When Does a Flood Map Become a Treatment? An Outcome-Blind Economic Admission Audit in Kuwait

Measurement and design audit with exploratory proxy associations

Mallooki

August 2026

Working paper | Measurement and research design

Abstract

Satellite data can measure environmental change where conventional records are thin, but a physical map becomes an economic treatment only when it aligns with the outcome geography and the mechanism being estimated. We develop an outcome-blind economic admission audit and apply it to Kuwait's November 2018 floods. A conservative Sentinel-2 rule identifies candidate persistent open water in 10 of 541 polygons linked through a reviewed historical geography bridge, below an author-defined feasibility floor of 20 exposed units. Rebuilding the same rule before that bridge mask finds positives in 15 of 232 held-out polygons, including Sabah Al-Ahmad. The binding problem is therefore not a Kuwait-wide absence of candidate water; it is the inability to connect much of the signal to stable housing-market units. Separate Sentinel-1 candidates reach 49, 26, and 11 linked polygons, but none passes the frozen radar-optical agreement rule, and exact-time validation of the late-November state is unavailable in the public sources audited through 12 August 2026. The confirmatory causal design stops there. A post-audit appendix then reports the full 36-cell radar-proxy family: only 12 cells are estimable, and alternative contrast definitions select different geographies and yield imprecise completed-home associations of different signs. The Kuwait case shows what standard map-accuracy checks do not: economic salience, administrative coverage, assignment-unit support, and outcome design determine whether pixels identify a causal parameter or only a proxy association.

1. Introduction

The central problem in a flood-capitalization study is not whether visibly affected places have different property prices. Elevation, drainage, coastal access, urban age, and local investment can affect both inundation and values long before a storm. The useful question is whether an unusual flood changes registered transaction prices in more exposed locations relative to comparable places that would otherwise have followed similar paths.

The mechanism is plausible in Kuwait. A rare rainfall sequence may reveal which roads drain slowly, which underpasses become impassable, and which structures face costly repairs. Buyers can update

beliefs about future losses, daily access, public mitigation, and resale difficulty. Direct damage can change current condition at the same time. Yet a satellite observes neither beliefs nor transactions. It records a sensor response to a physical surface at a particular time. A dark radar pixel may be calm water, pavement, roof shadow, or another low-backscatter surface; narrow street flooding can also disappear before the pass or brighten through urban double bounce (Lin et al., 2019). The measurement problem is part of identification, not a preliminary cleaning task.

Flood-capitalization evidence makes the distinction consequential. Realized disasters can change insurance demand and prices through learning, but responses differ with disclosure, prior beliefs, damage, coastal amenities, market selection, and time since the event (Gallagher, 2014; Beltrán et al., 2018; Bakkensen and Barrage, 2022; Hino and Burke, 2021). U.S. results do not transport mechanically to Kuwait's state-centered housing system or its registered private-transfer market. The question is therefore not whether the international literature contains a plausible sign. It is whether the Kuwait data measure the local treatment that the mechanism requires.

Remote-sensing accuracy assessment already distinguishes the map, the reference sample, the response design, and the estimator (Stehman and Czaplewski, 1998; Olofsson et al., 2014). Economics adds another problem. Satellite products can expand empirical work precisely where administrative measures are scarce, but measurement error can bias causal estimates and interact with the research design (Donaldson and Storeygard, 2016; Jain, 2020). Recent econometric work studies parameter recovery and downstream causal inference once a remotely sensed variable enters estimation (Proctor et al., 2026; Alsharif et al., 2026). This paper asks the prior urban-treatment question: whether the product measures the event-time state, maps to market assignment units, and carries the economic salience required to enter that estimation. A pixel-level confusion matrix does not answer those questions.

This paper turns that gap into an **economic admission audit**. The audit asks five questions in order: Was the product constructed reproducibly? Does it measure the intended physical state at the relevant time? Is that state economically meaningful to the mechanism? Can it be assigned across the administrative units observed in the outcome data? Does the resulting variation support the intended counterfactual and inference? Passing all five supports a causal treatment claim. Failing one does not ban estimation, but it moves the analysis to a weaker evidence tier and changes the language attached to the coefficient.

The Kuwait application produces a sharper result than the original design anticipated. Within the 541-polygon Central Statistical Bureau (CSB)-Ministry of Justice (MOJ) bridge approved for analysis, a frozen Sentinel-2 machine rule identifies candidate persistent open water in only 10 polygons. That fails the author-defined minimum of 20 exposed assignment units. But the source geography contains another 232 held-out polygons. An outcome-blind coverage audit finds optical positives in 15 of them, including

the prominently affected Sabah Al-Ahmad location discussed by Alghais (Alghais, 2019). The ten-polygon result is therefore conditional on the reviewed geography bridge, not a statement that only ten Kuwaiti places flooded.

Radar does not solve the problem. The three early-sequence Sentinel-1 candidates contain positives in 49, 26, and 11 approved polygons. The first two clear a crude count of 20, so the optical count is not a mathematical ceiling on radar support. None, however, clears the frozen cross-sensor machine-reference rule. The public record also does not provide a probability sample of positive and negative ground observations close enough to the 25 November radar pass to validate that later state. The early optical and radar diagnostics share evidence and are separate gates, not statistically independent proofs.

The achieved research object is design feasibility: whether the audited public-data candidates satisfy the measurement and support conditions required for a causal transaction-price estimand. They do not. This closes the confirmatory Sentinel-1/Sentinel-2-CSB/MOJ implementation, while leaving room for explicitly exploratory associations and for a new causal study built around a different exposure product, a repaired historical geography, or new validation evidence. The contribution is a worked economics application of an assignment-unit admission framework. Kuwait is the demonstration, not an impossibility theorem about public imagery or flood mapping.

The admission decision itself was outcome-blind; the completed paper is not, because a later appendix opens prices under an explicitly exploratory protocol. That appendix retains all 36 contrast-threshold-market-form cells. Twelve are estimable. For the permissive completed-home binary proxies, the associations are -0.020, 0.094, and 0.125 log points across the three contrast definitions, with wide intervals. The corresponding proxy-positive block sets barely overlap, so this pattern is joint timing-and-geography sensitivity rather than a clock-only comparison. It is the realized descriptive finding, not a causal flood estimate.

2. Kuwait’s market and the flood-capitalization mechanism

2.1 What the transaction market represents

Kuwait’s housing system is not a conventional mortgage-led private market. State land allocation, public housing, subsidized finance, and the expectation of government provision shape household options and urban expansion. Alshalfan, Alrasheed, and Albabtain describe a system in which housing operates partly as a channel for distributing oil wealth, the public waitlist is long, and private single-family housing is expensive and supply constrained (Alshalfan et al., 2022). Recent work also links public-housing delays

to spatially heterogeneous private-market prices, underscoring that policy expectations can enter the same outcomes a flood study would use (AlBader and AlSabah, 2025).

The empirical population in this paper is narrower than Kuwait's housing stock. The MOJ source records registered private-residential transfers obtained from the ministry's official live query (Kuwait Ministry of Justice, 2026b). It does not observe all residents, public-housing allocations, informal occupancy, citizenship, compensation receipt, financing, or buyer expectations. It also lacks a stable parcel or legal-deed identifier and exact coordinates. The paper therefore studies the feasibility of a design for registered transfers, not welfare or housing consumption across Kuwait.

The source distinguishes recorded land from completed homes, but both labels are administrative proxies. "Recorded land" is not independently verified as vacant at sale. A completed-home price divided by plot area is a registered value per square meter of plot, not a structure-adjusted price per square meter of built space. Structure age, quality, renovation, built area, frontage, and parcel history are missing. These constraints also complicate constant-quality price measurement in Kuwait (Alfalah et al., 2023). A valid exposure would not remove them; it would merely allow the downstream outcome and selection gates to be tested.

2.2 From rainfall to a market treatment

The economic chain has five links:

1. an unusual rainfall sequence produces physical inundation;
2. inundation causes direct damage, service interruption, or access disruption;
3. households and buyers observe or learn about those consequences;
4. current condition and beliefs about future loss, public repair, and compensation change; and
5. prices, sale probability, and the composition of registered transfers respond.

The satellite candidates measure only part of the first link. Persistent open water at a pass may proxy for drainage problems, but it does not establish damage, route disruption, or buyer salience. Short-lived street water, basement flooding, and damaged utilities may be economically important while absent from an open-water layer. Conversely, retained water in a peripheral depression can be accurately classified without mattering to residential buyers.

