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.

Local agglomeration and hydrogen investment decisions: evidence from a global project panel

Georges Ali Benyahia

ali.benyahia@um6p.ma

Université Mohammed VI Polytechnique

Tareq Ait Hsain

Université Mohammed VI Polytechnique

Research Article

Keywords: Agglomeration economies, Investment timing, Hydrogen, Discrete-time hazard, Equivalence testing, Spatial units

Posted Date: November 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10987057/v1

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

Additional Declarations: No competing interests reported.

Local agglomeration and hydrogen investment decisions: evidence from a global project panel

Abstract

Announced hydrogen capacity greatly exceeds the capacity that has reached a final investment decision, and European policy promotes local clusters, hydrogen valleys, that co-locate production, demand and infrastructure. Whether producers already committed pull other producers across the decision margin is a further hypothesis, which that policy does not state. We test it: are hydrogen projects already committed or operating within 50 km associated with the transition of a neighbouring pre-decision project to a final investment decision? Six vintages of the International Energy Agency hydrogen project database, harmonized across a changing status vocabulary and aggregated into documented investment units, yield a risk panel of 3,211 unit-period observations and 128 transitions between 2023 and 2026. A discrete-time hazard model with country and vintage fixed effects gives an association of -0.04 percentage points per standard deviation of local commitment, against a mean transition probability of 4.0%, with a 90% interval of $[-0.83, +0.74]$. The estimate is statistically equivalent to zero within a 1.08-point benchmark, under the primary specification and not under country-by-vintage effects, where the association moves to $+0.51$ points. The panel records status changes between vintages, not decision dates, and the association is not causal.

Keywords Agglomeration economies · Investment timing · Hydrogen · Discrete-time hazard · Equivalence testing · Spatial units

JEL Classification R12 · Q42 · C81

1 Introduction

Announced hydrogen capacity worldwide is several times the capacity that has reached a final investment decision, and the gap has widened (Odenweller and Ueckerdt 2025). The binding margin is not whether a project is announced but whether it is sanctioned: an announcement commits no capital, while a final investment decision converts an option to invest into an irreversible outlay (Arrow and Fisher 1974; McDonald and Siegel 1986; Dixit and Pindyck 1994; Kellogg 2014). Public programs respond in large part by concentrating projects in space. The European Union's strategy foresees local clusters it calls hydrogen valleys, described there as relying on local renewable production and local demand carried over short distances (European Commission 2020, sec. 2): what it envisages integrating is production, demand and infrastructure in one place. Whether producers already committed pull other producers across the decision margin is a further hypothesis, which the strategy does not state and which agglomeration theory suggests. It is the hypothesis this paper tests.

Regional economics supplies reasons why that hypothesis could hold, and reasons why it could fail. Neighbouring commitments can lower the cost of the marginal project through shared infrastructure and thicker markets for specialised suppliers, and lower its perceived risk by demonstrating that permitting, offtake and financing are attainable at that location (Marshall 1890; Duranton and Puga 2004; Ellison et al. 2010; Greenstone et al. 2010). They can also raise it, since projects at one location draw on the same grid connection, land and permits, and a firm holding an option to wait may treat a neighbour's commitment as a reason to observe its outcome before spending. Section 2 sets out both families; an observational panel does not separate them.

We therefore ask an empirical question with a definite sign to test: is the presence of hydrogen projects already committed or operating within 50 km associated with the transition of a neighbouring pre-decision project to a final investment decision? The economic object is the date of a decision. What the source records is a change of status between two

snapshots, so the estimand is the probability that a unit at risk in one vintage is recorded as committed in the next, and no decision date is observed. We estimate an association and claim no causal effect.

We use six vintages of the International Energy Agency hydrogen production project database, 2021 to 2026. Two features of the source must be handled first: the status vocabulary changes between vintages, and a single site frequently occupies several rows, since phases, trains and partner-specific entries are recorded separately. We harmonize the vocabulary and aggregate rows into investment units, then estimate a discrete-time hazard model on 3,211 unit-period observations and 128 transitions between 2023 and 2026. The association is -0.04 percentage points per standard deviation of local commitment, against a mean transition probability of 4.0%, with a 90% interval of $[-0.83, +0.74]$. It is statistically equivalent to zero within a benchmark of 1.08 points calibrated on the gap between committed and achievable 2030 capacity. Equivalence does not hold under country-by-vintage effects, where the estimate moves to $+0.51$ points at unchanged precision.

The paper contributes on three levels. It estimates the association between prior local commitment and the transition of hydrogen projects to a final investment decision, on a global panel and against an explicit equivalence benchmark, so that the absence of a detectable association is a bounded statement and not a failure to reject. It reports how far that estimate is precise and how far it is sensitive, to the fixed-effects scheme, to the definition of the investment unit and to eleven alternative exposure measures. And it supplies a documented construction of investment units from a project database, with the crosswalk, matching log and unit assignment that any spatial analysis of such a source requires. Sect. 4.3 shows that a spatial statistic computed on rows at the project level is partly a statistic about recording.

The remainder of the paper is organized as follows. Section 2 sets out the mechanisms through which neighbouring commitments could move the decision margin, the scale at which they have been measured and the limits of identification. Section 3 presents the data, the

construction of investment units and the empirical strategy. Sections 4 and 5 report the results and the final comments.

