



UNIVERSITÉ DU QUÉBEC EN OUTAOUAIS
DÉPARTEMENT DES SCIENCES ADMINISTRATIVES

WORKING PAPER NO. 22

SHARP-HPI: A Spatially Hierarchical, Adaptive and Robust Framework for Quality-Adjusted Housing Price Indices

Monthly municipal indices for Quebec from 770,000 registered sales, 2021–2026

Simon-Pierre Boucher

Département des sciences administratives
Université du Québec en Outaouais

simon-pierre.boucher@uqo.ca

Gatineau – Pavillon Alexandre-Taché
283, boulevard Alexandre-Taché
Gatineau, Québec, Canada J9A 1L8

October 2026

Version 1.0

Abstract

This paper develops and evaluates SHARP-HPI, a framework for monthly quality-adjusted housing price indices at the municipal level in thin markets. Stage one is a robust pooled hedonic regression with fine-scale location fixed effects, smoothly time-varying quality prices and provincial month effects; stage two models each geography's deviation from its parent (province, administrative region, regional county municipality, municipality) with heteroskedastic local-level state-space models whose innovation variances are calibrated by empirical Bayes across cells, so that the weight given to local sales adapts to their reliability. We apply it to 769,995 registered sales in Quebec from January 2021 to October 2026 matched to the assessment roll, with an audited population definition (632,006 sales), a measured publication lag and real-time vintages. Against eleven benchmark estimators (median, stratified median, pooled and rolling time dummies, hedonic imputation, GEKS, HMTS, repeat sales, gradient boosting and a production hierarchical index), Monte Carlo experiments on synthetic Quebec-like markets with known latent paths show that the adaptive hierarchy halves the error of municipal monthly changes relative to per-cell state-space estimation and keeps interval coverage near nominal, and two deliberate counter-examples show where it does not: heavy-tailed contamination favours a coarser location partition, and at the regional level region-specific hedonic coefficients beat pooled coefficients plus a smoothed deviation. The empirical panel covers 2,654 geography–class cells with reliability grades, effective sample sizes and local-data weights; aggregates agree closely with Teranet and CREA indices in Québec City and diverge by an explainable margin in Montréal. Every ingredient of the estimator exists in the literature; the contribution is their integration, the explicit reliability accounting and the documented performance.

Keywords: house price index, hedonic regression, state-space model, small-area estimation, empirical Bayes, real-time data, Quebec

JEL Classification: R31, C43, C11, C21, C55

1. Introduction

A house price index answers a deceptively simple question: by how much has the price of a dwelling of constant quality changed between two dates? For a province, the question is well posed and the data abundant. For a municipality of three thousand inhabitants that records four sales a month, the question is the same but the evidence is thin, the sold dwellings differ from month to month, and any estimator must decide how much to trust the local sales and how much to borrow from neighbouring and encompassing markets. Statistical agencies have answered this trade-off with rules of thumb (minimum counts, suppression) or with methods designed for liquid markets. This paper treats the trade-off as an estimation problem.

We propose SHARP-HPI, a two-stage framework for monthly, quality-adjusted price indices across a geographic hierarchy. The first stage is a robust pooled hedonic regression in which log prices are explained by structural characteristics, fine-scale location fixed effects (a nested metric grid whose resolution adapts to the density of sales), provincial month effects, and quality prices allowed to drift smoothly over time. The second stage treats each geography’s monthly mean residual as a noisy observation of a latent deviation from its parent’s path and estimates that deviation with a heteroskedastic local-level state-space model. The innovation variance of each cell, which governs how quickly the deviation is allowed to move and therefore how much weight a month of local sales receives, is calibrated by empirical Bayes across all cells of a level rather than estimated from each cell’s own handful of observations. The result is a panel of indices for the province, 103 regional county municipalities (MRCs) and 1,125 municipalities by property class, each with a count of actual sales, an effective sample size, a local-data weight, an interval, and a reliability grade.