The intended binary estimand was a within-neighborhood post-event contrast in log registered price per plot square meter between blocks admitted as meaningfully exposed and blocks admitted as zero or very low exposure. A continuous dose response was a separate estimand, not another label for the binary effect. The planned transaction-level specification would have included block fixed effects and neighborhood-by-time effects. Its coefficient would be weighted by observed sales and would describe registered transaction prices conditional on a transfer. Interpreting it as a property-stock value effect

would additionally require stable transaction selection and composition, or a design that models those margins jointly. Any-sale, transfer count, plot area, and transferred-share outcomes were therefore co-primary diagnostics rather than optional robustness checks.

Even with valid exposure, the reduced form would combine damage, information, access disruption, compensation expectations, and mitigation. It would not separately identify beliefs, the probability of loss, or household welfare. Shared roads and national news also create interference: a nominally dry block can receive information or lose access because a neighboring route floods. The appropriate estimand is a local exposure bundle conditional on the chosen mapping, not a pure hydrological parameter.

2.3 A sequence of events and responses

November 2018 contained several rain episodes rather than one clean shock. Alghais centers the first major event on 9 November and reports roughly 111 millimeters in under six hours, with case evidence from Kuwait City, Fahaheel, and Sabah Al-Ahmad (Alghais, 2019). Al-Hemoud et al. model the broader 4-14 November sequence on a 50-meter hydraulic grid and report spatially uneven hazard (Al-Hemoud et al., 2023). Official reporting records government mobilization, road closures, shelters, drainage, and reopening around the two main waves (Kuwait News Agency, 2018c; Kuwait News Agency, 2018b; Kuwait News Agency, 2018a; Kuwait News Agency, 2018d). These sources establish public disruption and response; they do not assign 20-meter treatment.

Government action is part of the plausible treatment bundle. Compensation intent was announced on 11 November, a fact-finding response followed on 13 November, and later claims and repair processes could alter expected losses. Aggregate announcements do not reveal block-level benefits, but controlling them away mechanically would be inappropriate in a total-event reduced form. Their timing also means that any eventual price response could reflect both revealed vulnerability and expected state protection.

The sensor clock differs from the public-event clock. Sentinel-1 acquired north and south scenes on 1, 13, and 25 November at about 14:50 UTC. Sentinel-2 observed the analysis area on 12 November at 07:31 UTC, roughly 31 hours before the early radar pass. The 13 November radar comparison captures the cumulative early-November sequence, not an isolated 9 November pulse. Later rainfall around 23-24 November prevents the 22 and 27 November optical observations from cleanly bracketing the 25 November radar state (Al-Yamani, 2021). Figure 1 keeps those clocks separate.

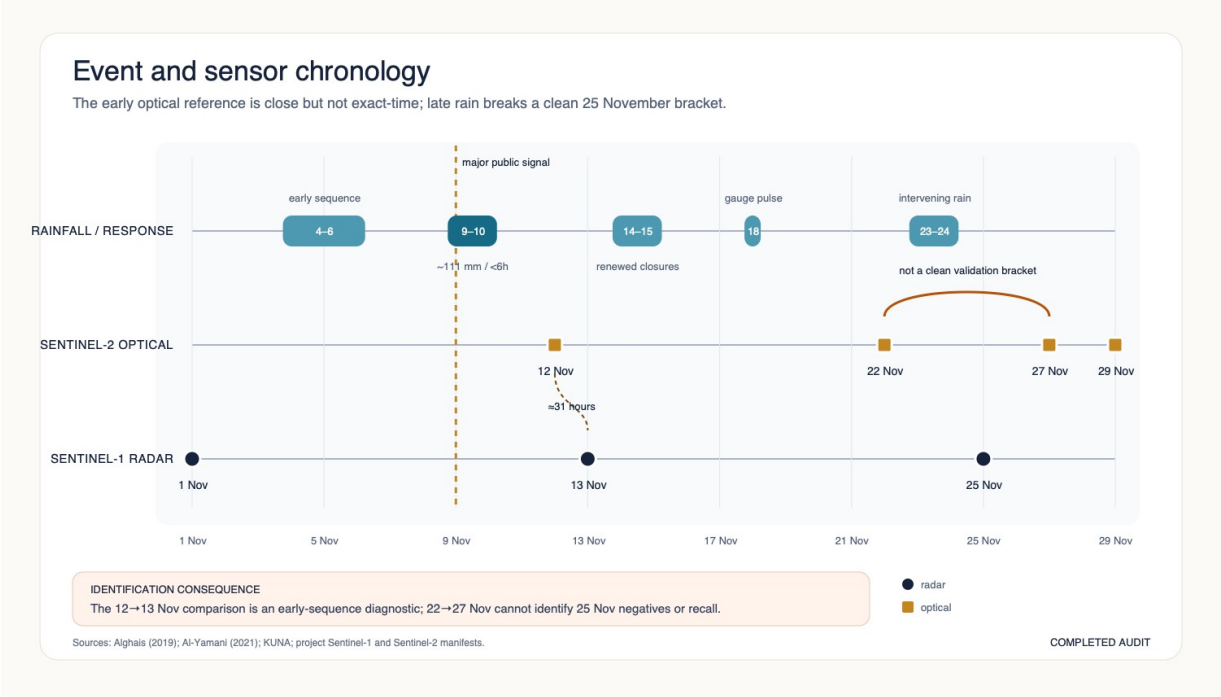

Figure 1. Public-event and satellite chronology. The 12 November optical observation is about 31 hours before the early radar pass; additional late-month rain prevents a clean optical bracket around 25 November.

Notes: The chronology uses the event literature, KISR rainfall compilation, official KUNA reports, and frozen satellite manifests (Alghais, 2019; Al-Yamani, 2021; Kuwait News Agency, 2018b; Kuwait News Agency, 2018a; Kuwait News Agency, 2018d). Official disruption and response reports are contextual evidence, not cell labels.

3. From map accuracy to economic admission

3.1 Five layers of evidence

Table 1 separates standard map questions from the additional work required when a raster becomes a causal regressor. The distinction is sequential. Technical reproducibility cannot substitute for target-state truth. Physical validity cannot establish economic salience. A salient exposure can still fail because it cannot be joined to outcomes or appears in too few assignment units. And geographic support does not prove parallel counterfactual paths or valid inference.

Table 1. The economic admission audit

Layer	Question	Evidence needed	Kuwait status
Technical and provenance QA	Were the intended scenes, calibration, grid, masks, and rules used reproducibly?	Scene receipts, processing hashes, coverage and seam checks	Passed for the audited products

Layer	Question	Evidence needed	Kuwait status
Physical target validity	Does the candidate measure the same surface state at the relevant time?	Temporally compatible reference design, error matrix, spatially valid uncertainty	Early radar rules fail the machine-reference rule; late exact-time truth unavailable
Economic relevance and salience	Does the physical state proxy damage, access, or information observed by market participants?	Incident, claims, damage, access, or salience evidence linked to assignment units	Not established
Administrative coverage and assignment support	Can exposure be joined to stable outcome units with dispersed roles?	Crosswalk audit, exposed and comparison counts, within-market variation	Approved bridge fails for the optical proxy; held-out positives expose a coverage problem
Outcome and causal design	Do supported units identify the stated selected-transaction estimand?	Outcome support, pre-event comparability, selection and interference analysis, credible inference	Not reached

Notes: The first two rows build on thematic map-accuracy and satellite-measurement practice (Stehman and Czaplewski, 1998; Olofsson et al., 2014; Donaldson and Storeygard, 2016; Jain, 2020). The last three are the economic admission extension. “Passed” at an early layer is necessary, never sufficient, for a housing regression.

3.2 Author-defined gates, not universal standards

The original workflow fixed numerical rules without reading housing prices. A candidate classifier had to clear precision of 0.80 and recall of 0.70 relative to the relevant reference, with the lower bounds of the reported 95 percent intervals also above those thresholds. A promoted exposure then needed at least 20 meaningfully exposed blocks, 20 comparison blocks, and 10 reviewed neighborhoods containing both roles. At least 10 blocks in each role had to meet the transaction-support rule, and no one canonical neighborhood could contribute more than 60 percent of exposed analysis blocks.

