Supporting Information

The fee schedule is an entry rule:

Ad valorem tariffs and market structure in Portugal's notarial reform

A Numerus Clausus Revision Check

This appendix reports a direct check of whether the 2014/2015 revision plausibly accounts for the medium-term inversion documented in the main analysis. The test compares the municipal configuration immediately before the revision, using 2013, with the post-revision configuration in 2016, allowing one year for implementation. Because the available DGPJ files used in the article do not include a separate official quota variable, the test uses the observed number of active public and private positions in each municipality as the operational measure of the effective entry constraint.

Table A1: Observed-position changes around the 2014/2015 numerus clausus revision

Statistic	Value
Municipalities in matched sample	273
Municipalities with any observed-position change	17
Share with any observed-position change (%)	6.23
Mean change, all municipalities	-0.037
Mean absolute change, changed municipalities	1.059
Median change, changed municipalities	-1.000
Minimum change	-2
Maximum change	1
Public/private/competitive reclassifications	9
Single/multi-entry reclassifications	9

Source: Author's calculations using DGPJ administrative data. The observed-position measure is computed as public plus private notarial positions in each municipality.

The results show that the change was limited in scale. Of the 273 municipalities matched between 2013 and 2016, 17 experienced an observed-position adjustment, corresponding to 6.23% of the matched sample. Among affected municipalities, the mean absolute change was 1.06 positions. This supports the interpretation that the revision was marginal relative to the national municipal structure.

The appendix also clarifies the scope of the claim. Using observed positions as a proxy for the effective constraint, nine municipalities shifted between single-entry and multi-entry status. Accordingly, the evidence supports a narrower formulation: the 2014/2015 revision was small in aggregate scope and cannot plausibly explain the broad 2009–2019 inversion documented in the paper, but the available operational data do not support the stronger claim that no municipality changed single-entry versus multi-entry status.

B Robustness and Diagnostics for the Multinomial Logit

This appendix provides diagnostic evidence in support of the multinomial logit results reported in Section ??. Two model properties require verification. First, several specifications in Table ?? display very large relative risk ratios, particularly for the Market category in 2009. These may reflect strong separation of the activity covariate across categories rather than a pathological estimation failure. Second, the multinomial logit imposes the independence of irrelevant alternatives (IIA), which is potentially restrictive given that private monopoly and competitive market structures share privately-operated provision. The diagnostics below address both concerns in turn.

The Appendix B diagnostics support the paper's substantive interpretation while qualifying the interpretation of coefficient magnitudes. The evidence indicates that the structural inversion documented between 2009 and 2019 is not driven by sparse outcome cells, influential outliers, or the IIA assumption. However, the very large RRRs in some multinomial logit specifications should be read as evidence of strong separation in the activity distribution rather than as precise effect-size estimates.

Cell counts and covariate overlap

Table B1: Cell counts by wave and market structure

Wave	Specification	Public monopoly	Private monopoly	Market	Total
2009	Total acts	79	117	82	278
2009	Real-estate acts	78	117	82	277
2019	Total acts	51	102	123	276
2019	Real-estate acts	**32**	102	123	257

Counts are the estimation samples, not the wave frames. The real-estate rows are smaller because twenty municipality-years record no property acts and the logarithm is undefined for them. In 2009 that removes one municipality (Vila Nova da Barquinha, a public monopoly with 7,376 residents and 67 notarial acts of any kind). In 2019 it removes nineteen, every one of them a public monopoly, which is why the public-monopoly cell falls from 51 to 32 while the other two are unaffected. Bold marks the cell that falls below 37% of its own wave frame.

Table B1 reports the cells of the estimation samples rather than of the wave frames, and the distinction matters here. Every cell contains more than twenty municipalities, so the large relative risk ratios in Table ?? do not arise from empty or near-empty categories. But the cells are not the same across specifications. The real-estate specification loses one municipality in 2009 and nineteen in 2019 to the logarithm of a zero act count, and all twenty are public monopolies, so its base category falls to 32 in the second wave against 51 in the total-acts specification on the same municipalities. That is the sample loss set out in Section ??, shown here cell by cell.