2 Local commitment and the decision to invest

2.1 Channels of acceleration and delay

Three channels could bring a decision forward. The first is cost: a committed neighbour justifies facilities the next project does not have to fund alone, from grid reinforcement and water treatment to storage, port and pipeline capacity. Sharing is the classical source of localisation economies (Marshall 1890; Duranton and Puga 2004; Rosenthal and Strange 2004), and coagglomeration patterns and plant-opening episodes supply the leading evidence that it operates (Ellison et al. 2010; Greenstone et al. 2010). The second is the market for specialised inputs. Electrolyser installation, high-pressure engineering and hydrogen-specific maintenance are thin markets, and a local pipeline of sanctioned work makes it worthwhile for suppliers and skilled labour to establish nearby, shortening the lead time a developer must price.

The third channel is informational, and the timing margin makes it salient. A commitment is taken under uncertainty about permitting, about offtake at a price that covers cost, and about technology at scale, and a neighbour that commits resolves part of that uncertainty for a specific location: it demonstrates that the regulator grants, that a buyer signs and that lenders accept the risk on those terms. In an option-to-wait framework, a fall in the variance of the payoff lowers the threshold at which committing dominates waiting (Arrow and Fisher 1974; McDonald and Siegel 1986; Dixit and Pindyck 1994; Kellogg 2014). Neighbour effects on discrete decisions are documented in this journal for industrial location (Alamá-Sabater et al. 2011) and for the adoption of a conservation practice (Gedikoglu et al. 2023).

Two channels work the other way. Projects at one location compete for scarce local

complements, above all a grid connection, but also land, water rights, skilled crews and the administrative capacity that issues permits, so a neighbour that commits first takes the cheapest slot and the next project faces a worse one. The informational channel also has a second edge: if a neighbour's commitment is expected to reveal something material soon, the option value of waiting rises, and the response to a committed neighbour is to postpone.

These are interpretations, not identified quantities: the estimate below is a net association and cannot attribute it, or its absence, to any one channel. The design tests a directional hypothesis, that the association is positive if the accelerating channels dominate at the scale the data resolve. A positive association is not a necessary condition for agglomeration economies to exist here, since the mechanisms may operate at a finer scale, offset one another, or be masked by the selection of projects into the panel.

2.2 Scale, unit of observation, and identification

The scale at which these mechanisms operate for hydrogen is not known. The nearest evidence measures other outcomes: Larsson (2014) obtains a wage-density elasticity that falls from 0.0122 for squares of 0.0625 km^2 to 0.0048 for 10 km^2 on Swedish geocoded data, and frames the ambiguity in unit size explicitly as a modifiable areal unit problem, while Almeida et al. (2026) find the effect of college-educated workers within one kilometre several times the effect in wider rings. These are wage and human-capital gradients, and they do not establish that the complementarities specific to hydrogen, which run through grid connections, pipelines, storage and offtake, decay over the same distances. We treat the relevant scale as uncertain and report the estimate at two radii, 50 and 100 km, without claiming that either is the scale at which the mechanism operates.

The unit of observation is a separate matter, and it is what separates decisions. A single site occupies several rows, so a neighbour pair 200 metres apart may be one decision recorded twice, and widening the radius does not remove such pairs, since a 50 km neighbourhood contains every pair below 500 metres as well. Separation has to come from the aggregation

rule applied before any distance is computed. We fix the unit first, and we report what the choice of unit changes, as regional science reports estimates against alternative aggregations (Díaz Dapena et al. 2016; Békés and Harasztosi 2018). The case here differs from that literature in that the ambiguity concerns the identity of a point and not the boundary of a zone, so the alternatives are not indexed by a scale parameter the analyst can move, and the distance-based measures built to escape given areal units (Kosfeld et al. 2011; Lang et al. 2020; Contreras 2022) do not escape this one, because they take the point pattern as given.

We claim no causal effect. Spatial correlation among projects can reflect conditions shared by neighbours, and the reflection problem (Manski 1993), the fragility of spatial econometric identification (Gibbons and Overman 2012) and the difficulty of interpreting peer coefficients (Angrist 2014) all apply. We do not estimate a model with an endogenous spatial lag. Andersson and Gråsjö (2009) show that accessibility variables on the right-hand side capture dependence a spatial lag leaves unmodelled, and our exposure measures are accessibility variables in that sense; Online Resource 1 reports why the impact decomposition such a model requires (Lacombe and LeSage 2018) could not be estimated credibly here.

3 Data and empirical strategy

3.1 Vintages, status harmonization, and coverage

We use six vintages of the International Energy Agency hydrogen production project database, dated 2021, 2022, 2023, 2024, September 2025 and June 2026, published under a CC BY 4.0 licence (IEA 2026b).[1] Each vintage is a snapshot of the world pipeline carrying, for every project, a status, a technology, a product, an announced capacity, source references and, from September 2025 onward, coordinates. Arbia et al. (2016) show by simulation that

[1]During the preparation of this work the authors used Claude Opus 4.8 to organize and reconcile the successive vintages of the project database, and to check consistency between the manuscript, the tables and the replication code. It was not used to generate results or text presented as findings. The authors reviewed and edited all output and take full responsibility for the content.

missing observations and positional error degrade both spatial estimates and the tests based on them, and recommend that any empirical analysis document them. Three corrections are required, each quantified here or in Online Resource 1.

