Supplementary material

Who sees the sea? Sea-view potential by façade and floor from cadastral LoD1 models on the Costa del Sol. Gabriel Rojo, CASAS.ES.

This supplement has seven parts. S1 and S2 describe secondary outputs of the same rays, the rest of the horizon and the orientation of each building; they are part of the data release (`data/buildings.csv.gz`, columns `{level}_mountain_deg`, `{level}_summit_*`, `{level}_openness`, `view_facade` and `dominant_facade`), are not validated as the sea views are, and are reported here so that users of the data know what they contain. S3 compares the model with a rule-of-thumb heuristic (summarised in Sections 6.1 and 7.3 of the paper), and S4 gives the sensitivity of the geometry in the pilot, variant by variant (Section 7.4). S5 documents the blinded relabelling of the listing photographs, S6 the diagnostics of the inference with few clusters, and S7 the comparison with the LiDAR surface by floor. All figures are from the version of the model described in the paper, with every polygon of each cadastral building read as a separate block.

S1 · Mountains and openness

- **Definitions** (Section 4.4 of the paper). *Mountain*: degrees of horizon where relief beyond 2 km rises above everything nearer and at least 3° above the horizontal; the highest such point is the dominant summit, named after the nearest NGBE mountain or summit within 1 km. *Openness*: the share of rays whose obstruction within 300 m stays below 15°, a ray-based analogue of spatial openness (Fisher-Gewirtzman & Wagner, 2003) and of view access (Ko et al., 2022).
- **Pilot (Marbella, 5,231 buildings).** From the top floor, 90 % of buildings (4,731) see relief beyond 2 km rising above everything nearer over at least 4°; for 3,667 of them the dominant summit has an NGBE name within 1 km, most often La Concha (2,466 buildings), followed by Loma de la Gaceta (407), Cerro del Lastonar (355) and Loma Pierna del Caballo (284). Mean openness is 0.54 at ground level and 0.78 at the top floor.
- **Whole coast (158,060 buildings).** From the top floor, 75 % of buildings see relief beyond 2 km rising above everything nearer over at least 4°. Mean openness is 0.51 at ground level and 0.71 at the top floor.
- **Limits.** The 3° threshold is fixed, haze and visibility distance are not modelled, and neither metric has been compared with an independent reference.

S2 · Orientation against energy certificates

Two candidates are computed for the main orientation of a building. The **dominant façade** is the octant with the most free façade length, the usual geometric choice. The **view façade** is the octant whose ground-floor viewpoint has the lowest mean horizon within 2 km, that is, the most open side.

For the 359 single-family houses in the pilot whose energy performance certificate (EPC) declares an orientation, we compared the declared octant with the two geometric candidates (Table S1). The view façade matches the declared orientation exactly in 42 % of houses, against 24 % for the dominant façade and 18 % for the dominant façade of the heuristic of S3. Neither beats the constant guess "south" within ±45° (62 %), a reminder that EPC fields carry errors and conventions of their own (Hardy & Glew, 2019; Crawley et al., 2019). In a detached house with four free façades, which façade a certifier calls the orientation is partly a convention, and the evidence suggests it is often the one with the view.

Table S1. Orientation of single-family houses against the orientation declared in their energy certificate (pilot).

Predictor	Houses	Exact octant	Within ±45°
View façade (this model)	359	42 %	59 %
View façade, with trees	359	23 %	36 %
Dominant façade (longest free façade)	359	24 %	35 %
Always south	359	37 %	62 %
Heuristic of S3: dominant façade	262	18 %	30 %
Heuristic of S3: garden side	177	23 %	56 %

EPCs come from the Registro de Certificados Energéticos de Andalucía, through the CASAS.ES API (953 buildings in the pilot, 359 of them single-family).

S3 · A rule-of-thumb heuristic

The heuristic is the one the CASAS.ES website used before this model, a proxy of the kind used by property websites: it reduces each neighbour within 300 m to its centroid, looks from the roof, and searches for the sea only along a corridor to the nearest mapped beach. It is not a GIS viewshed; the ablations of Section 7.3 of the paper, without buildings or without floors, play that role, and Section 7.4 compares the tracer with a standard viewshed algorithm. On 726 buildings drawn at random in the pilot, it finds a sea view for 60 (8 %). The model finds one from the top floor for 591 (81 %), or 401 (55 %) with opaque trees.