These cutoffs are conservative author-defined admission rules. They are not field standards, identification theorems, or power guarantees. A coefficient can be algebraically estimable with 10 treated blocks and still be too leveraged for credible inference; 20 exposed blocks do not guarantee overlap, enough clusters, or a plausible counterfactual. The 20/20 rule was a procedural floor intended to stop sparse, location-dominated designs before outcome inspection. Any future paper should supplement it with leverage diagnostics, design simulations, and a minimum-detectable-effect calculation tied to the actual cluster structure.

The rules were frozen in August 2026, after the event but before opening the Sentinel-1 comparison and before any exposure-outcome merge. “Outcome-blind” is the correct claim. It does not mean the design was publicly preregistered before imagery discovery, and it does not protect against every measurement-side choice. One later optical-radar intersection is accordingly reported below as a post-diagnostic sensitivity, not a prespecified scientific gate.

3.3 Identification after admission

Had a layer passed, the identifying variation would have been within-neighborhood differences across blocks over time. Block effects would absorb stable site attributes; neighborhood-by-time effects would absorb housing-market shocks common within the local area. The main assumption would be that, absent admitted local exposure, supported exposed and comparison blocks would have followed comparable registered-transfer paths. That assumption would still be incomplete without stable selection into sale, no exposure-correlated block shock, a defensible information date, an explicit treatment mapping for shared roads and repairs, and sufficient within-neighborhood variation.

Parallel-trend equivalence tests, leave-one-neighborhood diagnostics, small-cluster inference, and volume/composition outcomes were coded for a future validated exposure. Passing those checks would not repair treatment measurement; failing measurement appropriately stops the design earlier. This order also prevents a visually attractive threshold from being selected because it happens to produce a plausible event study.

4. Data, construction, and geography

4.1 Optical and radar candidates

The radar inventory contains six Sentinel-1A Interferometric Wide Swath acquisitions: north and south tiles on 1, 13, and 25 November 2018, all on ascending relative orbit 174 with VV and VH polarization. The production layer uses radiometrically terrain-corrected gamma-naught intensity from the Microsoft Planetary Computer and a common 20-meter grid (Microsoft Planetary Computer, 2026). The audited rules compare backscatter across dates and polarizations, remove small components, and preserve seam and coverage diagnostics. Same-sensor change produces candidates, not truth.

The conservative optical machine reference uses Sentinel-2 Level-2A imagery from 12 November and clear baselines from 18, 20, and 30 October. A positive cell must be classified as water by the Scene Classification Layer, exceed frozen NDWI and MNDWI thresholds, have dry support on at least two October dates, lie outside a buffer around baseline water, and belong to a contiguous component of at least 1,600 square meters (McFeeters, 1996; Xu, 2006; European Space Agency, 2022). This construction

favors persistent, spectrally clean, spatially contiguous open water. It is likely to miss short-lived street water, narrow underpasses, and mixed urban pixels.

The optical rule was frozen before inspecting the radar comparison and without housing outcomes. It is still a machine reference observed about 31 hours before the early radar pass, not human-adjudicated exact-time truth. Cross-sensor agreement therefore measures concordance with that specific physical proxy. It cannot validate the buyer-information or damage links in the economic chain.

4.2 The historical geography bridge

The CSB 2011 source contains 773 lower-area polygons (Kuwait Central Statistical Bureau, 2011). A reviewed crosswalk approves 541 polygons across 89 district correspondences for the MOJ analysis and holds out 232 polygons across 49 source district labels because of changed boundaries, merged units, or unresolved name conflicts. Seven held-out northern polygons fall outside the frozen optical grid; 225 are represented.

This distinction was missing from the first complete draft and materially changes the scope of the result. The canonical optical label raster was hard-masked to the 541 approved polygons, so it could not reveal what occurred in held-out areas. A bounded audit rebuilt the identical rule from the cached optical inputs before applying that mask. It reproduced the approved labels exactly and then aggregated the same candidate class over all source polygons. The audit did not read transactions or promote the held-out signal into treatment.

MOJ records identify a neighborhood and block label but provide no point coordinate or stable parcel identifier. The bridge normalizes integer block labels and links a transaction only when the label exists in an approved district correspondence. Exposure is never used to repair an unmatched outcome. This protects outcome blindness but makes administrative coverage a first-order design constraint.

4.3 Outcome-only preparation and limits

The outcome audit starts with 40,715 eligible MOJ rows from 2014-2022. Exact-source fingerprints identify 5,799 duplicate excess rows, leaving 34,916 duplicate-aware transaction proxies. Of those, 16,510 link to 347 approved blocks: 4,380 recorded-land and 12,130 completed-home proxies. Monthly support rules are met by 85 land and 316 completed-home blocks; the stricter primary quarterly construction yields 84 and 315 because all of 2018Q4 is excluded.

These counts establish administrative linkability, not an analysis-ready constant-quality price index. The fingerprints are proxy identifiers rather than legal deed numbers. Registration date may differ from bargaining date, and repeated parcels cannot be tracked. The completed-home denominator is plot area. Conditional exposed/comparison outcome support, mixed-neighborhood variation, selection into

sale, and pre-event comparability remain unknown because no exposure was admitted. Housing data were therefore not the first observed constraint; they may still have become a later one.

5. What the audit finds

5.1 The geography bridge omits real candidate support

Inside the approved bridge, the optical rule produces 45,239 positive cells across 10 of 541 polygons. Outside it, the same pre-mask rule produces another 13,459 positive cells across 15 of 232 held-out polygons, 225 of which intersect the grid. The held-out signal is not optically empty and is not random with respect to the event narrative.

The clearest case is the source polygon labeled `Saba Elahmed`, the CSB spelling used for Sabah Al-Ahmad. It is held out because of an unresolved name conflict, not lack of satellite coverage. The rule identifies 3,835 positive cells there, equal to 1.534 square kilometers and 1.476 percent of the rasterized polygon. Alghais independently identifies Sabah Al-Ahmad as a major case location (Alghais, 2019). The candidate pixels remain unvalidated and cannot be assigned to MOJ-compatible historical blocks, but their presence rules out the broad interpretation that the public imagery contains support in only 10 places.

Figure 2 makes the corrected result explicit. The approved optical proxy fails its frozen 20-unit feasibility floor. The full source geography contains more candidate-positive polygons, but much of that signal cannot enter the housing design because the historical crosswalk is unresolved. The bridge, not simply the sensor, is part of the treatment definition.

Figure 2. Optical candidate support inside and outside the reviewed geography bridge. Approved and held-out counts are assignment polygons; held-out positives remain unlinked and are not treatment.

Notes: The pre-mask reconstruction reproduces all positive and negative labels inside the approved mask. Held-out counts are outcome-blind coverage diagnostics. Seven held-out polygons do not intersect the fixed optical grid, and partial coverage of some large polygons limits area interpretation.

5.2 Nominal pixel abundance collapses at the assignment unit

Within the approved bridge, the 45,239 positive cells equal 18.096 square kilometers. Their distribution is highly uneven. The Bar Elahmadi support polygon contains 34,854 positive cells, or 13.942 square kilometers and 77.0 percent of all approved positive area. Yet those positives cover only 0.3995 percent of that very large polygon's supported cells. The Failaka source polygon contains 3.929 square kilometers and 21.7 percent of approved positive area, but its within-polygon positive share is 8.476 percent. The other eight positive polygons together contain 562 cells, or 0.225 square kilometers.

The 77 percent statistic is therefore a description of where the machine-reference area lies, not treatment intensity, residential damage, or a formal concentration-gate failure. Polygon size mechanically affects raw area. Figure 3 shows absolute candidate-positive area and the within-polygon share side by side. Both denominators are needed before interpreting a block dose.