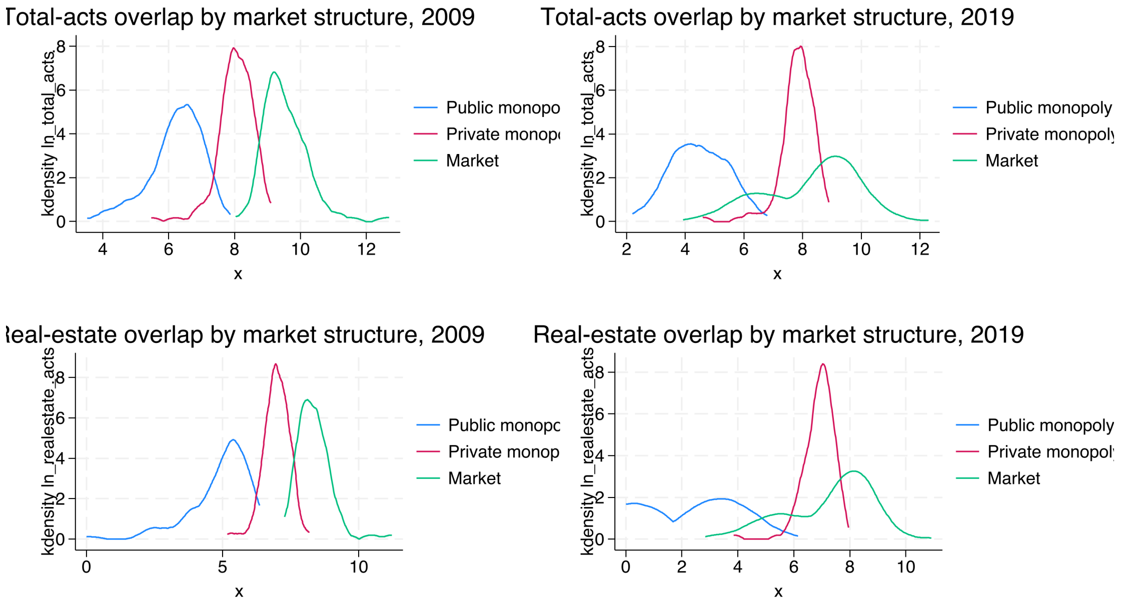

Figure B1: Overlap diagnostics for activity covariates by market structure

Figure B1 shows that activity distributions overlap across categories but unevenly. In 2009, the Market category is concentrated in higher-activity municipalities. By 2019, these distributions show the inversion documented in the main text. The partial separation confirms that the model is fitting a genuine distributional difference and not a numerical artifact, which explains why some RRRs are large.

Marginal effects in the real-estate specification

Figure **??** in the main text plots the total-acts specification, which is the one the cross-wave comparison is led with because it loses no municipality. Figure B2 plots the same quantities for the real-estate specification. The pattern is the one the main text describes: competitive structure rising monotonically with activity in 2009, and the private-monopoly curve peaking at the top of the distribution in 2019. The 2019 panel is estimated on the sample that drops nineteen public monopolies, so the level of the public-monopoly curve there is not comparable with the total-acts panel, while the crossing of the two private curves, which is what the inversion concerns, is unaffected.

Influence diagnostics and trimming robustness

Tables B2 to B5 identify the ten municipalities with the highest influence in each wave and specification. These are not data errors: they are municipalities located in parts of the covariate distribution that are most informative for the inversion mechanism. The appropriate robustness check is symmetric trimming, not selective removal.

Figure B2: Predicted probability of private monopoly and of competitive market, real-estate specification