Table S2. Sea view from the top floor, rule-of-thumb heuristic against the model (726 random buildings in the pilot).

	Model: sea	Model: no sea
Heuristic: sea	50	10
Heuristic: no sea	541	125

They disagree on 76 % of buildings, and in both directions (Figure 3b of the paper). Beyond 300 m from the beach, the heuristic finds the sea for 3 % of buildings, because it only looks towards one beach; the model finds it for 83 %, most of them on the hillside. Within 300 m, the heuristic marks 82 % with sea and the model 65 %: low buildings in the second line, behind the seafront blocks, are hidden in the model. Because the heuristic changes obstacles, observer height and search direction at once, the comparison shows that it is a different variable, not which of these changes matters. Against the listing photographs, its AUC is 0.55 (85 dwellings; Table 9 of the paper).

S4 · Sensitivity of the geometry in the pilot

Each modelling choice of Section 7.4 of the paper was changed in turn for the pilot's 5,231 buildings. Table S3 gives its effect on buildings and Table S4 carries it to dwellings, to the pilot's income gradient and, where the variant was also traced for the listing buildings, to the photographs. The field-of-view rows read the same rays as the default within ±60° or ±45° of the façade normal; the recessed viewpoint is traced again with the viewpoint 0.1 m from the wall and its own block opaque

beyond 1 m. The consequential variants were then run for the whole coast (Tables 10 and A4 of the paper).

Table S3. Sensitivity of the geometry in the pilot (5,231 residential buildings): share that see the sea (4° or more) or have an open view (20° or more) from the ground and top floors (Eq. 2), and share whose class changes against the default at either threshold.

Variant	Ground floor, sea	Ground floor, open view	Top floor, sea	Top floor, open view	Changes class, ground floor	Changes class, top floor
Default (Table 3)	56.9 %	36.1 %	77.7 %	62.0 %	0.0 %	0.0 %
Numerical: rays every 1°	55.6 %	35.5 %	76.6 %	61.5 %	1.7 %	1.4 %
Numerical: near-field step max(0.5 m, 0.002 d)	56.6 %	35.8 %	77.5 %	61.7 %	0.4 %	0.3 %
Numerical: sea mask without a 16 m edge strip	56.9 %	36.1 %	77.7 %	62.0 %	0.0 %	0.0 %
Structural: field of view ±60°	49.2 %	32.4 %	71.6 %	56.0 %	7.7 %	6.1 %
Structural: field of view ±45°	46.5 %	29.7 %	68.4 %	51.8 %	10.5 %	10.2 %
Metric: sea extent weighted by distance (relative to 1 km)	41.1 %	15.7 %	66.6 %	36.3 %	20.4 %	25.7 %
Structural: viewpoint on the plane of the wall, own block opaque beyond 1 m	53.7 %	34.1 %	76.8 %	61.7 %	5.0 %	2.7 %

Variant	Ground floor, sea	Ground floor, open view	Top floor, sea	Top floor, open view	Changes class, ground floor	Changes class, top floor
Structural: own block transparent only within 4 m	56.7 %	35.9 %	77.6 %	61.8 %	0.3 %	0.2 %
Structural: three viewpoints along the longest edge of each octant	59.7 %	37.9 %	79.0 %	64.2 %	2.8 %	2.2 %
Structural: floors of a façade within 1.5 m, not 6 m	56.9 %	36.1 %	71.2 %	52.2 %	0.0 %	9.7 %
Structural: flat roofs and footprint-level eye base	58.9 %	38.3 %	78.8 %	63.9 %	3.5 %	3.6 %
Structural: LiDAR heights where available	56.6 %	35.7 %	76.1 %	61.6 %	3.0 %	4.1 %
Vegetation: LiDAR trees at half height	34.6 %	17.3 %	73.5 %	56.2 %	22.3 %	5.8 %
Vegetation: LiDAR trees, opaque	12.1 %	3.9 %	55.4 %	36.1 %	44.8 %	25.9 %
Combined: 1.5 m radius + own block opaque beyond 4 m + flat roofs	58.7 %	38.1 %	72.8 %	54.4 %	3.6 %	10.5 %