Vintages are linked by the project reference the database carries and not by name, through a four-step hierarchy set out in Online Resource 1. Matching on the normalized name alone would declare 91 departures from the geolocated pre-decision cohort against 31 under the hierarchy, and treating those 60 as failures would manufacture a large share of the outcome.

The status vocabulary is not constant. The 2021 and 2022 vintages distinguish *FID* from *under construction*, which are merged from 2023 into a single *FID/construction* category. Taking the raw labels at face value records a relabeling as an investment decision and so manufactures 133 of 298 apparent events, together with a spurious peak of 24.5% in 2022 that has no counterpart in any other vintage.

Coordinates appear only in the two most recent vintages and are propagated backwards. Of the 2,254 projects geolocated in both, 2,224 have coordinates identical to the metre and 30 move; excluding all 30 leaves the conclusion unchanged. Propagation has a consequence for scope that matters economically: the panel is restricted to projects surviving into the recent vintages, so a project announced in 2021 and abandoned before 2025 is absent, and the estimate speaks to the maturation of projects still alive late in the window. Coverage is 54% to 89% by vintage, and we take 2023 to 2026, where it exceeds 72%, as the main window.

3.2 Investment units

A single site frequently appears as several rows, so the row is not the decision. We aggregate rows into investment units before computing any distance. The strict definition requires both proximity below 0.5 km and documentary or naming evidence of a single decision, namely a shared source reference, an identical base name once phase markers are stripped, or a near-identical name. Of the 1,051 co-located pairs, 472 satisfy this test, 390 of them

through a shared source, so the rule keeps 55% of co-located pairs apart. Units are the connected components of the retained pairs, none of which chains into a group wider than the threshold.

Neither condition suffices alone, and the data supply cases in both directions: three rows at Eemshaven in the Netherlands carry two different project names at identical coordinates and one status throughout, so a name-based rule would count three decisions; two rows in Abu Dhabi share coordinates because both carry a city-level geocode while describing unrelated projects at different stages, so proximity alone would merge two decisions into one. Both are set out in Online Resource 1.

We assess the rule on observed behaviour. Among co-located pairs, those the rule retains carry an identical status in 47.6% of the vintages in which both are observed, against 37.2% for the pairs it keeps apart. Conditioning on movement gives the sharper comparison: among intervals in which at least one row changes status, both change within the same interval in 35.7% of cases for retained pairs against 11.1% for pairs kept apart, the measure requiring no more than that each changes, and not that the two transitions coincide. These diagnostics support the grouping. They do not establish that a retained pair records one economic decision, which is why the definition is carried through the estimation as one of six: four purely geographic alternatives, at identical coordinates and at 0.5, 2 and 10 km, and one that drops the shared-source rule. The unconditional concordance measures and the full diagnostic tables are in Online Resource 1.

3.3 Risk set, event, and exposure

We organize the investment-unit panel as a unit-period risk set and estimate a discrete-time transition model (Allison 1982; Jenkins 1995); the database is a sequence of snapshots and the outcome is a status revision recorded between two of them. Let $s_{it} \in \{\text{PRE}, \text{COM}, \text{OP}\}$ denote the state of unit i at vintage t, for pre-decision, committed but not operating, and operating, a unit taking the most advanced state among its projects. Unit i belongs to the

risk set $\mathcal{R}_t$ when $s_{it} = \text{PRE}$, and the event is the transition out of that state,

$$y_{i,t\to t+1} = \mathbf{1}\{s_{i,t+1} \in \{\text{COM}, \text{OP}\}\}, \qquad i \in \mathcal{R}_t, \tag{1}$$

so the object of estimation is the hazard $h_{it} = \Pr(y_{i,t\to t+1} = 1 \mid i \in \mathcal{R}_t)$. A unit that cannot be matched in a later vintage is censored, never treated as a failure, and transitions to cancellation, dormancy or on-hold status remain non-events in this binary advancement outcome.

Exposure is built from the units other than i within radius r, with d_{ij} the great-circle distance. Writing $\mathcal{N}_{irt} = \{j \neq i : d_{ij} \leq r\}$, the count and capacity measures at vintage t are

$$N^k_{irt} = \log\Big(1 + \textstyle\sum_{j\in\mathcal{N}_{irt}} \mathbf{1}\{s_{jt} \in k\}\Big), \qquad K^k_{irt} = \log\Big(1 + \textstyle\sum_{j\in\mathcal{N}_{irt}} q_{jt}\,\mathbf{1}\{s_{jt} \in k\}\Big), \tag{2}$$

where q_{jt} is announced capacity and k ranges over the operating, the committed, and the committed or operating states for the counts, but only over the operating and the committed or operating states for the capacities. With the share of neighbours committed or operating, this gives six measures at each of $r \in \{50, 100\}$ km, hence twelve, each evaluated at t, before the transition it explains. The radii bracket the scale at which the mechanism might operate, which Sect. 2.2 treats as uncertain; the aggregation rule, and not the radius, separates decisions, and Sect. 4.4 takes up the cost of the choice.