The paper makes four contributions. First, it builds the index on an audited population: 769,995 registered sales published by the Quebec land registry feed between January 2021 and October 2026, matched to the provincial assessment roll, with an exclusion ledger (Table 2), a measured publication lag and a classification of dwellings into mutually exclusive classes. Second, it compares the proposed estimator with eleven benchmarks spanning the hedonic, multilateral, repeat-sales, machine-learning and production-index traditions on the same data. Third, it evaluates all estimators in Monte Carlo experiments on synthetic markets with known quality-constant price paths, including markets where the quality of sold homes rises while prices are flat, markets with heavy-tailed noise, and two counter-examples designed so that the adaptive hierarchy should not win. Fourth, it simulates the real-time publication process with monthly vintages and quantifies revisions, timeliness and publication policies.

We are explicit about novelty. The verified literature (Section 2) contains every ingredient: structural time-series index levels (Francke, 2010), hierarchical common-plus-deviation trends

with heavy-tailed errors (Francke and van de Minne, 2017), small-area shrinkage with estimated weights (Fay and Herriot, 1979; Rao and Molina, 2015), rolling-window hedonics with low-volume safeguards (Hill et al., 2022), and hedonic multilateral time-series smoothing for small domains in official statistics (Ishaak et al., 2024). What we did not find in a single published estimator is the combination of a single-sale hedonic observation equation, a multi-level administrative hierarchy, empirical-Bayes calibration of the pooling strength across cells, reliability accounting for every cell-month, and real-time vintage evaluation with a measured registration lag. The paper should be read as an integration and empirical assessment of these components, not as a new class of model.

Three findings matter for practice. In the Monte Carlo experiments the hierarchical state-space estimator reduces the root mean squared error of municipal monthly changes to 1.58 log points, against 2.91 for per-cell state-space estimation, 12.52 for a cell-level time dummy and 24.28 for the median; the empirical-Bayes calibration of the innovation variance matters for interval coverage rather than for point accuracy; the local-linear-trend and Gaussian variants do not improve on the robust local level; and two counter-examples show where the design loses: with heavy-tailed contamination a coarser location partition is more robust, and at the regional level, where sales are plentiful, estimating each region’s own hedonic coefficients beats pooled coefficients plus a smoothed deviation. On the Quebec data the provincial house index rises from 100 in 2021 to 154.9 in September 2026, condominiums to 136.4 and plexes to 151.2; the dispersion across municipalities is wide and largely reliable only for cells graded A to C. In real time, first-release monthly changes at the provincial level are revised by a mean absolute 0.16 log points (houses), and municipal first releases by 1.72; publishing without revision costs about half a log point per month at the municipal level.

The remainder of the paper follows the protocol in which it was pre-specified. Section 2 reviews the literature and states the gap. Section 3 describes the data, the institutional setting and the population definition. Section 4 defines the estimand and the measurement challenges. Section 5 presents the SHARP-HPI model and identification. Section 6 covers estimation, inference, linking and vintages. Section 7 reports the Monte Carlo evidence. Section 8 presents the Quebec indices, Section 9 the benchmarks, validation, ablations and robustness, Section 10 the implications, and Section 11 the limitations and conclusion. Code, the index panel, the vintage database and every table and figure are generated by the accompanying repository.

2. Literature and research gap

2.1. Quality adjustment: hedonic time dummies, imputation and multilateral methods

The hedonic approach to constant-quality price measurement goes back to Court (1939) and Griliches (1961), with the theoretical foundation in Rosen (1974). For housing, Hill (2013)

surveys and classifies the hedonic index methods, and the Eurostat et al. handbook (Eurostat et al., 2013) codifies them for official statistics. Two families dominate. The *time-dummy* method pools transactions across periods and reads the index off period dummies; the *hedonic imputation* method estimates a separate hedonic function per period and imputes the price of a fixed basket of dwellings at each period’s prices, yielding Laspeyres, Paasche and Fisher variants (Hill and Melser, 2008; Diewert et al., 2009; de Haan, 2010; Silver and Heravi, 2007). Hill and Scholz (2018) show that flexible geospatial controls materially improve hedonic imputation indices. The time-dummy method revises its history whenever new periods are added; the *rolling time dummy* (RTD) of Hill et al. (2022) estimates the model on a rolling window and chains only the newest movement, with options for choosing the window length, the linking rule and for handling periods with few transactions. Multilateral methods designed for scanner data have been transferred to housing: GEKS on bilateral hedonic links with rolling windows and splicing (Ivancic et al., 2011; de Haan and Krsinich, 2014; Diewert and Fox, 2022) and the hedonic multilateral time-series method (HMTS) of Ishaak et al. (2024), which smooths the series of imputed average prices of each base period with a state-space model and splices a window of them, producing stable and limited-revisable indices for small domains of commercial property in the Netherlands. Hill and Steurer (2026) review the developments in real estate price measurement since the global financial crisis.