Variant	Ground floor, sea	Ground floor, open view	Top floor, sea	Top floor, open view	Changes class, ground floor	Changes class, top floor
Combined, with LiDAR trees at half height	40.0 %	22.4 %	65.8 %	46.3 %	19.5 %	17.9 %
Combined, with opaque LiDAR trees	16.4 %	6.2 %	41.9 %	23.6 %	41.6 %	39.6 %

Table S4. The same variants carried to dwellings and to the two references: share of the pilot's 37,708 dwellings on a floor with open-view potential under the best-façade and length-share rules (Eqs. 3 and 4, nearest floor); income coefficient across the pilot's 49 census sections (log income and log distance, HC1 standard error in brackets); and AUC against the first-round labels of the listing photographs, scored as in Table 9 of the paper (from the coast runs of the same variant, or re-traced for the 92 buildings, whose default re-traced the same way scores 0.81; — where neither was run).

Variant	Dwellings, open: best façade	Dwellings, open: length share	Income coefficient, pilot	AUC, photographs
Default (Table 3)	50.4 %	27.6 %	120 (28)	0.81
Numerical: rays every 1°	50.2 %	27.1 %	119 (29)	—
Numerical: near-field step max(0.5 m, 0.002 d)	50.0 %	27.3 %	121 (29)	—
Numerical: sea mask without a 16 m edge strip	50.4 %	27.5 %	120 (28)	—
Structural: field of view ±60°	45.6 %	19.9 %	123 (25)	0.79
Structural: field of view ±45°	42.0 %	15.2 %	125 (24)	0.78
Metric: sea extent weighted by distance (relative to 1 km)	34.2 %	17.5 %	99 (32)	0.82

Variant	Dwellings, open: best façade	Dwellings, open: length share	Income coefficient, pilot	AUC, photographs
Structural: viewpoint on the plane of the wall, own block opaque beyond 1 m	48.7 %	24.5 %	118 (25)	0.82
Structural: own block transparent only within 4 m	49.8 %	26.5 %	122 (26)	0.81
Structural: three viewpoints along the longest edge of each octant	52.6 %	27.5 %	119 (28)	0.81
Structural: floors of a façade within 1.5 m, not 6 m	45.4 %	24.3 %	112 (31)	0.83
Structural: flat roofs and footprint-level eye base	53.2 %	28.8 %	121 (25)	—
Structural: LiDAR heights where available	49.7 %	26.7 %	126 (28)	—
Vegetation: LiDAR trees at half height	44.7 %	24.0 %	85 (25)	0.83
Vegetation: LiDAR trees, opaque	33.5 %	17.7 %	32 (20)	0.80
Combined: 1.5 m radius + own block opaque beyond 4 m + flat roofs	47.0 %	24.1 %	123 (25)	—
Combined, with LiDAR trees at half height	42.3 %	20.5 %	92 (22)	0.81
Combined, with opaque LiDAR trees	30.2 %	14.3 %	26 (17)	0.81

S5 · Blinded relabelling of the listing photographs

Why. In the first labelling (Section 7.3 of the paper), three AI models and the author labelled the photographs and only the final label of each listing was recorded, so its agreement cannot be computed, and the author who decided ties built the model. The relabelling documents a protocol,

keeps every individual label, removes access to anything but the photographs, and adds a second developer, so that the reference does not rest on the models of one developer.

Material. 148 listings, each under an anonymous code assigned in a random order: L001–L115, the 115 clearly labelled listings (seed 20261022), and L116–L148, 33 of the 34 listings the first round left undetermined (seed 110; one had no photographs left). For each, contact sheets with every photograph as a numbered thumbnail, in the order the seller published them (up to 30 per sheet), and each photograph at 640 × 480. Sheets made during the first labelling were not used, because the enlargements chosen then could reveal its judgement. No listing text, price, address, map position, model output or earlier label was available, and the labellers were told not to open any other file.

Labellers. Claude Opus 5.5 (claude-opus-5-5), Claude Sonnet 5 (claude-sonnet-5) and Claude Haiku 4.5 (claude-haiku-4-5-20251001), Anthropic, run on 22 October 2026 as separate agents on batches of 16–29 listings that opened the contact sheets and enlarged the photographs they needed; Gemini 3.1 Pro (requested as gemini-pro-latest, served as gemini-3.1-pro-preview) and Gemini 3.8 Flash (gemini-3.8-flash), Google, run on the same day through the Gemini API, one request per listing with the protocol and every photograph of the listing inline, numbered, and a structured answer with the same fields. They are models of two developers, so their errors need not be independent, and none is a human judgement.