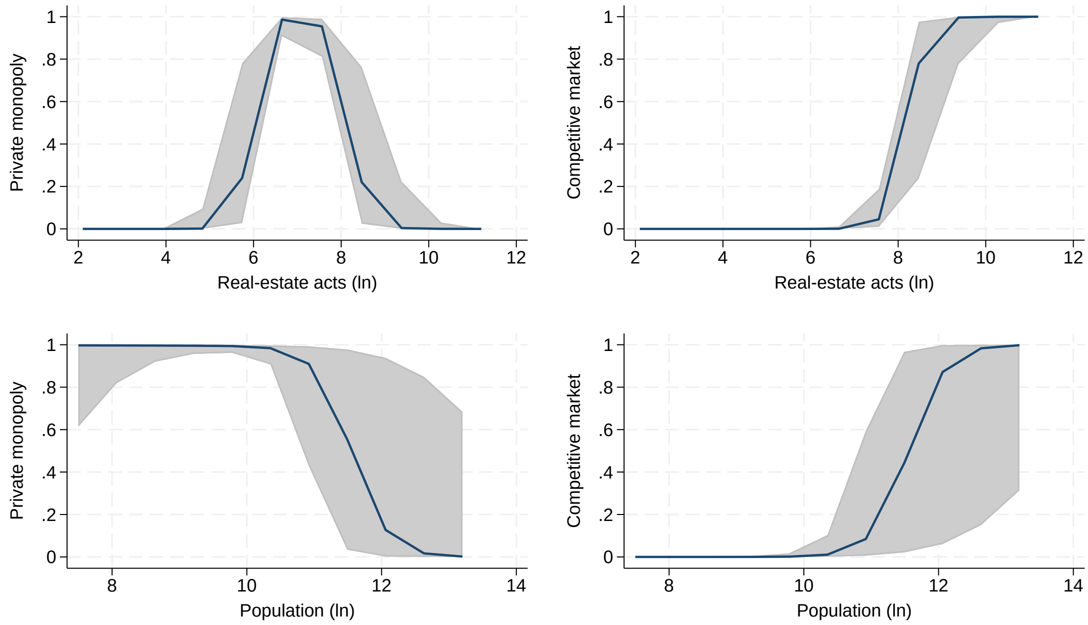

(a) 2009

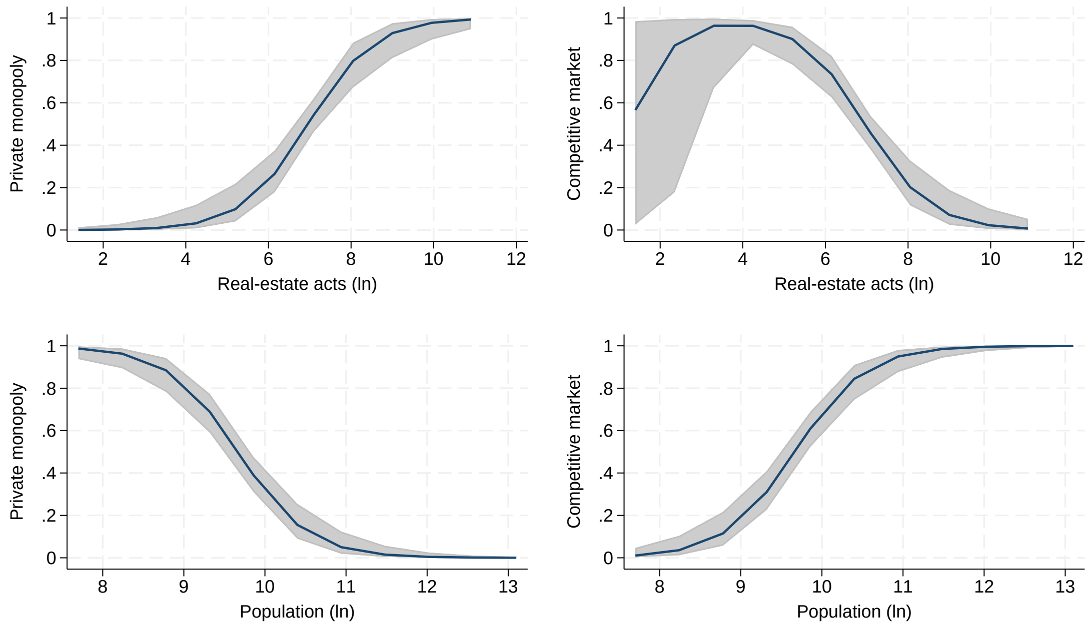

(b) 2019

Notes: As Figure **??** in the main text, for the real-estate specification. The upper row of each panel plots the two private structures against log real-estate acts and the lower row against log population. Public monopoly is the base outcome and is not plotted. Bands are 95% intervals formed on the log-odds scale and mapped back through the inverse logit.

Table B2: Influence diagnostics, total acts, 2009

Rank	Municipality	Outcome	Activity	Pr(public)	Pr(private)	Pr(market)
1	Paredes de Coura	2	5.472	0.998	0.002	0.000
2	Constância	2	6.217	0.947	0.053	0.000
3	Portel	2	6.837	0.792	0.208	0.000
4	Mondim de Basto	2	6.989	0.834	0.166	0.000
5	Montalegre	1	7.492	0.187	0.813	0.000
6	Alijó	1	6.720	0.670	0.330	0.000
7	Monção	2	6.946	0.452	0.548	0.000
8	Ferreira do Alentejo	1	6.938	0.599	0.401	0.000
9	Torre de Moncorvo	1	7.018	0.491	0.509	0.000
10	Sabrosa	1	6.962	0.601	0.399	0.000

Table B3: Influence diagnostics, total acts, 2019

Rank	Municipality	Outcome	Activity	Pr(public)	Pr(private)	Pr(market)
1	Chamusca	3	3.932	0.745	0.008	0.247
2	Sernancelhe	3	5.106	0.846	0.051	0.103
3	Pedrógão Grande	3	5.587	0.799	0.150	0.051
4	Golegã	3	5.124	0.702	0.107	0.191
5	Arraiolos	3	4.828	0.641	0.061	0.298
6	Vila Viçosa	2	4.595	0.559	0.048	0.393
7	Castro Verde	1	2.833	0.816	0.004	0.181
8	Idanha-a-Nova	1	5.257	0.378	0.119	0.504
9	Ferreira do Alentejo	1	4.625	0.615	0.040	0.345
10	Borba	1	6.168	0.104	0.493	0.403

Table B4: Influence diagnostics, real-estate acts, 2009

Rank	Municipality	Outcome	Activity	Pr(public)	Pr(private)	Pr(market)
1	Constância	2	5.176	0.929	0.071	0.000
2	Paredes de Coura	2	5.472	0.954	0.046	0.000
3	Vale de Cambra	1	6.078	0.202	0.798	0.000
4	Mondim de Basto	2	6.100	0.704	0.296	0.000
5	Vila Nova de Paiva	1	6.250	0.297	0.703	0.000
6	Penela	2	6.426	0.844	0.156	0.000
7	Golegã	1	5.883	0.464	0.536	0.000
8	Montalegre	1	6.358	0.348	0.652	0.000
9	Sabrosa	1	6.019	0.599	0.401	0.000
10	Mora	1	5.283	0.868	0.132	0.000

Table B5: Influence diagnostics, real-estate acts, 2019

Rank	Municipality	Outcome	Activity	Pr(public)	Pr(private)	Pr(market)
1	Arraiolos	3	2.833	0.713	0.014	0.272
2	Sernancelhe	3	4.007	0.873	0.037	0.089
3	Pedrógão Grande	3	4.477	0.803	0.139	0.058
4	Chamusca	3	2.833	0.455	0.015	0.531
5	Idanha-a-Nova	1	3.434	0.525	0.034	0.440
6	Tabuaço	1	3.526	0.604	0.061	0.335
7	Torre de Moncorvo	1	3.850	0.535	0.060	0.406
8	Vila Nova da Barquinha	3	3.784	0.449	0.066	0.485
9	Castro Verde	1	0.000	0.910	0.000	0.089
10	Miranda do Douro	1	3.912	0.673	0.061	0.266

Table B6: Trimming robustness: total-acts specification

	(1) 2009 full	(2) 2009 trimmed	(3) 2019 full	(4) 2019 trimmed
Private_monopoly				
Log total acts	34.640***	113.912***	19.411***	19.198***
	(24.422)	(106.772)	(10.535)	(10.864)
Log municipal population	4.181	2.384	4.906	4.022
	(3.667)	(2.213)	(6.120)	(4.921)
Log municipal income	63.867	255.526	112.634	1100.709
	(186.533)	(874.592)	(390.808)	(5364.283)
Constant	0.000***	0.000***	0.000**	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)
Market				
Log total acts	4107.653***	13521.748***	5.319***	5.055***
	(6071.691)	(21663.494)	(2.618)	(2.611)
Log municipal population	149.211***	85.139***	56.500***	43.160***
	(205.192)	(119.917)	(69.442)	(51.770)
Log municipal income	167.191	669.446	73.378	533.585
	(744.896)	(3213.356)	(245.965)	(2584.866)
Constant	0.000***	0.000***	0.000***	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)
Observations	278	250	276	248
Log likelihood	-60.77	-53.58	-149.66	-146.93

Relative risk ratios. Standard errors in parentheses.
Trimmed columns exclude the bottom and top 5 percent of the activity distribution, with cut points computed within wave.
Public monopoly is the reference category and is not printed.