3.4 Specification, inference, and equivalence

The primary estimand is the coefficient on the log count of committed or operating units within 50 km, under the strict investment unit, in the 2023 to 2026 window; the other eleven measures are secondary and form a single multiple-testing problem. We estimate

$$y_{i,t\to t+1} = \beta E_{irt} + X'_{it}\gamma + \alpha_t + \mu_c + \varepsilon_{it} \tag{3}$$

by linear probability, where E_{irt} is one of the measures of (2), α_t are vintage fixed effects and μ_c fixed effects for all 90 countries. X_{it} contains unit capacity, technology and product shares, siting and resource proxies, isolation, the number of projects in the unit and the age of the oldest announcement, listed with the estimation sample in Online Resource 1.

The default variance estimator is the spatio-temporal analogue of Conley (1999), with a 200 km distance cutoff and a one-vintage lag cutoff; Online Resource 1 gives the expression. We report alongside it standard errors clustered on the investment unit, two-way clustering (Cameron et al. 2011), homoskedastic standard errors, and a Firth penalized logit (Firth 1993; King and Zeng 2001; Heinze and Schemper 2002).

Absence of significance does not establish that an association is small; the remedy is to state a magnitude the design would have detected and to test against it (Rainey 2014; Hartman and Hidalgo 2018). Two equivalence margins are used, and both are interpretive rather than preregistered. The first is calibrated on the committed-to-achievable 2030 capacity gap (IEA 2026a): with a baseline hazard of 3.99% and four annual transitions remaining, closing half the gap requires the hazard to rise by 1.08 percentage points per transition, a derivation set out with its two assumptions in Online Resource 1. The second is set conventionally at a quarter of the baseline hazard, 1.00 percentage points. Equivalence is assessed by two one-sided tests, and we report the smallest margin at which they reject at the 5% level, a property of the estimate that depends on neither benchmark. The family of twelve is corrected by Holm-adjusted p values (Anderson 2008) and a Romano-Wolf stepdown with a joint test (Romano and Wolf 2005, 2016), both from a bootstrap resampling 140 k-means spatial blocks whose median diameter exceeds the kernel cutoff (Bester et al. 2011; Colella et al. 2023; Lee and Shaikh 2014).

4 Results

4.1 Prior local commitment and the decision to invest

On the investment-unit panel, the coefficient on the log count of committed or operating units within 50 km is -0.0008 with a Conley spatio-temporal standard error of 0.0087 (Table 1 and Fig. 1). A one standard deviation increase in exposure is associated with a change of -0.04 percentage points in the probability of reaching a final investment decision between two vintages, against a mean transition probability of 4.0%, with a 90% interval of $[-0.83, +0.74]$. The estimate is close to zero and the interval is narrow relative to the magnitudes at stake: it excludes positive associations of one percentage point or larger per standard deviation of local commitment, which is the order of magnitude that would matter for the 2030 pipeline. The smallest margin at which the two one-sided tests reject non-equivalence at the 5% level is 0.83 percentage points, below both benchmarks, which reject at $p = 0.015$ and $p = 0.023$.

Read economically, prior local commitment does not measurably bring a neighbouring hydrogen project across the decision margin at the scale of 50 km. The accelerating channels of Sect. 2.1, if they operate at this scale, do not dominate by an amount the design would have detected.

4.2 Sensitivity of the estimate

The equivalence statement is established on one specification, and one alternative overturns it. Replacing country and vintage effects by 244 country-by-vintage effects raises the coefficient to $+0.0094$ (0.0082), a per-standard-deviation association of $+0.51$ percentage points with a 95% interval of $[-0.36, +1.39]$. The coefficient remains statistically indistinguishable from zero at $p = 0.252$, but the interval no longer lies inside the 1.08-point benchmark and the smallest margin at which equivalence holds rises to 1.25 points. What changes is the location of the estimate and not its precision: the standard error is 0.0082 against 0.0087 in the primary specification, so the interval is no wider. Identification then comes only from

Table 1 Primary estimand and equivalence tests

Primary estimand: log committed or operating units within 50 km	
Coefficient (Conley spatio-temporal s.e.)	-0.0008 (0.0087)
Effect of one standard deviation (p.p.)	-0.044
90% confidence interval (p.p.)	[-0.827, 0.739]
95% confidence interval (p.p.)	[-0.977, 0.889]
Baseline hazard	0.040
Equivalence	
Smallest margin at which equivalence holds (p.p.)	0.827
Margin from the IEA ambition gap (p.p.)	1.076
two one-sided tests, p	0.0151
Margin at 25% of the baseline hazard (p.p.)	0.997
two one-sided tests, p	0.0226
Family of twelve exposure estimates	
Joint test of the twelve coefficients, spatial-block bootstrap p	0.608
Smallest Holm-adjusted p in the family	0.774
Observations	3211
Events	128
Investment units	1428
Countries	90

Estimates of Eq. 3 on the strict investment-unit risk panel, 2023 to 2026, with 3,211 observations and 128 events, the full control vector, vintage fixed effects and fixed effects for 90 countries. Standard errors are the spatio-temporal Conley estimator of Sect. 3.4 unless otherwise stated. The two benchmarks are interpretive, not preregistered. Source: prepared by the authors.