Protocol (released as `validation/PROTOCOL.md`, with the instruction as `validation/PROMPT.txt` and the script for the Gemini requests). Question: is the sea visible from the dwelling that is for sale? Labels: SEA, the sea clearly visible from the dwelling itself (windows, balconies, terraces, private garden or roof terrace); PARTIAL, only a sliver or a far, small sea; NO_SEA, the photographs show the dwelling's outside views and no sea; UNDETERMINED, interiors only, renderings or views not from the dwelling. Rules: only the sea counts (not pools, lakes, rivers or golf courses); promotional photographs of the area do not count; if the sea is visible only from a communal area, label what the dwelling itself sees and flag it (yes, no or unclear); decide each listing on its own. Also recorded: the photographs that show the sea, whether the photographs are mostly renderings, and a confidence.

Reference and ties. One vote per developer: the majority of the three Anthropic models (positive, negative, undetermined, or split) and Gemini 3.1 Pro; Gemini 3.8 Flash is a sensitivity check. Agreement, or a single deciding developer, gives the reference; a disagreement, or a split Anthropic majority with a deciding Gemini, is a tie. There were 15 ties, all a sliver of sea seen by Gemini and not by the Anthropic models. The author decided them from one page per listing with the full contact sheet and, enlarged, the photographs in which any labeller had marked sea, without saying which, and without the model's scores; under the protocol, 6 were labelled PARTIAL and 9 NO_SEA. The author built the model, so these labels are not blind to it in the way the AI labels are; leaving out the 9 ties in the main sample (6 without sea, 3 partial) gives an AUC of 0.81 (84 dwellings).

Results (Table S5). The four labellers that vote agree on 68 % of the 148 listings across the four labels (Fleiss' κ 0.76) and on 81 % for sea, partial included, against no sea among the 110 that all four decide (κ 0.78); the three Anthropic models alone give κ 0.79 and 0.85. Gemini 3.1 Pro labels partial more often than the others (37 listings against 25–32 for the Anthropic models), which is why the AUC against its labels alone is the lowest; Gemini 3.8 Flash agrees more with the Anthropic models. The reference agrees with the first labelling on 92 % of the listings both call determinate (Cohen's κ 0.84). The released validation table gives every label, under the anonymous code of the release.

Table S5. Blinded relabelling of 148 listings (the 115 clearly labelled ones and 33 of the first round's

undetermined ones) by five AI models of two developers, each run separately under the protocol of this section, against the final labels of the first round. Labels: sea, partial, no sea, undetermined. Undet.: undetermined. Agree (agreement with the first round) and Cohen's κ are for sea (partial included) against no sea, on the listings both call determinate; the next two columns count the disagreements in each direction. AUC: the default model's score against each set of labels in the main sample, on the dwellings that set decides. Reference: one vote per developer (the Anthropic majority and Gemini 3.1 Pro), ties decided by the author. * The Anthropic majority is sea (partial included) or no sea; its 69 positives are not split.

Labeller	Sea	Partial	No sea	Undet.	Agree	κ	First no, blinded yes	First yes, blinded no	AUC (n)
First round (115 listings)	53	12	50	—	—	—	—	—	0.81 (93)
Claude Opus 5.5	39	32	47	30	93.9 %	0.88	6	1	0.81 (93)
Claude Sonnet 5	41	31	43	33	95.5 %	0.91	5	0	0.81 (89)
Claude Haiku 4.5	38	25	50	35	93.5 %	0.87	3	4	0.81 (87)
Gemini 3.1 Pro	48	37	34	29	84.3 %	0.67	18	0	0.76 (93)
Gemini 3.8 Flash (sensitivity)	48	19	51	30	93.0 %	0.86	5	3	0.84 (93)
Anthropic majority	69 *	*	46	33	95.5 %	0.91	4	1	0.82 (90)
Reference	40	36	43	29	92.2 %	0.84	9	0	0.80 (93)

S6 · Inference with few, unequal clusters

Design. 860 census sections in 21 municipalities; Málaga holds 51 % of the sections and 32 % of the dwellings. The partial leverage of a municipality is its share of the variation of log income that remains after the other regressors, the quantity that drives the weight of each cluster in the coefficient (MacKinnon et al., 2023b). The effective number of clusters G* follows Carter et al. (2017) with no within-cluster correlation; with G* near 2, cluster-robust inference at the municipality level is unreliable whatever the method.

