible. The framework travels where its ingredients recur, a tourism-specialised or peripheral region, a recognisable

place-based tradable that visitors encounter, and shipments identifiable by destination market, a class that includes wines, spirits, geographical-indication foods, and other heritage goods, and what travels is the framework and the measurement problem rather than the coefficient. Closing these gaps requires data the public series do not yet provide, and collecting them is itself within the reach of an export promotion agency or a regional institute: channel-level and direct-to-consumer sales by buyer country, travel-retail and cellar-door volumes, scanner evidence in origin markets linkable to prior travel, and monthly rather than annual frequency. Until then, the honest summary is that tourism may sell the product where the visitor stands, but the delayed home-market shipment channel is not detectable in annual administrative shipment data.

References

Karin Alant and Johan Bruwer. Wine tourism behaviour in the context of a motivational framework for wine regions and cellar doors. *Journal of Wine Research*, 15(1):27–37, 2004.

Rinaldo Brau and Anna Maria Pinna. Movements of people for movements of goods? *The World Economy*, 36(10):1318–1332, 2013. doi: 10.1111/twec.12104.

Juan Gabriel Brida, Isabel Cortés-Jiménez, and Manuela Pulina. Has the tourism-led growth hypothesis been validated? A literature review. *Current Issues in Tourism*, 19 (5):394–430, 2016.

Johan Bruwer and Karin Alant. The hedonic nature of wine tourism consumption: an experiential view. *International Journal of Wine Business Research*, 21(3):235–257, 2009.

Johan Bruwer, Larry Lockshin, Anthony Saliba, and Martin Hirche. Trial-purchase-repurchase of the brand: how does a cellar door visit impact future sales? *Wine and Viticulture Journal*, 30(1):60–63, 2015.

Olivier Cadot, Ana M. Fernandes, Julien Gourdon, and Aaditya Mattoo. Are the benefits of export support durable? Evidence from Tunisia. *Journal of International Economics*, 97(2):310–324, 2015.

Steve Charters and Jane Ali-Knight. Who is the wine tourist? *Tourism Management*, 23(3):311–319, 2002.

Erica Croce and Giovanni Perri. *Food and Wine Tourism: Integrating Food, Travel and Territory*. CABI, Wallingford, 2010.

Zouheir El-Sahli. The role of inbound tourist flows in promoting exports. *The World Economy*, 41(5):1457–1475, 2018. doi: 10.1111/twec.12597.

Christian Fischer and Luis A. Gil-Alana. The nature of the relationship between international tourism and international trade: the case of German imports of Spanish wine. *Applied Economics*, 41(11):1345–1359, 2009. doi: 10.1080/00036840601019349.

Donald Getz and Graham Brown. Critical success factors for wine tourism regions: a demand analysis. *Tourism Management*, 27(1):146–158, 2006. doi: 10.1016/j.tourman.2004.08.002.

Mar Gómez, Marlene A. Pratt, and Arturo Molina. Wine tourism research: a systematic review of 20 vintages from 1995 to 2014. *Current Issues in Tourism*, 22(18):2211–2249, 2019.

C. Michael Hall, Liz Sharples, Richard Mitchell, Niki Macionis, and Brock Cambourne. *Food Tourism Around the World: Development, Management and Markets*. Butterworth-Heinemann, Oxford, 2003.

Keith Head and John Ries. Do trade missions increase trade? *Canadian Journal of Economics*, 43(3):754–775, 2010.

Nada Kulendran and Kenneth Wilson. Is there a relationship between international trade and international travel? *Applied Economics*, 32(8):1001–1009, 2000. doi: 10.1080/000368400322057.

Mukesh Kumar, Sanjeev Prashar, and R. K. Jana. Does international tourism spur international trade and output? Evidence from wavelet analysis. *Tourism Economics*, 25(1):22–33, 2019.

Richard Mitchell and C. Michael Hall. Wine tourism research: the state of play. *Tourism Review International*, 9(4):307–332, 2006.

María Santana-Gallego, Francisco J. Ledesma-Rodríguez, and Jorge V. Pérez-Rodríguez. International trade and tourism flows: an extension of the gravity model. *Economic Modelling*, 52:1026–1033, 2016. doi: 10.1016/j.econmod.2015.10.043.

David J. Telfer and Geoffrey Wall. Linkages between tourism and food production. *Annals of Tourism Research*, 23(3):635–653, 1996. doi: 10.1016/0160-7383(95)00087-9.

Rebecca Torres. Linkages between tourism and agriculture in Mexico. *Annals of Tourism Research*, 30(3):546–566, 2003.

Johannes Van Biesebroeck, Jozef Konings, and Christian Volpe Martincus. Did export promotion help firms weather the crisis? *Economic Policy*, 31(88):653–702, 2016.

Online Resource 1

Online Resource 1 (ESM_1.pdf) contains the auxiliary and robustness specifications (Table S1), the leave-one-country-out estimates (Table S2), the guest-based symmetric dynamics (Table S3), and the per-country cross-correlation profiles (Figure S1).

Data availability

All data and code underlying this paper are open. The wine series come from the public administrative commercialisation statistics compiled by DREM and IVBAM, covering volume and value by destination market. The tourism series are the retrospective statistics of DREM, with residence-country breakdowns.

The replication package contains: the compiled census panel and all the estimation pipeline. Running the pipeline regenerates every number, table, and figure in the paper, asserting validation anchors and continuity guards before estimation. A clean-directory rebuild reproduces the manuscript's results in full.

The complete replication package is available to the editorial office and referees during review and will be archived in a public repository under a persistent identifier upon acceptance, with no access restrictions. No on-request or restricted material is used anywhere in this paper.