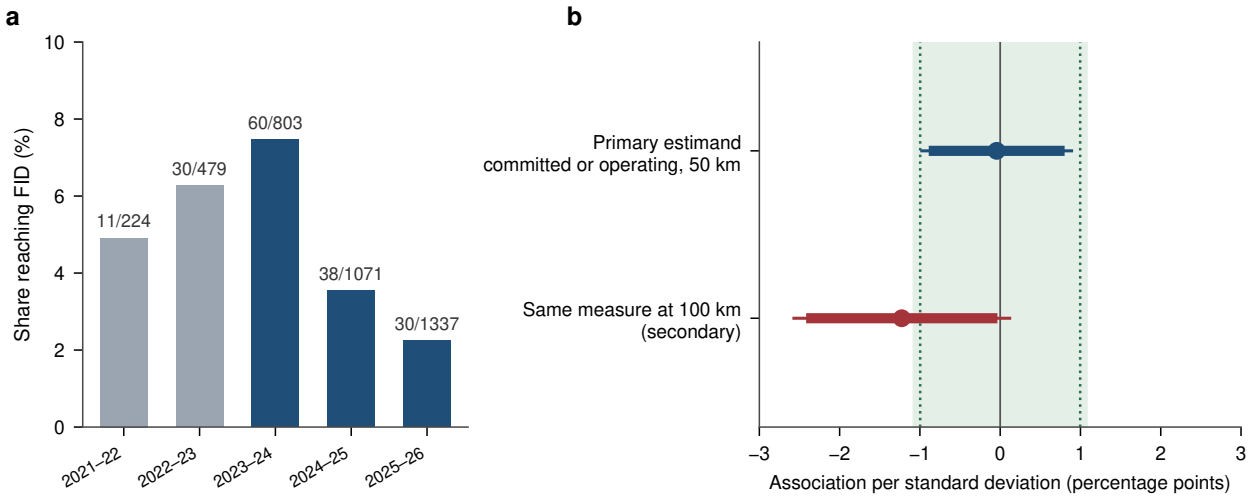


Fig. 1 The transition to a final investment decision, and the primary estimand against the equivalence benchmarks. **a** Share of at-risk investment units reaching a decision between consecutive vintages, events over units at risk above each bar, the first two bars being the extension window where coverage is below 65%. **b** Association per standard deviation of exposure in percentage points, thick bars 90% intervals and thin bars 95% intervals, the shaded band the 1.08-point benchmark and the dotted lines the 1.00-point benchmark

variation within a country and a vintage, and 66 of the 244 cells are singletons, which is why we take country and vintage effects as the primary specification. The equivalence conclusion is conditional on that choice, and we do not extend it to the other specifications: they are not significantly different from zero, which is a weaker statement.

Table 2 re-estimates the primary specification under six definitions of the investment unit and one restriction of the sample. Coefficients and standard errors both move with the definition, from −0.0008 (0.0087) under the strict rule to −0.0280 (0.0172) at the widest purely geographic aggregation, 10 km, and one definition changes the sign, the rule without the shared-source condition returning +0.0018 (0.0084). Every estimate remains statistically indistinguishable from zero at the 10% level, the widest coming closest at $p = 0.104$.

The remaining checks are in Online Resource 1 and none disturbs the estimate. Across the twelve exposure measures the per-standard-deviation associations run from −0.11 to +0.28 points at 50 km; at 100 km all five count and capacity estimates are negative and two reach the 10% level in raw terms, but neither survives correction and the joint test does not reject at $p = 0.608$. Changing the window, the variance estimator or the estimator itself,

Table 2 Alternative investment units and alternative specifications

Specification	Obs.	Events	Coefficient (s.e.)
Definition of the investment unit			
Strict investment unit	3211	128	-0.0008 (0.0087)
Identical coordinates	2935	127	-0.0036 (0.0109)
Geographic, 0.5 km	2861	128	-0.0075 (0.0112)
Geographic, 2 km	2699	124	-0.0182 (0.0122)
Geographic, 10 km	2270	111	-0.0280 (0.0172)
Investment unit without the shared-source rule	3308	129	0.0018 (0.0084)
Projects with stable coordinates only	3177	126	-0.0029 (0.0089)
Country and time fixed effects			
All 90 country fixed effects	3211	128	-0.0008 (0.0087)
Country fixed effects, ten-observation threshold	3211	128	-0.0001 (0.0087)
Small countries dropped from the sample	3031	122	0.0000 (0.0087)
Country and region-by-vintage fixed effects	3211	128	0.0017 (0.0088)
Country-by-vintage fixed effects	3211	128	0.0094 (0.0082)

This table re-estimates Eq. 3 with exposure measured as the log count of committed or operating units within 50 km. Every row includes the full control vector; rows in the unit block use vintage and country fixed effects. The unit block varies the rule that individuates investment units, holding fixed the criterion that defines what is counted. Standard errors are the spatio-temporal Conley estimator. The estimation window, variance estimator, rare-event and selection blocks are in Online Resource 1. Source: prepared by the authors.

and weighting by the inverse probability of appearing in the panel (Wooldridge 2007), leave the coefficient between -0.0042 and $+0.0017$; six heterogeneity interactions are individually insignificant and a joint Wald test does not reject at $p = 0.489$.