Table S6. Cluster diagnostics for the income coefficient (main outcome): partial leverage of each municipality (its share of the residual variation of log income after the other regressors) and the coefficient without that municipality, for M2 and M5. G* = 2.2 (M2) and 2.3 (M5).

Municipality	Partial leverage, M2	Coefficient without it, M2	Partial leverage, M5	Coefficient without it, M5
Málaga	0.66	73.5	0.65	16.6
Torremolinos	0.08	52.3	0.07	14.6

Municipality	Partial leverage, M2	Coefficient without it, M2	Partial leverage, M5	Coefficient without it, M5
Marbella	0.07	50.8	0.08	18.1
Vélez-Málaga	0.04	52.1	0.04	18.5
Benalmádena	0.04	51.9	0.02	17.1
Mijas	0.02	51.8	0.04	15.9
Alhaurín de la Torre	0.02	53.9	0.02	17.5
Fuengirola	0.02	52.0	0.03	18.8
Rincón de la Victoria	0.02	50.9	0.02	14.6
Estepona	0.02	50.7	0.01	16.6
Manilva	0.01	51.9	0.01	15.8
Alhaurín el Grande	0.00	52.6	0.00	16.0
Benahavís	0.00	52.3	0.00	16.4
Nerja	0.00	52.5	0.00	16.9
Torrox	0.00	52.4	0.01	16.6
Algarrobo	0.00	52.3	0.00	16.4
Casares	0.00	52.3	0.00	16.6
Frigiliana	0.00	52.5	0.00	16.5
Istán	0.00	52.5	0.00	16.5
Ojén	0.00	52.5	0.00	16.6
Totalán	0.00	52.5	0.00	16.5
Full sample	—	52.5	—	16.5

The M5 discrepancy. With building controls, the coefficient hardly moves when any municipality is dropped (Table S6), so CV3, which is built from those leave-one-out estimates, is close to the clustered error, and its t-based interval excludes zero. WCR imposes the null hypothesis and resamples municipality residuals with Webb weights; with Málaga holding about two-thirds of the partial leverage, its bootstrap distribution of t is wide and the test is conservative, while the unrestricted bootstrap (WCU) is narrow and over-rejects. Neither method is reliable with $G^* \approx 2$ (MacKinnon et al., 2023a), so the paper reports both and draws no firm conclusion about M5.

A p-value that is not monotone. For the terrain-only coefficient of any sea with the best façade (Tables A3 and A4 of the paper, marked †), the WCR p-value at zero is 0.017, but it rises again below zero: values from -8.7 to about -1.3 are not rejected (p up to 0.09), and neither are values from about 0.7 to 20.3. With $G^* \approx 2$, the bootstrap distribution of the t statistic changes with the hypothesised value, and the set of values not rejected need not be an interval. The p-values on a grid of hypothesised values are released with the data.

Spatial dependence. Municipal clustering does not cover correlation between neighbouring sections of different municipalities. Conley (1999) errors with a Bartlett kernel on the distance between dwelling-weighted section centroids are reported in Table A1 and Table 10 of the paper. Square spatial blocks were resampled with replacement, with the whole block entering each replicate, and the regression with municipality fixed effects re-estimated (Table S7). Moving the grid by half a block moves the bounds by up to about 15 points, mostly the upper ones; none of the intervals

includes zero. Draws in which a municipality fixed effect cannot be estimated are skipped, so each interval rests on 311–478 of the 499 draws.

Table S7. Spatial-block bootstrap intervals for the M2 income coefficient (main outcome) by block size and grid origin (499 draws; blocks are resampled with replacement and every section in a drawn block enters with its block; municipality fixed effects are re-estimated in each replicate, and draws in which one cannot be estimated are skipped; in brackets, blocks and draws kept).