Full-sample versus trimmed estimates, total acts. Public Monopoly is the reference category.

Table B7: Trimming robustness: real-estate specification

	(1) 2009 full	(2) 2009 trimmed	(3) 2019 full	(4) 2019 trimmed
Private_monopoly				
Log real-estate acts	410.698***	410.652***	20.154***	25.844***
	(556.601)	(556.581)	(12.435)	(18.706)
Log municipal population	0.791	0.791	3.775	2.525
	(0.946)	(0.946)	(5.760)	(4.048)
Log municipal income	35.504	35.503	7560.841	9007.315
	(144.701)	(144.693)	(47002.102)	(64165.025)
Log real-estate value	4.214**	4.214**	1.807	1.422
	(2.478)	(2.478)	(1.182)	(0.946)
Constant	0.000***	0.000***	0.000**	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)
Market				
Log real-estate acts	47038.931***	47024.853***	5.720***	7.525***
	(85705.740)	(85686.481)	(3.260)	(5.161)
Log municipal population	34.432**	34.429**	37.881**	22.574**
	(57.187)	(57.183)	(57.018)	(35.683)
Log municipal income	53.046	53.032	2835.949	2283.101
	(293.702)	(293.624)	(17534.977)	(16194.486)
Log real-estate value	14.665**	14.664**	2.015	1.553
	(15.655)	(15.654)	(1.294)	(1.020)
Constant	0.000***	0.000***	0.000**	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)
Observations	277	249	257	231
Log likelihood	-40.81	-40.81	-140.70	-136.22

Relative risk ratios. Standard errors in parentheses.
Trimmed columns exclude the bottom and top 5 percent of the activity distribution, with cut points computed within wave.
Public monopoly is the reference category and is not printed.

Full-sample versus trimmed estimates, real-estate. Public Monopoly is the reference category.

Tables B6 and B7 exclude the bottom and top five percent of the activity distribution, with the cut points computed within each wave rather than on the pooled data. The trim removes 28 municipalities from each wave in the total-acts specification, 28 from 2009 and 26 from 2019 in the real-estate one, and the resulting sample sizes are printed in the table so that it can be seen to have bound.

The signs and substantive ordering are preserved: activity and population continue to distinguish private monopoly and market outcomes relative to public monopoly. The magnitude of the largest relative risk ratios is sensitive to trimming, as expected under tail-driven separation, since the 2009 private-monopoly activity ratio moves from 34.6 to 113.9 in the total-acts specification. The core empirical pattern is not eliminated. The interpretation in the main text relies on signs, predicted probabilities, and average marginal effects rather than on the literal magnitude of the extreme ratios.

Penalized maximum-likelihood robustness

A multinomial Firth estimator was not available, so we apply penalization to binary problems. Two distinct penalized estimators appear in this article and they answer different questions, so we state the difference here rather than leave it to be inferred from the numbers.

The main text reports *contrasts against the public-monopoly baseline*. Each is a penalized

Table B8: Penalized maximum-likelihood status

Wave	Specification	Method	Status
2009	total	firthlogit	binary Firth one-vs-rest estimated
2009	realestate	firthlogit	binary Firth one-vs-rest estimated
2019	total	firthlogit	binary Firth one-vs-rest estimated
2019	realestate	firthlogit	binary Firth one-vs-rest estimated

logit on the sub-sample containing only the baseline category and the contrast category, which is the binary problem corresponding to a multinomial contrast. Those estimates are what the main text quotes when it asks whether the ordering of the multinomial coefficients survives the correction for separation.

The tables below report *one-vs-rest* estimates. Each contrasts a single market structure against the other two on the full wave. These are not multinomial contrasts, they do not share a reference category, and their coefficients are not comparable in level to the ones in the main text. They are reported because the one-vs-rest form uses every observation in the wave and therefore provides a separation diagnostic that does not depend on which pair of categories happens to be retained. If quasi-complete separation were driving the main results, the penalized estimates would be expected to reverse signs or lose significance under either form.

They do not, with one exception that the main text reports and that this appendix must not paper over. In the total-acts specification both forms place the 2009 competitive coefficient above the private-monopoly one and reverse that ordering in 2019, which is the inversion the article reports, and the one-vs-rest form delivers it more starkly. The one-vs-rest form does the same in the real-estate specification. The contrast form does not: on the real-estate sub-samples the 2009 activity coefficient is 5.176 for private monopoly against 2.881 for the competitive category, so the 2009 ordering is already reversed there before any comparison with 2019 is drawn. That is why Section **??** rests H1.3 only weakly on the first wave and rests the two-wave contrast on 2019. What survives across all three readings is the 2019 configuration and the direction of the change, not the 2009 ordering in the real-estate specification.

The Firth estimates preserve the substantive direction of the main results. Higher activity lowers the probability of public monopoly and is positively associated with private monopoly in both waves. The association between activity and the market outcome is strong in 2009 and weakens sharply by 2019, consistent with the structural inversion documented in the paper. Quasi-complete separation is not generating the headline pattern.

Table B9: Binary Firth one-vs-rest estimates: total-acts specification

	(1) 2009 Public	(2) 2009 Priv. mon.	(3) 2009 Compet.	(4) 2019 Public	(5) 2019 Priv. mon.	(6) 2019 Compet.
xb						
Log total acts	-3.278***	1.731***	4.306***	-1.978***	1.772***	-0.241
	(0.660)	(0.293)	(1.203)	(0.447)	(0.251)	(0.158)
Log municipal population	-1.294	-2.281***	3.200***	-2.519**	-2.597***	1.627***
	(0.831)	(0.370)	(0.973)	(1.114)	(0.366)	(0.307)
Log municipal income	-4.092	-0.079	0.922	-5.207	0.282	-0.826
	(2.997)	(1.008)	(3.319)	(3.186)	(1.205)	(1.193)
Constant	62.322***	8.679	-77.325***	68.074***	9.351	-8.580
	(20.585)	(6.212)	(25.390)	(23.355)	(7.585)	(7.593)
Observations	278	278	278	276	276	276
Log likelihood	-34.27	-155.65	-26.00	-27.72	-130.48	-137.53

Penalized (Firth) logit coefficients, not relative risk ratios.
Standard errors in parentheses. Each column contrasts one market structure against the other two on the full wave, so the columns are not multinomial contrasts and do not share a reference category.