4.3 Sensitivity to the definition of the investment unit

One diagnostic shows why the unit had to be built before any of this was estimated. Take the cohort of projects that were pre-decision and geolocated in the September 2025 vintage, and the residual of a linear probability model of cancellation on project covariates and country fixed effects. Computed at the project level on a ten-nearest-neighbour graph, that residual is concentrated in space: the Moran statistic (Moran 1950) is $+0.037$ against a within-country randomization null of -0.018, $z = 5.0$, so cancellations would appear contagious among neighbours.

The concentration is carried by links inside sites. Only 7.9% of the links in the graph join two projects of one site, yet deleting those links alone moves the statistic to -0.009,

inside the randomization null, and the concentration is absent under every definition of the investment unit, including the purely geographic ones. Documentary evidence points the same way: 80% of co-located canceled pairs share a declared source reference, against 38.5% of the 697 co-located pairs in the same cohort. These figures identify the links that carry the statistic; they do not establish that every such pair is one economic decision. For the estimates above, sensitivity to the unit is a property of this particular statistic, computed on rows at the project level, and not of the panel of Sect. 4.1, whose estimate the aggregation rule leaves statistically indistinguishable from zero throughout. Statistics, unit counts, z values and the corresponding figure are in Online Resource 1.

4.4 Interpretation and limitations

The result speaks to how local project concentrations mature. Announced projects cluster, and the hypothesis under test is that clustering pulls them across the decision margin; on this panel, at 50 and 100 km, the pull is not detectable at magnitudes the design can exclude. Three readings remain open and the data do not separate them.

The first is scale, which Sect. 2.2 leaves uncertain. If the complementarities that matter operate over a few hundred metres, a pull at that range is averaged away by radii of 50 and 100 km. Narrowing the radius would not settle the question, because at short range the estimate becomes sensitive to the aggregation rule itself: a pair the rule failed to group enters as two units at any radius. That is why we report a bound and not an estimate of a mechanism.

The second is offsetting channels. Committed neighbours may de-risk through demonstration and shared infrastructure while competing for the same grid connection, land and permits, and a net association near zero is what two channels of similar size would produce. Every count and capacity estimate at 100 km is negative, the sign rivalry would produce, but none survives correction, so this reading is suggested and not established. The result is close to what the evaluation literature on cluster programs reports outside energy (Martin

et al. 2011; Duranton 2011; Falck et al. 2010; Alecke et al. 2006; Calignano and Fitjar 2017; Grashof and Fornahl 2021).

The third is that the relevant complementarity is not spatial at this scale. The contracting literature locates the binding constraint in the gap between production cost and willingness to pay (Vazquez and Hallack 2024), which distance does not represent and our measures do not observe.

Four features limit the scope of the result. The association is conditional on the fixed-effects scheme, as Sect. 4.2 shows. The panel is measured on projects surviving into the recent vintages, so unobserved selection is not bounded, although the main window, the weighting and a calibration in Online Resource 1 leave the conclusion standing. Interaction is measured indirectly, by proximity to units in a given state, which captures shared conditions as well as any influence between them. And exposure has a large mass point at zero, 65% of at-risk observations having no committed or operating unit within 50 km.

5 Concluding remarks

Prior local commitment shows no measurable association with the transition of neighbouring hydrogen projects to a final investment decision. On 3,211 unit-period observations and 128 transitions between 2023 and 2026, a one standard deviation increase in committed or operating units within 50 km is associated with -0.04 percentage points in the transition probability, against a mean of 4.0%, and the estimate is statistically equivalent to zero within a benchmark of 1.08 points calibrated on the gap between committed and achievable 2030 capacity. Equivalence is a property of the primary specification: under country-by-vintage effects the estimate moves to $+0.51$ points at unchanged precision and the interval exceeds the benchmark.

The scope of that answer is set by what the design can see. It is an association, measured at 50 and 100 km while the scale at which hydrogen complementarities operate is not known,

and estimated on projects surviving into the recent vintages. Within that scope the finding is informative: the design would have detected a pull of the magnitude that would matter for the 2030 pipeline, and does not find one.

For policy, this does not evaluate cluster programs, whose causal effect is not identified here, and European strategy promotes the co-location of production, demand and infrastructure, which is not what this design tests. What it tests is the further hypothesis that producers already committed pull neighbouring producers across the decision margin, and it finds no support for that at the scale the data resolve. A design reaching shorter distances will need a source that individuates decisions directly, such as award records from support programs or grid connection queues.

The construction of investment units is the complementary contribution. A project database records rows whose relation to decisions is set by a recording convention, and Sect. 4.3 shows that a spatial statistic computed on rows can be carried entirely by links inside sites. The crosswalk, matching log and unit assignment accompanying this paper are supplied so that such statistics can be reported against alternative definitions of the unit, as regionally estimated elasticities are now reported against alternative aggregations (Békés and Harasztosi 2018).