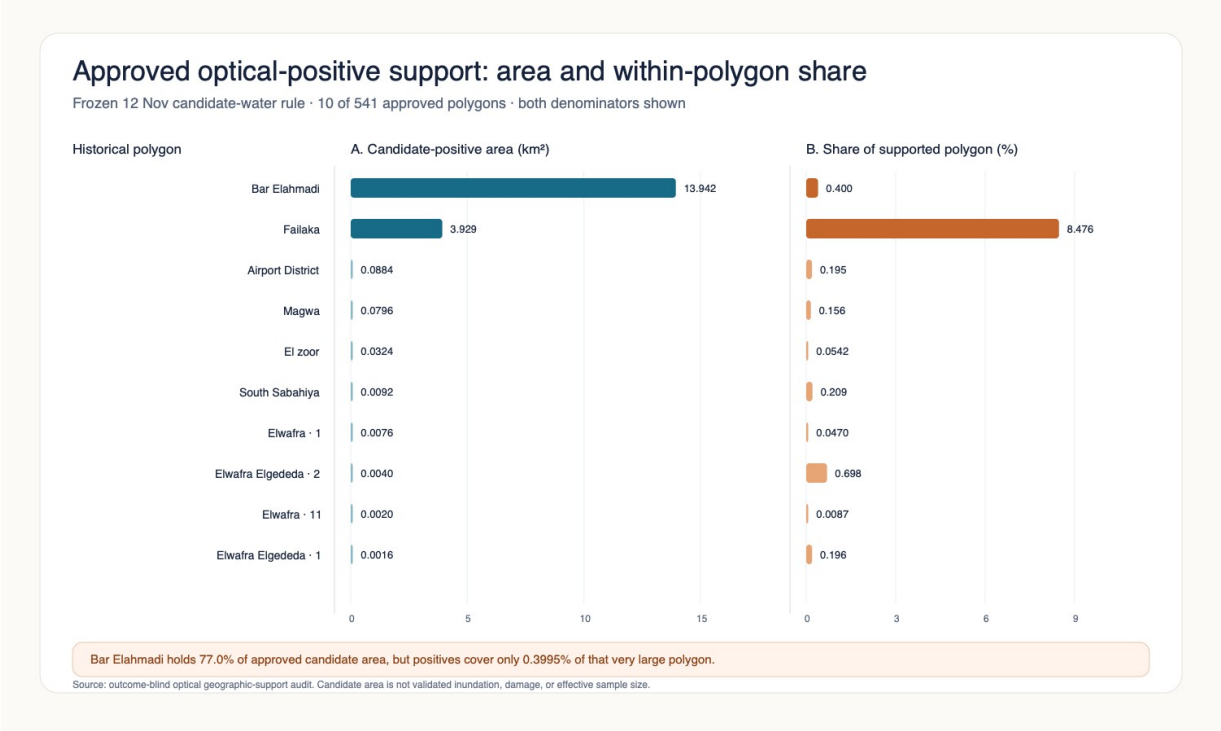

Figure 3. Approved polygons with any optical candidate-positive cell. Absolute candidate area and within-polygon share tell different stories because the historical polygons vary sharply in size.

An outcome-blind land-cover context check adds caution but does not validate salience. ESA WorldCover 2020 classifies 96.1 percent of approved optical positives as bare or sparse vegetation and only 0.47 percent as built-up; 74.5 percent lie more than 500 meters from a 2020 built-up cell (Zanaga et al., 2021). This is a two-years-later land-cover diagnostic. It cannot establish 2018 land use, and generic built-up proximity does not measure road access. The result motivates claims or road-network data, not relabeling.

5.3 Radar has broader raw support but fails promotion

The three early-sequence radar rules produce any-candidate cells in 49, 26, and 11 approved polygons. This corrects an important logical error: 10 optical-positive polygons do not cap a separately generated radar candidate at 10. Radar can predict positives where the optical machine reference is negative or ambiguous. Those extra predictions could be real water missed by the optical rule, temporal change during the 31-hour gap, or false positives.

The relevant next question is measurement performance. A 1,460-cell sample was frozen before opening the radar diagnostic, oversampling optical-positive and optical-negative strata across 175 five-kilometer spatial clusters. Two observations lack a radar prediction, leaving 1,458 evaluations. Inverse-probability weighting restores the prespecified sample-frame mix; cluster bootstrap intervals reflect the spatial resampling design. The full raster pair is also observed, so finite-frame agreement can be

computed exactly conditional on the machine labels, but it supplies neither independent truth nor uncertainty about label construction and temporal mismatch.

Table 2. Early-sequence Sentinel-1 agreement with the optical machine reference

Radar rule	Approved polygons with any candidate	IPW precision [95% spatial-cluster CI]	IPW recall [95% spatial-cluster CI]	Machine-reference rule
t1	49	0.491 [0.252, 0.960]	0.566 [0.447, 0.693]	fail
t2	26	0.714 [0.408, 1.000]	0.484 [0.391, 0.606]	fail
t3	11	0.688 [0.331, 1.000]	0.380 [0.258, 0.541]	fail

Notes: Precision is the share of radar-positive predictions classified positive by the optical reference; recall is the share of optical-reference positives recovered by radar. The author-defined rule required precision of at least 0.80 and recall of at least 0.70, including interval lower bounds. Full-frame descriptive precision/recall are 0.429/0.624 for t1, 0.550/0.543 for t2, and 0.697/0.452 for t3. All views fail. These are radar-optical concordance measures, not ground-truth accuracy.

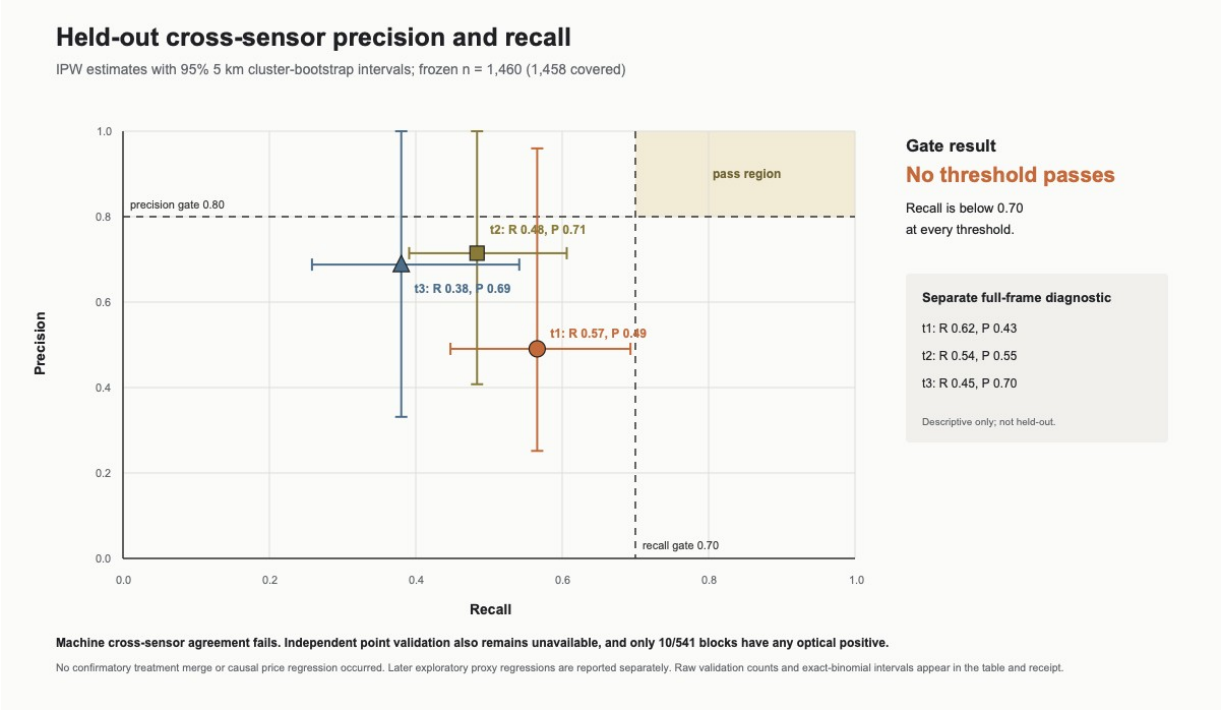

Figure 4. Weighted radar-optical precision and recall. The shaded upper-right region shows the author-defined machine-reference rule; all three candidates remain outside it.

Raw block count alone cannot promote t1 or t2. Both clear 20 any-candidate polygons but fail the measurement rule. Meaningfully exposed, transaction-supported radar blocks were therefore not defined, and the later outcome gates were not reached. Pixel-level agreement would also not be

sufficient by itself: a future validation contract must show how spatially correlated classification error changes block doses and treatment roles.

5.4 The late state lacks an admissible reference design

The intended 25 November target encounters a different problem. Precision requires an admissible sample of predicted-positive cells adjudicated as water or non-water at the radar time. Recall requires a reference frame or probability sample of true-water cells across radar predictions. The public sources audited through 12 August 2026 supply neither denominator at 20-meter precision near 14:50 UTC.

Road closures, shelters, photographs, and reopening records provide sampling anchors and establish disruption. They do not label an arbitrary cell at the pass. A road reopened on 24 November establishes usability, not the absence of adjacent water a day later. Hydraulic depth surfaces are models rather than independent observations. A same-sensor reclassification is also not a reference design. Intervening rainfall prevents nearby optical dates from supplying defensible negatives.