Block side	Origin (0, 0)	Origin shifted half a block east	Shifted half a block north	Shifted both ways
2.5 km	37.3 to 67.7 (125 blocks, 439 draws)	38.4 to 70.1 (131 blocks, 475 draws)	36.7 to 70.7 (125 blocks, 439 draws)	37.8 to 70.5 (130 blocks, 478 draws)
5 km	37.1 to 70.4 (65 blocks, 426 draws)	40.8 to 69.9 (65 blocks, 431 draws)	35.2 to 82.5 (67 blocks, 446 draws)	39.6 to 74.8 (63 blocks, 440 draws)
10 km	40.3 to 91.2 (30 blocks, 440 draws)	42.5 to 76.1 (27 blocks, 440 draws)	42.1 to 82.0 (29 blocks, 311 draws)	39.7 to 76.6 (28 blocks, 322 draws)

Functional form. The fractional logit (Papke & Wooldridge, 1996) is a binomial quasi-likelihood with a logit link and weights equal to dwellings, with municipality fixed effects; the average marginal effect of log income is 58.7 percentage points with buildings and 1.6 on terrain alone, against 52.5 and 3.2 in the linear model, so the share of the coefficient that remains on terrain alone is smaller, not larger, than in the linear model.

S7 · The LiDAR surface by floor

Table S8 splits the comparison of Section 7.2 of the paper, with the full LiDAR building surface, by traced floor. The construction of the reference surface is described there; the code is `vistas_v3.py --lidar <building nDSM> --sup-lidar` (with `--sup-lidar-huellas` for the surface restricted to the cadastral parts and `--sup-lidar-filtro 50` for the one without small patches).

Table S8. The model against the full LiDAR building surface (no trees) by traced floor: mean degrees of sea, model / reference, mean difference and agreement at 20° (κ), pilot viewpoints.

Traced floor	Viewpoints	Mean sea, model / reference (°)	Mean difference (°)	Same class at 20° (κ)
Ground floor	26,178	12.3 / 7.8	4.5	91.9 % (0.72)
Middle floor	8,787	11.1 / 7.6	3.5	93.2 % (0.75)
Top floor	18,006	25.2 / 19.8	5.4	90.7 % (0.80)

References

Carter, A. V., Schnepel, K. T., & Steigerwald, D. G. (2017). Asymptotic Behavior of a *t*-Test Robust to Cluster Heterogeneity. *The Review of Economics and Statistics*, 99(4), 698–709. https://doi.org/10.1162/rest_a_00639

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics*, 92(1), 1–45. https://doi.org/10.1016/s0304-4076(98)00084-0

Crawley, J., Biddulph, P., Northrop, P. J., Wingfield, J., Oreszczyn, T., & Elwell, C. (2019). Quantifying the Measurement Error on England and Wales EPC Ratings. *Energies*, 12(18), 3523. https://doi.org/10.3390/en12183523

Fisher-Gewirtzman, D., & Wagner, I. A. (2003). Spatial Openness as a Practical Metric for Evaluating Built-up Environments. *Environment and Planning B: Planning and Design*, 30(1), 37–49. https://doi.org/10.1068/b12861

Hardy, A., & Glew, D. (2019). An analysis of errors in the Energy Performance certificate database. *Energy Policy*, 129, 1168–1178. https://doi.org/10.1016/j.enpol.2019.03.022

Ko, W. H., Kent, M. G., Schiavon, S., Levitt, B., & Betti, G. (2022). A Window View Quality Assessment Framework. *LEUKOS*, 18(3), 268–293. https://doi.org/10.1080/15502724.2021.1965889

MacKinnon, J. G., Nielsen, M. Ø., & Webb, M. D. (2023a). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, 232(2), 272–299. https://doi.org/10.1016/j.jeconom.2022.04.001

MacKinnon, J. G., Nielsen, M. Ø., & Webb, M. D. (2023b). Leverage, influence, and the jackknife in clustered regression models: Reliable inference using summclust. *The Stata Journal*, 23(4), 942–982. https://doi.org/10.1177/1536867x231212433

Papke, L. E., & Wooldridge, J. M. (1996). Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics*, 11(6), 619–632. https://doi.org/10.1002/(sici)1099-1255(199611)11:6%3C619::aid-jae418%3E3.0.co;2-1