References

Alamá-Sabater L, Artal-Tur A, Navarro-Azorín JM (2011) Industrial location, spatial discrete choice models and the need to account for neighbourhood effects. Ann Reg Sci 47(2):393–418. https://doi.org/10.1007/s00168-010-0383-7

Alecke B, Alsleben C, Scharr F, Untiedt G (2006) Are there really high-tech clusters? The geographic concentration of German manufacturing industries and its determinants. Ann Reg Sci 40(1):19–42. https://doi.org/10.1007/s00168-005-0014-x

Allison PD (1982) Discrete-time methods for the analysis of event histories. Sociol Methodol 13:61–98. https://doi.org/10.2307/270718

Almeida E, Silveira Neto RM, Rocha RM (2026) The spatial extent of human capital spillovers: evidence from the Brazilian labor market. Ann Reg Sci 75:36. https://doi.org/10.1007/s00168-026-01455-0

Anderson ML (2008) Multiple inference and gender differences in the effects of early intervention: a reevaluation of the Abecedarian, Perry Preschool, and Early Training projects. J Am Stat Assoc 103(484):1481–1495. https://doi.org/10.1198/016214508000000841

Andersson M, Gråsjö U (2009) Spatial dependence and the representation of space in empirical models. Ann Reg Sci 43(1):159–180. https://doi.org/10.1007/s00168-008-0211-5

Angrist JD (2014) The perils of peer effects. Labour Econ 30:98–108. https://doi.org/10.1016/j.labeco.2014.05.008

Arbia G, Espa G, Giuliani D (2016) Dirty spatial econometrics. Ann Reg Sci 56(1):177–189. https://doi.org/10.1007/s00168-015-0726-5

Arrow KJ, Fisher AC (1974) Environmental preservation, uncertainty, and irreversibility. Q J Econ 88(2):312–319. https://doi.org/10.2307/1883074

Békés G, Harasztosi P (2018) Grid and shake: spatial aggregation and the robustness of regionally estimated elasticities. Ann Reg Sci 60(1):143–170. https://doi.org/10.1007/s00168-017-0849-y

Bester CA, Conley TG, Hansen CB (2011) Inference with dependent data using cluster covariance estimators. J Econom 165(2):137–151. https://doi.org/10.1016/j.jeconom.2011.01.007

Calignano G, Fitjar RD (2017) Strengthening relationships in clusters: how effective is an indirect policy measure carried out in a peripheral technology district? Ann Reg Sci 59(1):139–169. https://doi.org/10.1007/s00168-017-0821-x

Cameron AC, Gelbach JB, Miller DL (2011) Robust inference with multiway clustering. J Bus Econ Stat 29(2):238–249. https://doi.org/10.1198/jbes.2010.07136

Colella F, Lalive R, Sakalli SO, Thoenig M (2023) acreg: arbitrary correlation regression. Stata J 23(1):119–147. https://doi.org/10.1177/1536867X231162031

Conley TG (1999) GMM estimation with cross sectional dependence. J Econom 92(1):1–45. https://doi.org/10.1016/S0304-4076(98)00084-0

Contreras SA (2022) One size does not fit all: evaluating the impact of microenterprise measurement on policy evaluation. Ann Reg Sci 68(3):587–613. https://doi.org/10.1007/s00168-021-01094-7

Díaz Dapena A, Fernández Vázquez E, Rubiera Morollón F (2016) The role of spatial scale in regional convergence: the effect of MAUP in the estimation of β-convergence equations. Ann Reg Sci 56(2):473–489. https://doi.org/10.1007/s00168-016-0750-0

Dixit AK, Pindyck RS (1994) Investment under uncertainty. Princeton University Press, Princeton. ISBN 978-0-691-03410-2

Duranton G (2011) California dreamin': the feeble case for cluster policies. Rev Econ Anal 3(1):3–45. https://doi.org/10.15353/rea.v3i1.1375

Duranton G, Puga D (2004) Micro-foundations of urban agglomeration economies. In: Henderson JV, Thisse JF (eds) Handbook of regional and urban economics, vol 4. Elsevier, Amsterdam, pp 2063–2117. https://doi.org/10.1016/S1574-0080(04)80005-1

European Commission (2020) A hydrogen strategy for a climate-neutral Europe. COM(2020) 301 final. European Commission, Brussels. CELEX 52020DC0301. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:52020DC0301. Accessed 9 September 2026

Ellison G, Glaeser EL, Kerr WR (2010) What causes industry agglomeration? Evidence from coagglomeration patterns. Am Econ Rev 100(3):1195–1213. https://doi.org/10.1257/aer.100.3.1195

Falck O, Heblich S, Kipar S (2010) Industrial innovation: direct evidence from a cluster-oriented policy. Reg Sci Urban Econ 40(6):574–582. https://doi.org/10.1016/j.regsciurbeco.2010.03.007

Firth D (1993) Bias reduction of maximum likelihood estimates. Biometrika 80(1):27–38. https://doi.org/10.1093/biomet/80.1.27

Gedikoglu H, Tandogan S, Parcell J (2023) Neighbor effects on adoption of conservation practices: cases of grass filter systems and injecting manure. Ann Reg Sci 70(3):723–756. https://doi.org/10.1007/s00168-019-00963-6

Gibbons S, Overman HG (2012) Mostly pointless spatial econometrics? J Reg Sci 52(2):172–191. https://doi.org/10.1111/j.1467-9787.2012.00760.x