Binary Firth (penalized) one-vs-rest estimates, total acts. Public Monopoly is the reference category.

Table B10: Binary Firth one-vs-rest estimates: real-estate specification

	(1) 2009 Public	(2) 2009 Priv. mon.	(3) 2009 Compet.	(4) 2019 Public	(5) 2019 Priv. mon.	(6) 2019 Compet.
xb						
Log real-estate acts	-5.176***	1.727***	4.063***	-1.829***	1.569***	-0.322**
	(1.159)	(0.282)	(1.067)	(0.481)	(0.238)	(0.148)
Log municipal population	0.291	-2.339***	3.299***	-2.139	-2.422***	1.608***
	(1.089)	(0.375)	(1.042)	(1.314)	(0.365)	(0.320)
Log municipal income	-3.563	0.124	0.594	-7.320	1.026	-0.604
	(4.054)	(1.129)	(3.632)	(5.899)	(1.263)	(1.274)
Log real-estate value	-1.226**	0.254	1.035	-0.572	-0.018	0.028
	(0.533)	(0.210)	(0.804)	(0.590)	(0.244)	(0.226)
Constant	64.822**	7.179	-80.692***	81.504*	5.919	-9.867
	(26.992)	(6.498)	(26.694)	(41.798)	(7.688)	(7.925)
Observations	277	277	277	257	257	257
Log likelihood	-18.70	-148.58	-21.97	-20.65	-122.03	-131.99

Penalized (Firth) logit coefficients, not relative risk ratios.
Standard errors in parentheses. Each column contrasts one market structure against the other two on the full wave, so the columns are not multinomial contrasts and do not share a reference category.

Binary Firth (penalized) one-vs-rest estimates, real-estate. Public Monopoly is the reference category.

IIA diagnostics and multinomial probit

Table B11: IIA diagnostics. Note: p-values and log-likelihood statistics are available in the replication output.

Wave	Specification	Hausman–McFadden	Small–Hsiao
2009	total	Invalid statistic	See replication log
2019	total	See replication log	See replication log
2009	realestate	Invalid statistic	See replication log
2019	realestate	Invalid statistic	See replication log

The Hausman–McFadden test is known to produce invalid statistics in finite samples; the multinomial probit (Table B12) is the primary IIA robustness check.

Table B11 is the weakest diagnostic in this appendix. The Hausman–McFadden test fails or returns invalid statistics in several specifications, a known pathology in finite samples (Cheng & Long, 2007). We therefore rely on the multinomial probit comparison below as the primary robustness criterion. The probit relaxes IIA by allowing correlated errors across alternatives. Where logit and probit average marginal effects point in the same direction and are of comparable magnitude, the IIA assumption does not affect the paper's substantive inference.

Table B12: Average marginal effects: multinomial logit versus multinomial probit, total acts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log total acts	0.00107	0.00205	0.130***	0.130***	0.226***	0.217***	-0.154***	-0.145***
	(0.03)	(0.06)	(5.00)	(4.81)	(7.34)	(7.58)	(-5.00)	(-5.03)
Log municipal population	-0.0446	-0.0368	0.0976***	0.103***	-0.325***	-0.319***	0.419***	0.406***
	(-1.14)	(-0.92)	(4.21)	(4.56)	(-7.39)	(-7.65)	(9.19)	(9.35)
Log municipal income	0.128	0.116	0.0264	0.0274	0.127	0.152	0.0137	-0.0138
	(0.91)	(0.85)	(0.29)	(0.30)	(0.63)	(0.72)	(0.07)	(-0.07)
Observations	278	278	278	278	276	276	276	276

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Average marginal effects, multinomial logit vs. multinomial probit, total acts. Model fit statistics are reported in the replication output.

Table B13: Average marginal effects: multinomial logit versus multinomial probit, real-estate acts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log real-estate acts	0.00902	0.00229	0.106***	0.116***	0.205***	0.201***	-0.151***	-0.148***
	(0.32)	(0.09)	(5.25)	(5.81)	(7.17)	(7.44)	(-5.23)	(-5.39)
Log municipal population	-0.0890***	-0.0903***	0.0845***	0.0855***	-0.314***	-0.318***	0.382***	0.380***
	(-2.84)	(-2.90)	(3.93)	(3.83)	(-7.42)	(-7.76)	(8.43)	(8.51)
Log municipal income	0.0591	0.0675	0.00901	0.0121	0.207	0.203	-0.0528	-0.0631
	(0.52)	(0.65)	(0.11)	(0.15)	(1.03)	(0.98)	(-0.26)	(-0.30)
Log real-estate value	-0.000462	0.00217	0.0279	0.0268	-0.0100	-0.000590	0.0268	0.0179
	(-0.02)	(0.10)	(1.43)	(1.43)	(-0.27)	(-0.02)	(0.70)	(0.46)
Observations	278	278	278	278	276	276	276	276

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Average marginal effects, multinomial logit vs. multinomial probit, real-estate.

The average marginal effects tables above show that the multinomial logit and multinomial probit yield the same substantive reading. Across both specifications, the marginal effects for activity and population point in the same direction for the private monopoly and market outcomes. The inference that the relationship between municipal demand and market structure changed materially between 2009 and 2019 does not depend on the logit functional form.