The appropriate status is validation unavailable, not zero accuracy. A separate outcome-blind, post-diagnostic sensitivity intersects the optical reference with the permissive radar t1 rule. Under its implemented 50 percent support and 4,000-square-meter area conventions, it produces four exposed-dose polygons. The sensitivity is useful documentation, but its block thresholds were not frozen before the radar diagnostic, it is mechanically tied to the optical reference, and absolute area is a poor dose for unequal polygons. It is not an additional independent gate or an admissible treatment.

5.5 The stopping decision

The audited implementation stops for a combination of linked and separate reasons. The standalone optical proxy fails the author-defined support floor inside the approved bridge. The held-out audit shows why that result cannot be generalized across Kuwait and reveals that the historical outcome crosswalk is itself incomplete. Radar t1 and t2 have broader raw block support, but all three radar rules fail the machine-reference agreement criterion. The later radar target lacks an admissible exact-time reference design. Economic salience is also unobserved.

These findings close the frozen Sentinel-1/Sentinel-2-CSB/MOJ design as a confirmatory causal study. They do not make every exposure-outcome comparison meaningless. After the admission audit was complete, we ran a separate exploratory specification family using all nine existing radar proxy definitions rather than selecting one from housing outcomes. The family crosses three contrast definitions, three thresholds, two markets, and binary or continuous forms. Twelve of its 36 average-post cells are estimable. For the permissive completed-home binary proxies, the early-sequence association is -0.020 log points with a 95 percent interval [-0.417, 0.378]. The whole-November and

post-13-November definitions produce 0.094 [-0.156, 0.343] and 0.125 [-0.108, 0.357]. These are machine-proxy associations, not flood effects.

The three completed-home definitions do not hold geography fixed. After fixed-effect singleton removal, their proxy-positive block sets overlap by only 3 blocks for early versus whole November, 1 for early versus post-13 November, and 0 for whole versus post-13 November. Their Jaccard overlaps are 0.231, 0.053, and 0.000. The sign pattern can therefore reflect timing, terrain, persistent surfaces, spatial composition, or transaction selection. It shows that treatment construction changes the observed-sale association; it does not isolate the role of the proxy clock. This is not a zero effect result: the intervals admit economically meaningful associations in both directions, and the proxies do not measure verified flooding.

Appendix A reports the design, full 36-cell table, completed-home coefficient plot, support, overlap, pretrend, and influence diagnostics. The transaction-level models use registered transfers from 2014 through 2022, retain block and canonical-neighborhood-by-quarter fixed effects, omit 2018Q4, pool 2019Q1-2022Q4 as post, and use CRV3 inference clustered by canonical neighborhood. Support is computed after singleton removal: the early permissive completed-home model uses 12 proxy-positive blocks across 10 mixed neighborhoods, rather than the pre-singleton count of 13 across 11. None of the estimable pretrend intervals establishes the frozen ±0.05-log-point-per-year equivalence condition. The appendix keeps non-estimable specifications visible and does not promote the coefficient with the most attractive sign or interval.

6. Implications for economic work with satellite treatments

6.1 The downstream assignment unit belongs in validation

The main methodological lesson is not that economists discovered map accuracy. The remote-sensing literature already provides rigorous sampling, response-design, error-matrix, and area-estimation guidance. The economics extension is downstream: classification error must be evaluated after the raster is translated into the units, roles, and counterfactuals used by the estimator.

Millions of cells can collapse into a handful of treatment-assignment units. A classifier can also have reasonable pixel accuracy while misranking block doses, switching a block across a binary threshold, or concentrating all usable variation in one market. The right reporting object is therefore a joint admission table: target state and time, reference design, assignment unit, administrative coverage, block-role uncertainty, salience, outcome support, and stop rule. Pixel performance and economic support answer different questions.

The Kuwait audit also shows why geography is not a neutral merge. The approved bridge excludes places with candidate positives, including Sabah Al-Ahmad, because the historical administrative relationship to MOJ blocks is unresolved. Treating missing geography as untreated would create measurement error correlated with the event narrative. The honest result is a coverage failure. In many low-data settings, repairing the crosswalk may add more economic value than optimizing a new classifier threshold.

6.2 What the audit does not imply

The audit does not measure the full 2018 flood extent, rank blocks for drainage investment, identify damaged homes, evaluate the emergency response, or show that buyers did or did not update beliefs. The exploratory appendix estimates sale-conditional registered-price associations with machine proxies. It does not estimate price, volume, composition, compensation, or welfare effects of verified flooding. The audit also does not show that Sentinel-1 is unsuitable for Kuwait, that public imagery can never support the question, or that 20 exposed blocks would automatically create credible inference.

Nor does the land-cover comparison establish that candidate water was economically irrelevant. Open water in a drainage channel or underpass can disrupt access without falling on a built-up pixel. The current audit lacks the road, structure, incident, claims, and information-salience evidence needed to distinguish such cases. Public repair and compensation expectations would remain part of any future treatment bundle.

The transferable object is the workflow: physical target, temporally compatible reference, economic salience, administrative coverage, assignment-unit support, and only then outcomes. The numeric cutoffs, classifier performance, and geography failures are specific to this event, these sensors, this historical bridge, and these candidate rules.

6.3 A reopening ladder

A new design should begin with the physical target. For a late-state study, same-day or tightly bracketed commercial optical imagery, authenticated timestamped incident records, or another independent event-water source would need positive and dry controls sampled to estimate both false positives and missed water. For an early-sequence study, an urban method that detects narrow or bright double-bounce water may be more relevant than another dark-water threshold. Either approach would be a new registered measurement design, not a repair selected from the failed housing results.

Second, the study must connect physical exposure to damage, access, or information. Geocoded claims, structure-damage inspections, road closures with exact segments and times, or validated route disruption would establish a closer economic construct. Third, the historical geography must be repaired so affected source polygons--including Sabah Al-Ahmad--map unambiguously to stable transaction units.

Fourth, exposed and comparison roles must pass support, influence, and mixed-neighborhood checks without one location dominating. Outcome support, transaction selection, pre-event comparability, interference, and spatial inference can be audited in parallel and can inform exploratory work, but prices cannot choose the physical proxy or determine whether it passes causal admission.

This sequence yields a constructive data implication. Event-time incident systems should preserve timestamps, coordinates, raw provenance, explicit non-inundated controls, damage or access type, and stable geographic identifiers compatible with property records. Those fields would improve later evaluation. This paper does not claim that collecting them would itself reduce flood losses.

7. Conclusion

A flood map becomes an economic treatment through a chain of evidence, not through visual plausibility. The Kuwait audit constructs reproducible optical and radar candidates, then asks whether they measure the right physical state, connect to market salience, cover the administrative outcome geography, and generate supported assignment-unit variation.

The answer for the frozen causal implementation is no. The conservative optical proxy reaches 10 of 541 approved polygons, while a newly completed pre-mask audit finds positives in 15 held-out polygons and exposes a non-random geography-bridge problem. Radar t1 and t2 reach more than 20 raw candidate polygons, but none of the three rules passes the frozen radar-optical criterion. Exact-time late-state truth and economic salience remain unavailable. The complete exploratory family adds a descriptive result: only 12 of 36 cells are estimable, the surviving completed-home associations depend on contrast-specific geography, and the early estimate changes sign when one active neighborhood is removed. Those patterns do not repair the missing causal links.

This does not imply that the floods had no effect on Kuwait’s property market. It identifies the evidence missing before that effect can be estimated credibly. The contribution is the admission logic: validate the target and its timing, verify the economic link, count the units that carry outcomes, and preserve outcome blindness when deciding whether a proxy may enter the confirmatory causal design. Outcome diagnostics and explicitly exploratory regressions remain separate and can proceed without changing that admission decision.

Data availability, rights, and ethics

The aggregate measurement receipts, source audits, non-geographic paper figures, outcome-audit counts, and aggregate exploratory estimates are documented for reproducibility. Sentinel-1 and Sentinel-2 metadata are public through Microsoft Planetary Computer; ESA WorldCover is used only as

later land-cover context. CSB-derived geometry remains subject to reuse review and is not redistributed in the circulable paper bundle. MOJ transaction rows, deed proxies, exact locations, and near-row-level derivatives remain confidential under the ministry portal terms (Kuwait Ministry of Justice, 2026a). Public aggregate estimates and figures use the required attribution: Source: Kuwait Ministry of Justice; author's calculations. No transaction locations, candidate treatment map, or restricted block-level price output is published.