Grashof N, Fornahl D (2021) "To be or not to be" located in a cluster? A descriptive meta-analysis of the firm-specific cluster effect. Ann Reg Sci 67(3):541–591. https://doi.org/10.1007/s00168-021-01057-y

Greenstone M, Hornbeck R, Moretti E (2010) Identifying agglomeration spillovers: evidence from winners and losers of large plant openings. J Polit Econ 118(3):536–598. https://doi.org/10.1086/653714

Hartman E, Hidalgo FD (2018) An equivalence approach to balance and placebo tests. Am J Polit Sci 62(4):1000–1013. https://doi.org/10.1111/ajps.12387

Heinze G, Schemper M (2002) A solution to the problem of separation in logistic regression. Stat Med 21(16):2409–2419. https://doi.org/10.1002/sim.1047

International Energy Agency (2026a) Global hydrogen review 2026. IEA, Paris. Published 18 June 2026, licence CC BY 4.0. https://www.iea.org/reports/global-hydrogen-review-2026. Accessed 9 September 2026

International Energy Agency (2026b) Hydrogen production and infrastructure projects database. IEA, Paris. Updated 18 June 2026, licence CC BY 4.0. https://www.iea.org/data-and-statistics/data-product/hydrogen-production-and-infrastructure-projects-database. Accessed 9 September 2026

Jenkins SP (1995) Easy estimation methods for discrete-time duration models. Oxf Bull Econ Stat 57(1):129–136. https://doi.org/10.1111/j.1468-0084.1995.tb00031.x

Kellogg R (2014) The effect of uncertainty on investment: evidence from Texas oil drilling. Am Econ Rev 104(6):1698–1734. https://doi.org/10.1257/aer.104.6.1698

King G, Zeng L (2001) Logistic regression in rare events data. Polit Anal 9(2):137–163. https://doi.org/10.1093/oxfordjournals.pan.a004868

Kosfeld R, Eckey HF, Lauridsen J (2011) Spatial point pattern analysis and industry concentration. Ann Reg Sci 47(2):311–328. https://doi.org/10.1007/s00168-010-0385-5

Lacombe DJ, LeSage JP (2018) Use and interpretation of spatial autoregressive probit models. Ann Reg Sci 60(1):1–24. https://doi.org/10.1007/s00168-015-0705-x

Lang G, Marcon E, Puech F (2020) Distance-based measures of spatial concentration: introducing a relative density function. Ann Reg Sci 64(2):243–265. https://doi.org/10.1007/s00168-019-00946-7

Larsson JP (2014) The neighborhood or the region? Reassessing the density–wage relationship using geocoded data. Ann Reg Sci 52(2):367–384. https://doi.org/10.1007/s00168-014-0590-8

Lee S, Shaikh AM (2014) Multiple testing and heterogeneous treatment effects: re-evaluating the effect of PROGRESA on school enrollment. J Appl Econom 29(4):612–626. https://doi.org/10.1002/jae.2327

Manski CF (1993) Identification of endogenous social effects: the reflection problem. Rev Econ Stud 60(3):531–542. https://doi.org/10.2307/2298123

Marshall A (1890) Principles of economics. Macmillan, London

Martin P, Mayer T, Mayneris F (2011) Public support to clusters: a firm level study of French "local productive systems". Reg Sci Urban Econ 41(2):108–123. https://doi.org/10.1016/j.regsciurbeco.2010.09.001

McDonald R, Siegel D (1986) The value of waiting to invest. Q J Econ 101(4):707–727. https://doi.org/10.2307/1884175

Moran PAP (1950) Notes on continuous stochastic phenomena. Biometrika 37(1-2):17–23. https://doi.org/10.1093/biomet/37.1-2.17

Odenweller A, Ueckerdt F (2025) The green hydrogen ambition and implementation gap. Nat Energy 10(1):110–123. https://doi.org/10.1038/s41560-024-01684-7

Rainey C (2014) Arguing for a negligible effect. Am J Polit Sci 58(4):1083–1091. https://doi.org/10.1111/ajps.12102

Romano JP, Wolf M (2005) Exact and approximate stepdown methods for multiple hypothesis testing. J Am Stat Assoc 100(469):94–108. https://doi.org/10.1198/016214504000000539

Romano JP, Wolf M (2016) Efficient computation of adjusted p-values for resampling-based stepdown multiple testing. Stat Probab Lett 113:38–40. https://doi.org/10.1016/j.spl.2016.02.012

Rosenthal SS, Strange WC (2004) Evidence on the nature and sources of agglomeration economies. In: Henderson JV, Thisse JF (eds) Handbook of regional and urban economics, vol 4. Elsevier, Amsterdam, pp 2119–2171. https://doi.org/10.1016/S1574-0080(04)80006-3

Vazquez M, Hallack M (2024) Auctions to reveal consumers' willingness to pay for low-carbon hydrogen projects: combining lessons from renewables and natural-gas industries. Clean Energy 8(4):103–110. https://doi.org/10.1093/ce/zkae040

Wooldridge JM (2007) Inverse probability weighted estimation for general missing data problems. J Econom 141(2):1281–1301
. https://doi.org/10.1016/j.jeconom.2007.02.002

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- supplementarymaterial.zip