C Spatial Dependence

Municipal observations are not independent, and the paper's mechanism makes the point sharper rather than softer: entry thresholds are calibrated on market size, market size is spatially clustered, and licensing boundaries are drawn on a map. This appendix reports the diagnostic in full.

C1. Construction

Municipal centroids are taken from the shapefile-derived file used for the maps in Section ??, joined to the analysis file through the same municipality crosswalk. The join covers all 278 mainland municipalities; the 30 unmatched records are the Azores and Madeira, which are outside every estimation sample in the article. Three row-standardized weight matrices are built from those centroids: the five nearest neighbors, the ten nearest neighbors, and inverse distance over all pairs. Moran's I is computed as $I = (n/S_0)\, z'Wz/z'z$ with z the demeaned variable, and its p-value comes from 999 random permutations of the variable across municipalities, which avoids assuming normality of quantities that are plainly not normal, such as binary indicators and multinomial residuals.

Two estimation samples are involved and each carries its own weight matrix. The entry model uses population, income and property value (278 municipalities in 2009, 276 in 2019, as in Table ??); the structure models add real-estate activity (277 and 257, as in Table ??). The reported entry threshold ratios reproduce Table ?? exactly in each block, which confirms the samples were not silently changed by the geographic join.

C2. Everything observed is clustered

Table C1: Moran's I for covariates and raw outcomes, five nearest neighbors

	2009		2019	
	I	p	I	p
$Pop_{(ln)}$	0.567	0.001	0.584	0.001
$Income_{(ln)}$	0.450	0.001	0.414	0.001
$Property\ Value_{(ln)}$	0.656	0.001	0.708	0.001
$Total\ Acts_{(ln)}$	0.406	0.001	0.352	0.001
Number of private offices	0.434	0.001	0.397	0.001
Public monopoly	0.299	0.001	0.184	0.001
Private monopoly	0.171	0.001	0.145	0.001
Competitive market	0.325	0.001	0.169	0.001

Notes: Row-standardized five-nearest-neighbor weights. p-values from 999 permutations, one-sided against positive autocorrelation. The covariate rows use the entry-model sample and the activity and structure rows the structure-model sample.

Property value is the most clustered variable in the data in both waves. That matters for what follows: it is the variable through which the *ad valorem* channel operates.

C3. The entry model leaves no residual dependence

Table C2: Residual Moran's I, ordered-probit entry model

	2009		2019	
Weights	I	p	I	p
Five nearest neighbors	−0.033	0.795	0.004	0.419
Ten nearest neighbors	−0.006	0.528	0.019	0.177
Inverse distance	−0.002	0.317	−0.000	0.255
Entry threshold ratios with the neighbors' covariates added				
ETR_2, baseline	2.180		2.245	
ETR_2, with neighbors	2.111		2.141	
ETR_3, baseline	1.631		2.332	
ETR_3, with neighbors	1.569		2.238	

Notes: Residuals are the observed count of private offices minus its predicted expectation. The lower panel adds the spatial lags of population, income and property value under five-nearest-neighbor weights.

Table C3 reports the estimates behind the lower panel of Table C2. None of the three spatial lags is distinguishable from zero in either wave: −0.172, −1.419 and 0.084 in 2009 ($p = 0.354$, 0.290 and 0.758), and −0.054, −1.741 and 0.207 in 2019 ($p = 0.739$, 0.164 and 0.402). The own-municipality population coefficient is unchanged in size and precision, at 2.895 and 1.814, both with z above ten, which is why the threshold ratios move by at most 0.104 when the lags are added. Neighboring demand does not enter the admission decision, which is consistent with a licensing rule written on municipal aggregates.

Table C3: Spatial-lag specification, ordered-probit entry model

	2009	2019
$Pop_{(ln)}$	2.895*** (0.268)	1.814*** (0.166)
$Income_{(ln)}$	1.497* (0.879)	0.484 (0.724)
$Property\ Value_{(ln)}$	0.185 (0.209)	−0.188 (0.203)
$W\ Pop_{(ln)}$	−0.172 (0.186)	−0.054 (0.162)
$W\ Income_{(ln)}$	−1.419 (1.341)	−1.741 (1.253)
$W\ Property\ Value_{(ln)}$	0.084 (0.273)	0.207 (0.247)
Observations	278	276
Log likelihood	−120.73	−180.51
Pseudo R^2	0.660	0.484
LR $\chi^2(6)$	467.67	339.12

Notes: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. W denotes the row-standardized five-nearest-neighbor spatial lag of the variable. Cut points are estimated but not reported. The specification without the lags is the one reported in Table **??**. Stata flags one completely determined observation in the 2009 column.

C4. The structure equations, and what removes the dependence

Three readings of Table C4 are worth separating, and Table C5 reports the estimates on

Table C4: Residual Moran's I, multinomial structure models, five nearest neighbors

	2009		2019	
Specification and category	I	p	I	p
Baseline				
Public monopoly	0.010	0.308	0.059	0.044
Private monopoly	0.016	0.259	0.083	0.016
Competitive market	−0.002	0.493	0.069	0.030
Baseline plus the neighbors' covariates				
Private monopoly	0.002	0.419	0.083	0.009
Competitive market	−0.004	0.519	0.067	0.027
With the activity × property value interaction				
Public monopoly	0.004	0.404	0.047	0.072
Private monopoly	0.007	0.353	0.032	0.182
Competitive market	−0.007	0.555	0.020	0.255
Interaction plus the neighbors' covariates				
Private monopoly	−0.031	0.782	0.023	0.239
Competitive market	−0.032	0.839	0.005	0.403
Wald test of the interaction, baseline sample	2.81 ($p = 0.245$)		32.13 ($p < 0.001$)	
Wald test, with the neighbors' covariates	4.09 ($p = 0.129$)		28.19 ($p < 0.001$)	