Reproducibility note

The public paper package includes source CSVs and deterministic builders for each displayed figure, aggregate validation tables, a bibliography, a claim-boundary audit, and aggregate exploratory multi-proxy outputs. Package checks reconcile the headline counts to frozen artifacts and verify that exposure did not enter outcome preparation or choose an outcome cleaning rule. The design decision and exploratory coefficient family can be reproduced without redistributing confidential transaction-level records.

Appendix A. Exploratory price patterns across unvalidated flood proxies

The confirmatory admission decision was made without housing outcomes. It found that the existing satellite candidates did not establish physical validity, economic salience, administrative coverage, and assignment support for a causal flood-price estimand. A failed admission decision does not make every descriptive comparison uninformative. It changes the estimand and the language attached to it.

After completing that audit, we merged every already-constructed radar candidate layer with the prepared Ministry of Justice transaction panel. This appendix reports the complete frozen family rather than selecting a candidate from the price results. The outcome-blind claim applies to the admission decision and to the construction of the proxy family. It does not describe this later exploratory outcome analysis.

The realized finding is joint timing-and-geography sensitivity. Twelve of 36 average-post cells are estimable. Among the permissive completed-home binary definitions, the early-sequence association is -0.020 log points, while whole-November and post-13-November definitions produce 0.094 and 0.125. The estimates have wide intervals, and the three proxy-positive block sets scarcely overlap. The pattern therefore does not isolate a clock effect or a flood effect. It shows that alternative machine definitions select different locations and yield different sale-conditional price paths.

A.1 Estimand, sample, and estimator

For proxy definition (p), market (m), registered transfer (i), block (b), canonical neighborhood (n), and quarter (q), the binary average-post model is

$$y_{ibnq} = \alpha_b + \lambda_{nq} + \beta_{pm}\left(P_b^p \times \mathbb{1}\left[q \geq 2019Q1\right]\right) + \varepsilon_{ibnq}.$$

The outcome is log nominal registered transfer price per recorded plot square meter. The binary block indicator equals one when a block contains at least one candidate-positive radar cell under contrast-threshold proxy definition p. The continuous model replaces that indicator with candidate-positive share divided by 0.001, so its coefficient is per 0.1 percentage-point increase in that share.

The sample contains observed registered transfers from 1 January 2014 through 31 December 2022. All of 2018Q4 is omitted as a transition period. The average-post coefficient pools 2019Q1 through 2022Q4; 2018Q3 is the omitted reference quarter in the internal event-time diagnostic. Models are estimated separately for completed homes and recorded land, include block and canonical-neighborhood-by-quarter fixed effects, and use CRV3 inference clustered by canonical neighborhood. Fixed-effect singletons are removed iteratively before support is reported.

The remaining variation is a transaction-weighted, within-neighborhood difference in pre/post registered-price paths between proxy-positive and other blocks. The estimand is conditional on an observed sale. It is not an ATT, a property-stock capitalization effect, or a constant-quality price-index effect. The proxy can also capture persistent water, pavement, terrain, drainage infrastructure, later rainfall, or classification error. Sale probability, transaction volume, and sold-property composition remain additional explanations for the coefficient.

A.2 The complete average-post family

Table A1 reports all three contrasts, all three thresholds, both markets, and binary and continuous forms. Non-estimable rows remain visible. Twelve of 36 cells contain a coefficient; 24 lose the relevant within-neighborhood variation. NE is an algebraic support result, not a zero.

Table A1. Complete exploratory average-post specification family

Contrast-specific proxy	Rule	Market	Form	Association [95% CRV3 CI]	Candidate blocks / mixed neighborhoods
Early sequence (1-13 Nov)	t1	Completed home	Binary	-0.020 [-0.417, 0.378]	12 / 10

Contrast-specific proxy	Rule	Market	Form	Association [95% CRV3 CI]	Candidate blocks / mixed neighborhoods
Early sequence (1-13 Nov)	t1	Completed home	Continuous	-0.010 [-0.064, 0.043]	12 / 10
Early sequence (1-13 Nov)	t1	Recorded land	Binary	0.376 [-0.453, 1.205]	9 / 7
Early sequence (1-13 Nov)	t1	Recorded land	Continuous	0.081 [-0.246, 0.407]	9 / 7
Early sequence (1-13 Nov)	t2	Completed home	Binary	-0.078 [-0.538, 0.382]	2 / 2
Early sequence (1-13 Nov)	t2	Completed home	Continuous	-0.028 [-0.167, 0.110]	2 / 2
Early sequence (1-13 Nov)	t2	Recorded land	Binary	NE	1 / 1
Early sequence (1-13 Nov)	t2	Recorded land	Continuous	NE	1 / 1
Early sequence (1-13 Nov)	t3	Completed home	Binary	NE	1 / 1
Early sequence (1-13 Nov)	t3	Completed home	Continuous	NE	1 / 1
Early sequence (1-13 Nov)	t3	Recorded land	Binary	NE	0 / 0
Early sequence (1-13 Nov)	t3	Recorded land	Continuous	NE	0 / 0
Whole November (1-25 Nov)	t1	Completed home	Binary	0.094 [-0.156, 0.343]	4 / 4
Whole November (1-25 Nov)	t1	Completed home	Continuous	0.044 [-0.053, 0.141]	4 / 4
Whole November (1-25 Nov)	t1	Recorded land	Binary	0.097 [-8.463, 8.656]	2 / 2
Whole November (1-25 Nov)	t1	Recorded land	Continuous	-0.012 [-3.840, 3.816]	2 / 2
Whole November (1-25 Nov)	t2	Completed home	Binary	NE	1 / 1
Whole November (1-25 Nov)	t2	Completed home	Continuous	NE	1 / 1
Whole November (1-25 Nov)	t2	Recorded land	Binary	NE	1 / 1

Contrast-specific proxy	Rule	Market	Form	Association [95% CRV3 CI]	Candidate blocks / mixed neighborhoods
Whole November (1-25 Nov)	t2	Recorded land	Continuous	NE	1 / 1
Whole November (1-25 Nov)	t3	Completed home	Binary	NE	0 / 0
Whole November (1-25 Nov)	t3	Completed home	Continuous	NE	0 / 0
Whole November (1-25 Nov)	t3	Recorded land	Binary	NE	0 / 0
Whole November (1-25 Nov)	t3	Recorded land	Continuous	NE	0 / 0
Post-13 November (13-25 Nov)	t1	Completed home	Binary	0.125 [-0.108, 0.357]	8 / 8
Post-13 November (13-25 Nov)	t1	Completed home	Continuous	0.044 [-0.046, 0.134]	8 / 8
Post-13 November (13-25 Nov)	t1	Recorded land	Binary	NE	4 / 3
Post-13 November (13-25 Nov)	t1	Recorded land	Continuous	NE	4 / 3
Post-13 November (13-25 Nov)	t2	Completed home	Binary	NE	0 / 0
Post-13 November (13-25 Nov)	t2	Completed home	Continuous	NE	0 / 0
Post-13 November (13-25 Nov)	t2	Recorded land	Binary	NE	0 / 0
Post-13 November (13-25 Nov)	t2	Recorded land	Continuous	NE	0 / 0
Post-13 November (13-25 Nov)	t3	Completed home	Binary	NE	0 / 0
Post-13 November (13-25 Nov)	t3	Completed home	Continuous	NE	0 / 0
Post-13 November (13-25 Nov)	t3	Recorded land	Binary	NE	0 / 0
Post-13 November (13-25 Nov)	t3	Recorded land	Continuous	NE	0 / 0

Notes: The outcome is log nominal registered transfer price per recorded plot square meter, conditional on an observed sale. Binary rows compare blocks with any candidate-positive cell with other blocks;

continuous rows are per 0.1 percentage-point increase in candidate-positive share. All models use the 2014-2022 transaction sample, omit 2018Q4, pool 2019Q1-2022Q4 as post, include block and canonical-neighborhood-by-quarter fixed effects, and report 95 percent CRV3 intervals clustered by canonical neighborhood. Support is computed after iterative fixed-effect singleton removal. NE means the proxy interaction is absorbed or there is insufficient within-neighborhood variation; it is not a zero association. These are unvalidated machine-proxy associations, not flood effects. Twelve of 36 cells are estimable. Source: Kuwait Ministry of Justice; author's calculations.