Notes: Residuals are $\mathbf{1}\{y_i = j\} - \hat{p}_{ij}$ for each category. The neighbors' covariates are the spatial lags of activity, population, income and property value under five-nearest-neighbor weights. In 2019 the pattern is identical under ten-nearest-neighbor and inverse-distance weights: the baseline residuals give 0.073 and 0.066 ($p = 0.003$ and 0.009) and 0.013 and 0.012 ($p = 0.006$ and 0.009) for private monopoly and competitive market, and the interaction residuals are insignificant under both. In 2009 the private-monopoly residual is marginally significant under those two weightings (0.046, $p = 0.042$; and 0.007, $p = 0.035$) although not under five nearest neighbors, and the interaction leaves it at 0.041 ($p = 0.043$) and 0.005 ($p = 0.067$); the first wave is therefore not quite as clean as the five-nearest-neighbor column alone suggests, which is one further reason the article rests the two-wave contrast on 2019.

which the second and third rest. First, the 2009 configuration is spatially clean under the baseline specification, so nothing in the first wave's estimates is a regional artifact.

Second, the 2019 baseline is not, and the table separates two statements that the residual diagnostic alone conflates. Neighboring activity does enter the 2019 equations: $W\,Total\ Acts_{(ln)}$ carries 2.012 ($p = 0.046$) in the private-monopoly equation and 1.859 ($p = 0.062$) in the market equation, while the lags of population, income and property value are individually indistinguishable from zero in both. But entering the specification is not the same as absorbing the dependence. With all four lags included the residual I stands at 0.083 and 0.067, exactly where the baseline left it, so the clustering in the 2019 residuals is not the level of demand in adjacent municipalities. The 2009 columns show neither effect: no lag reaches conventional significance in either equation, and the largest, $W\,Property\ Value_{(ln)}$ at 3.339 ($p = 0.096$) in the private-monopoly equation, is estimated on a fit in which 21 observations are completely determined.

Third, the interaction removes the residual dependence, and it does so without displacing the neighbors. In the interaction specification the activity lag is essentially unchanged, at 2.267 and 2.044 ($p = 0.051$ and 0.070), while the residual I falls to 0.023 and 0.005 and the interaction remains jointly significant ($\chi^2(2) = 28.19$, $p < 0.001$). The clustering is absorbed by the functional form, not by neighboring demand.

The interpretation is the one the paper's mechanism implies. What the baseline specification omits is not a regional effect but a functional form: under an *ad valorem* ceiling, structure responds to the product $\tau(V)Q$, and a model that enters V and Q additively misses exactly the variation that property value, the most spatially clustered variable in the data, carries across neighboring municipalities. The residual clustering in the 2019 baseline is that omission showing up in space. This does not establish that the residual process is non-spatial in truth, only that the specification the article reports is not leaving detectable spatial structure behind, and that the specification which does leave it is the one the model says is misspecified.

C5. Magnitude, penalization and collinearity

Three quantities bear on how Tables C3 and C5 should be read, and none of them is a coefficient.

The first is magnitude. Multinomial coefficients are log-odds against the base category, so the 2.012 that $W\,Total\ Acts_{(ln)}$ carries in the 2019 private-monopoly equation is not a quantity on the scale of the outcome. Averaged over the sample, the effect of the activity lag on the predicted probability of private monopoly is 0.040 (SE 0.048, $p = 0.412$) and on the probability of competitive market 0.000 (SE 0.049, $p = 0.995$). The only category on which it is significant is public monopoly, at -0.040 (SE 0.018, $p = 0.029$). The same specification puts the own-municipality effect at 0.219 (SE 0.031) on private monopoly

Table C5: Spatial-lag specifications, multinomial structure models

	2009		2019	
	Baseline	Interaction	Baseline	Interaction
Private monopoly				
$Total\ Acts_{(ln)}$	7.422*** (1.883)	29.93** (15.21)	3.418*** (0.746)	17.50*** (6.746)
$Property\ Value_{(ln)}$	0.833 (0.743)	14.08 (8.553)	−0.194 (1.118)	8.465** (3.413)
$Total\ Acts \times Prop.\ Value$	—	−2.109 (1.343)	—	−1.225* (0.658)
$Pop_{(ln)}$	−1.216 (1.518)	−1.345 (1.631)	1.532 (1.862)	1.142 (1.914)
$Income_{(ln)}$	5.147 (5.018)	5.727 (5.048)	12.04 (8.149)	12.62 (8.408)
$W\,Total\ Acts_{(ln)}$	1.858 (1.667)	1.809 (1.635)	2.012** (1.007)	2.267* (1.164)
$W\,Pop_{(ln)}$	−2.087 (2.733)	−0.923 (2.953)	−4.196 (2.808)	−5.215 (3.387)
$W\,Income_{(ln)}$	−6.776 (11.33)	−7.934 (11.31)	17.48 (10.62)	19.72* (11.07)
$W\,Property\ Value_{(ln)}$	3.339* (2.006)	2.412 (2.017)	1.893 (2.094)	2.638 (2.576)
Constant	−58.00 (60.35)	−195.8* (115.9)	−220.2** (91.98)	−332.4*** (98.66)
Competitive market				
$Total\ Acts_{(ln)}$	13.85*** (2.559)	−5.629 (35.93)	2.166*** (0.708)	−1.232 (6.482)
$Property\ Value_{(ln)}$	3.644** (1.812)	−12.42 (24.48)	−0.212 (1.106)	−2.290 (3.143)
$Total\ Acts \times Prop.\ Value$	—	1.727 (3.299)	—	0.388 (0.654)
$Pop_{(ln)}$	3.548 (2.315)	3.307 (2.448)	3.889** (1.846)	3.624* (1.876)
$Income_{(ln)}$	12.99* (7.665)	16.32* (8.406)	11.63 (8.116)	11.99 (8.280)
$W\,Total\ Acts_{(ln)}$	−1.091 (2.231)	−1.441 (2.210)	1.859* (0.995)	2.044* (1.128)
$W\,Pop_{(ln)}$	2.232 (3.441)	3.956 (3.675)	−4.042 (2.794)	−5.084 (3.345)
$W\,Income_{(ln)}$	−13.82 (13.21)	−16.38 (13.43)	15.58 (10.54)	18.82* (10.83)
$W\,Property\ Value_{(ln)}$	0.865 (2.540)	−0.471 (2.642)	2.069 (2.075)	2.556 (2.529)
Constant	−188.8** (76.65)	−12.49 (273.8)	−220.9** (91.62)	−220.4** (95.03)
Observations	277	277	257	257
Log likelihood	−32.87	−30.48	−137.15	−114.01
Pseudo R^2	0.890	0.898	0.455	0.547
Wald test of the interaction	—	4.09 (0.129)	—	28.19 (<0.001)
Residual I, private monopoly	0.002	−0.031	0.083	0.023
Residual I, competitive market	−0.004	−0.032	0.067	0.005

Notes: Public monopoly is the base outcome. Standard errors in parentheses; $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$. W denotes the row-standardized five-nearest-neighbor spatial lag. The Wald test is the joint test of the two interaction coefficients, with its p-value in parentheses. The residual I rows repeat the corresponding entries of Table C4 and carry the permutation p-values reported there. The specifications without the lags are those in Table ??. Stata flags 21 and 26 completely determined observations in the two 2009 columns and none in the 2019 columns, so the first wave's standard errors here should be read with that in mind.

and -0.164 (SE 0.032) on competitive structure, so neighboring activity operates at roughly a fifth of the local magnitude and on a different margin: it displaces public offices rather than reallocating between the private structures. The pattern is identical under the interaction specification: 0.036 ($p = 0.381$) on private monopoly, 0.007 ($p = 0.868$) on competitive market and -0.043 ($p = 0.046$) on public monopoly. The significance in the log-odds column and the insignificance on the two private categories are not in conflict: the first is a statement about the ratio to the public-monopoly probability, and that probability is what moves.

The second is the reliability of the 2009 columns. Twenty-one and twenty-six observations are completely determined there, with pseudo-R^2 of 0.89 and 0.90, which is quasi-complete separation and makes those standard errors untrustworthy. We therefore re-estimated the two binary contrasts with Firth's penalty, as in the robustness reported in Table **??**. In the private-monopoly contrast (195 municipalities) no lag is significant: $W\,Total\ Acts_{(ln)}$ is 1.300 (SE 1.438, $p = 0.366$), $W\,Pop_{(ln)}$ -1.674 (SE 2.415), $W\,Income_{(ln)}$ -5.371 (SE 13.660) and $W\,Property\ Value_{(ln)}$ 2.514 (SE 1.924, $p = 0.191$), while own activity remains 5.394 (SE 1.382, $p < 0.001$). In the competitive contrast (160 municipalities) no lag is significant either. The 2009 verdict survives penalization; the single marginal result in that wave does not, since the $W\,Property\ Value_{(ln)}$ coefficient of 3.339 ($p = 0.096$) in Table C5 falls to 2.514 ($p = 0.191$) once the penalty is applied. That is the expected direction under separation, and it removes the one entry in the 2009 columns that could have been read as evidence.

The third is collinearity among the lags themselves. Every lagged covariate is the same neighborhood smoothed, so the lags are more collinear with each other than the own-municipality covariates are: the correlation between $W\,Total\ Acts_{(ln)}$ and $W\,Pop_{(ln)}$ is 0.910 in 2009 and 0.893 in 2019, against 0.859 and 0.823 for the corresponding own-municipality pair. On the full regressor set the mean variance inflation factor rises from 3.16 to 5.43 in 2009 and from 2.85 to 5.07 in 2019 when the lags are added, and $W\,Pop_{(ln)}$ exceeds the conventional threshold of ten in both waves (11.41 and 10.13). This is why the population and activity lags should not be read separately: they are close to a single neighborhood-size term. It is not, however, the explanation for the imprecision of the income lag, whose standard errors of 10.6 to 11.1 in Table C5 might have suggested it. $W\,Income_{(ln)}$ has the lowest variance inflation of the four lags in both waves (2.84 and 2.80). Its standard errors are large because log municipal income varies little across municipalities, not because the lag is collinear with anything else, and the same holds for the own-municipality income coefficient in Tables **??** and **??**.

C6. What was not done

District fixed effects are not reported because they are not estimable here: a multinomial logit with eighteen district dummies on fewer than 280 municipalities is completely separated and does not converge, returning a pseudo-R^2 of 0.997 and a log-likelihood of -0.57 in the 2009 wave. Nothing should be read into that beyond the arithmetic of the sample. What is reported here is a spatial-lag-of-X specification: the neighbors' covariates enter the right-hand side, but the outcome is not allowed to depend on neighboring outcomes. That restriction is deliberate. Under a *numerus clausus* fixed centrally for each municipality, the licensing constraint is set jurisdiction by jurisdiction, so an endogenous-lag term would attribute to strategic interdependence what the allocation rule assigns administratively. A fully spatial discrete-choice estimator, such as a spatial multinomial probit carrying the autoregressive term on the outcome, would relax it, and we leave it to further work; the specifications reported here are sufficient to establish that the estimates in the body are not carrying undetected spatial structure.