Figure A1 shows all 18 completed-home cells on the same log-point scale. Eight are estimable. The early t1 model has the broadest surviving support, but its conventional interval remains wide. The strict completed-home definitions depend on only one or two proxy-positive blocks before the interaction disappears.

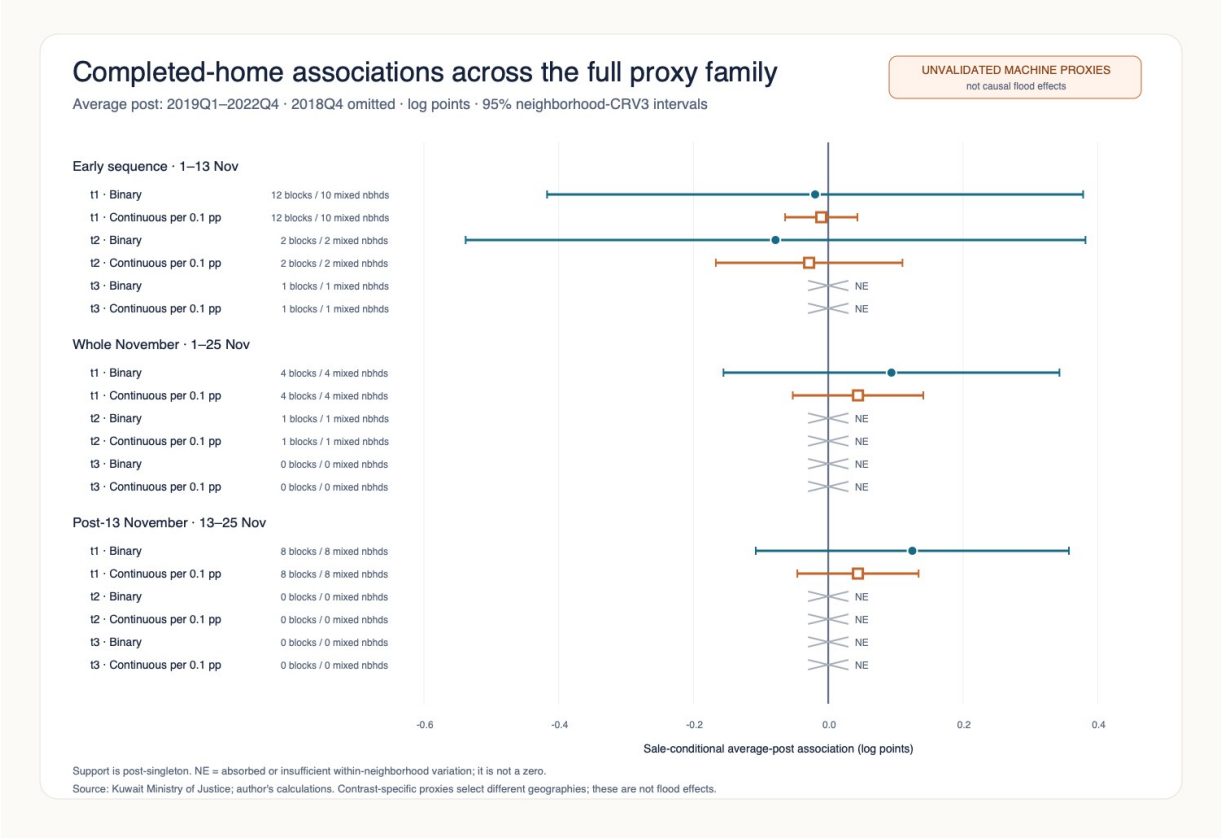

Figure A1. Completed-home average-post associations across every contrast, threshold, and form. Crosses marked NE preserve specifications that are absorbed or lack within-neighborhood variation.

Notes: All intervals use CRV3 clustered by canonical neighborhood. Support is post-singleton. Alternative contrast definitions select different geographies, so cross-panel differences combine timing and spatial composition. Source: Kuwait Ministry of Justice; author's calculations.

A.3 Support and singleton removal

Raw merged-sample counts overstate the support attached to the fitted coefficients. In the completed-home models, 185 observations are fixed-effect singletons; the recorded-land models remove 359. The early t1 completed-home merge contains 13 proxy-positive blocks across 11 neighborhoods before removal, but the fitted sample retains 12 blocks across 10 neighborhoods. All 12 have pre and post transactions, and the 10 neighborhoods contain both proxy-positive and other blocks. The fitted model uses 11,541 transactions across 338 blocks and 56 neighborhoods.

The other headline completed-home definitions are thinner. Whole November uses four proxy-positive blocks in four mixed neighborhoods, while the post-13-November definition uses eight in eight. The early t2 estimate depends on only two blocks in two neighborhoods. Land is more fragile: the early t1 coefficient uses seven mixed neighborhoods, and the whole-November coefficient uses two.

These counts explain why the nominal sample-wide cluster totals are not the relevant information totals. Most neighborhoods contribute outcome observations, but only mixed neighborhoods carry the proxy-by-post contrast.

A.4 The proxy clocks select different geographies

The three permissive completed-home definitions do not compare the same proxy-positive blocks under alternative clocks. Table A2 reports overlap after singleton removal.

Table A2. Pairwise overlap of completed-home t1 proxy-positive blocks

Pair	Block intersection	Jaccard overlap
Early sequence vs whole November	3	0.231
Early sequence vs post-13 November	1	0.053
Whole November vs post-13 November	0	0.000

Notes: Early, whole-November, and post-13-November sets contain 12, 4, and 8 post-singleton proxy-positive completed-home blocks. Radar candidate layers define the proxy-positive sets. Source: Kuwait Ministry of Justice; author's calculations.

The negative-to-positive sign pattern is therefore joint timing-and-geography sensitivity. A later contrast can select different terrain, persistent surfaces, drainage systems, neighborhoods, and transacting properties. These regressions show that the economic answer depends on how candidate pixels are

translated into treatment-like units. They do not show that one fixed geography changes sign when only the clock changes.

A.5 Pre-period paths and influence

The completed-home t1 binary pretrend estimates are 0.195 [-0.144, 0.535] log points per year for the early sequence, 0.117 [-0.147, 0.381] for whole November, and -0.013 [-0.184, 0.157] for the post-13-November definition. The corresponding continuous estimates are 0.030 [-0.021, 0.082], 0.044 [-0.023, 0.111], and -0.023 [-0.118, 0.072]. None of these intervals lies inside the frozen equivalence band of ±0.05 log points per year. More generally, none of the 14 estimable pretrend intervals in the complete family establishes equivalence. The analysis therefore cannot rule out pre-existing differential paths.

Table A3 reports influence diagnostics for the completed-home t1 binary estimates. The neighborhood shares allocate squared residualized treatment variation after the same fixed effects. Leave-one-out ranges delete each active proxy-positive neighborhood in turn. These are descriptive fragility checks, not alternative causal estimates.

Table A3. Completed-home t1 identifying-support diagnostics

Contrast-specific proxy	Association	Active mixed neighborhoods	Largest neighborhood share	Inverse-Herfindahl diagnostic	Leave-one-active-neighborhood range
Early sequence	-0.020	10	0.206	7.47	[-0.137, 0.098]
Whole November	0.094	4	0.407	2.86	[0.033, 0.166]
Post-13 November	0.125	8	0.260	6.34	[0.042, 0.169]

Notes: The early estimate changes sign under active-neighborhood deletion. The inverse-Herfindahl value summarizes the concentration of residualized treatment variation; it is not an inferential degrees-of-freedom correction. Source: Kuwait Ministry of Justice; author's calculations.

Land support is too leveraged for substantive interpretation. One neighborhood carries 63.3 percent of residualized treatment variation in the early t1 land model and 93.4 percent in the whole-November t1 model. Their inverse-Herfindahl diagnostics are 2.32 and 1.14, and their CRV3 intervals are extremely wide. The complete table retains these coefficients so that the family is auditable, but the paper draws no market conclusion from them.

A.6 Why event studies are not displayed

The internal pipeline retains binary event-time diagnostics with 2018Q3 as the reference quarter and 2018Q4 excluded. Those paths are not publication exhibits. They require CRV1 because CRV3 becomes singular in sparse cells, retained relative-time terms differ across definitions, and seven of 18 proxy-market cells return no event coefficients. Displaying a selected path would therefore look more complete and inferentially stable than the event-study family actually is. The complete average-post table, support audit, overlap matrix, pretrend intervals, and influence diagnostics are the more honest public summary.

A.7 Claim boundary and reproduction

The appendix supports three statements. First, the prepared MOJ transactions can be joined to the approved satellite geography for exploratory within-neighborhood comparisons. Second, only a minority of the complete proxy family retains identifying variation, and the surviving sale-conditional associations depend on contrast-specific geography. Third, the current data do not isolate a stable registered-price response to verified flooding.

It does not support the claim that the 2018 floods raised, reduced, or had no effect on Kuwaiti property prices. The complete proxy-family audit instead shows that most strict cells lose identifying variation, while the surviving associations remain vulnerable to geography choice, pre-existing paths, influence, and transaction selection.

Run `econometrics/run_exploratory_proxy_analysis.py` in the maintained Python 3.12/pyfixest environment. Machine-readable post associations, pretrends, support, overlap, influence diagnostics, internal event paths, and a hash receipt are in `econometrics/exploratory_results/`. The receipt records the measurement, transaction, crosswalk, analysis-config, script, and output hashes plus runtime versions. Public outputs contain aggregate coefficients and counts only. MOJ transaction rows, deed proxies, exact locations, and near-row-level derivatives remain confidential. Public aggregate estimates and figures use the attribution: **Source: Kuwait Ministry of Justice; author's calculations.** CSB geometry and raster maps remain excluded pending their separate reuse review.

References

Al-Hemoud, Ali, Al-Enezi, Abdulla, Al-Dashti, Hassan, Petrov, Peter, Misak, Raafat, AlSaraf, Manar, and Malek, Mariam (2023). Hazard Assessment and Hazard Mapping for Kuwait. International Journal of Disaster Risk Science. https://doi.org/10.1007/s13753-023-00473-2.

Al-Yamani, Faiza Yousef (2021). Fathoming the Northwestern Arabian Gulf: Oceanography and Marine Biology. Kuwait Institute for Scientific Research. https://www.kisr.edu.kw/media/filer_public/d7/f7/d7f77a5e-9ab4-4339-b3f3-057c990d45e1/fathoming_book_20_sep_optimized.pdf.

AlBader, Suleiman H. and AlSabah, Khaled W. (2025). Waiting for a Home: The Price of Public Housing Delays in Kuwait. LSE Middle East Centre. https://eprints.lse.ac.uk/129678/.

Alfalah, Abdullah Adel, D'Arcy, Eamonn, and Stevenson, Simon (2023). Constructing House Price Indices in an Emerging Market: The Case of Kuwait. Journal of Real Estate Literature. https://doi.org/10.1080/09277544.2023.2273050.

Alghais, Nayef (2019). Evaluation of the Impacts of Lack of Geoinformation Data in Crisis Management during the 2018 Kuwait Flood. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences. https://doi.org/10.5194/isprs-archives-XLII-3-W8-17-2019.

Alshalfan, Sharifa, Alrasheed, Dhari S., and Albabtain, Barrak (2022). Housing Kuwaitis: An Overview of the Current Model and Its Implications on Affordability and Quality of Life. Kuwait Foundation for the Advancement of Sciences. https://www.kfas.org/Publications/Housing-Kuwaitis.

Alsharif, Haya, Rambachan, Ashesh, Singh, Rahul, and Viviano, Davide (2026). Causal Inference with Satellite Imagery: A Comparison of Methods for Forest Conservation Data. AEA Papers and Proceedings. https://doi.org/10.1257/pandp.20261019.

Bakkensen, Laura A. and Barrage, Lint (2022). Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics. Review of Financial Studies. https://doi.org/10.1093/rfs/hhab122.

Beltrán, Allan, Maddison, David, and Elliott, Robert J. R. (2018). Is Flood Risk Capitalised into Property Values?. Ecological Economics. https://doi.org/10.1016/j.ecolecon.2017.12.015.

Donaldson, Dave and Storeygard, Adam (2016). The View from Above: Applications of Satellite Data in Economics. Journal of Economic Perspectives. https://doi.org/10.1257/jep.30.4.171.

European Space Agency (2022). Sen2Cor 2.12.03 Configuration and User Manual. https://step.esa.int/thirdparties/sen2cor/2.12.0/docs/OMPC.TPZ.SUM.002%20-%20i1r1%20-%20Sen2Cor%202.12.03%20Configuration%20and%20User%20Manual.pdf.

Gallagher, Justin (2014). Learning about an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States. American Economic Journal: Applied Economics. https://doi.org/10.1257/app.6.3.206.

Hino, Miyuki and Burke, Marshall (2021). The Effect of Information about Climate Risk on Property Values. Proceedings of the National Academy of Sciences. https://doi.org/10.1073/pnas.2003374118.

Jain, Meha (2020). The Benefits and Pitfalls of Using Satellite Data for Causal Inference. Review of Environmental Economics and Policy. https://doi.org/10.1093/reep/rez023.

Kuwait Central Statistical Bureau (2011). Kuwait Census 2011 Geoportal: Lower-Area Geography. https://gis.csb.gov.kw/en/.

Kuwait Ministry of Justice (2026a). Electronic Portal Terms of Use. https://www.moj.gov.kw/AR/pages/Footer03.aspx.

Kuwait Ministry of Justice (2026b). Live Registered Real-Estate Sale-Price Query. https://www.realestate.moj.gov.kw/live/Moj_Rs_11.aspx.

Kuwait News Agency (2018a). Civil Defense Provides Shelters for People Hit by Downpour. https://www.kuna.net.kw/ArticlePrintPage.aspx?id=2758482&language=en.

Kuwait News Agency (2018b). Interior Ministry Closes Roads Because of Accumulated Water. https://www.kuna.net.kw/ArticleDetails.aspx?Language=ar&id=2758443.

Kuwait News Agency (2018c). Kuwait Government Bodies Mobilized to Deal with Heavy Rainfall. https://www.kuna.net.kw/ArticleDetails.aspx?id=2757334&language=en.

Kuwait News Agency (2018d). Torrential Rains Swamp Kuwait Roads. https://www.kuna.net.kw/ArticleDetails.aspx?id=2758589&language=en.

Lin, Yunung Nina, Yun, Sang-Ho, Bhardwaj, Alok, and Hill, Emma M. (2019). Urban Flood Detection with Sentinel-1 Multi-Temporal Synthetic Aperture Radar (SAR) Observations in a Bayesian Framework: A Case Study for Hurricane Matthew. Remote Sensing. https://doi.org/10.3390/rs11151778.

McFeeters, Stuart K. (1996). The Use of the Normalized Difference Water Index (NDWI) in the Delineation of Open Water Features. International Journal of Remote Sensing. https://doi.org/10.1080/01431169608948714.

Microsoft Planetary Computer (2026). Sentinel-1 Radiometrically Terrain Corrected Collection. STAC collection and technical documentation. https://planetarycomputer.microsoft.com/api/stac/v1/collections/sentinel-1-rtc.

Olofsson, Pontus, Foody, Giles M., Herold, Martin, Stehman, Stephen V., Woodcock, Curtis E., and Wulder, Michael A. (2014). Good Practices for Estimating Area and Assessing Accuracy of Land Change. Remote Sensing of Environment. https://doi.org/10.1016/j.rse.2014.02.015.

Proctor, Jonathan, Carleton, Tamma, and Sum, Sandy (2026). Parameter Recovery Using Remotely Sensed Variables. Journal of the Association of Environmental and Resource Economists. https://doi.org/10.1086/742595.

Stehman, Stephen V. and Czaplewski, Raymond L. (1998). Design and Analysis for Thematic Map Accuracy Assessment. Remote Sensing of Environment. https://doi.org/10.1016/S0034-4257(98)00010-8.

Xu, Hanqiu (2006). Modification of Normalised Difference Water Index (NDWI) to Enhance Open Water Features in Remotely Sensed Imagery. International Journal of Remote Sensing. https://doi.org/10.1080/01431160600589179.

Zanaga, Daniele, Van De Kerchove, Ruben, Daems, Dirk, De Keersmaecker, Wanda, Brockmann, Carsten, Kirches, Grit, Wevers, Jan, Cartus, Oliver, Santoro, Maurizio, Fritz, Steffen, Lesiv, Myroslava, Herold, Martin, Tsendbazar, Nandin-Erdene, Xu, Panpan, Ramoino, Fabrizio, and Arino, Olivier (2021). ESA WorldCover 10 m 2020 v100. Zenodo. https://doi.org/10.5281/zenodo.5571936.