2.2. Repeat sales, hybrids and thin markets

The repeat-sales method (Bailey et al., 1963; Case and Shiller, 1987, 1989; Shiller, 1991) avoids the need for characteristics by comparing successive sales of the same dwelling, at the cost of discarding single sales and of selection and renovation biases (Clapp et al., 1991; Gatzlaff and Haurin, 1997, 1998; Hwang and Quigley, 2004). Hybrid models combine single and repeated sales (Case and Quigley, 1991; Quigley, 1995), and Wallace and Meese (1997) and Dorsey et al. (2010) compare the families empirically. For thin markets, Schwann (1998) and Francke (2010) replace period dummies by structural time-series models in which the index level follows a random walk or a local linear trend, so that periods with no sales still receive an estimate. Francke and van de Minne (2017) extend this to a hierarchical model with one common trend and cluster-specific deviations and Student-*t* errors, obtaining reliable indices at a granular level from few observations; Francke et al. (2025) add spatially varying loadings on latent factors to produce census-tract indices; Bogin et al. (2019) and Contat and Larson (2024) build large panels of local repeat-sales indices; Ren et al. (2017) cluster sparse tract series in a Bayesian nonparametric model; Lopez and Hewings (2018) propose a matching Fisher index for small areas; Bourassa et al. (2013, 2016) study robust repeat-sales and hedonic indices. Nowak and Smith (2020) document that the constant-quality assumption of repeat-sales indices fails in a time-varying way.

2.3. Small-area estimation, state-space models and reconciliation

The statistical treatment of “borrowing strength” is the small-area estimation literature: the area-level model of [Fay and Herriot \(1979\)](#) shrinks a direct estimate toward a synthetic one with a weight equal to the ratio of model variance to total variance, and its mean squared error is well understood ([Datta and Lahiri, 2000](#); [Datta et al., 2005](#); [Pfeffermann, 2013](#); [Rao and Molina, 2015](#)); the empirical-Bayes roots are [Efron and Morris \(1973, 1975\)](#), and [Gelman \(2006\)](#) discusses priors for variance parameters in hierarchical models. State-space methods ([Kalman, 1960](#); [Harvey, 1989](#); [Durbin and Koopman, 2012](#)) provide filtered (one-sided) and smoothed (two-sided) estimates of latent levels with exact variances. Coherent combination of estimates across a hierarchy is the subject of forecast reconciliation ([Hyndman et al., 2011](#); [Wickramasuriya et al., 2019](#); [Athanasopoulos et al., 2024](#)).

2.4. Timeliness, revisions and real-time evaluation

Indices built from deeds are published with a lag and revised as late registrations arrive. [Clapp and Giaccotto \(1999\)](#) and [Clapham et al. \(2006\)](#) analyse revisions in repeat-sales indices; [van de Minne et al. \(2020\)](#) use revisions as a quality criterion; [Sayag et al. \(2022\)](#) nowcast late-reported transactions to cut revisions; [Anenberg and Laufer \(2017\)](#) date transactions at the contract date using listings; [Croushore \(2011\)](#) surveys real-time data analysis. In Canada, the Teranet–National Bank index is a repeat-sales index on registry transfers ([Teranet and National Bank of Canada, 2017](#)), the CREA MLS Home Price Index is a benchmark hedonic index on brokered resales ([Canadian Real Estate Association, 2023](#)), and Statistics Canada’s resale index component used an arithmetic repeat-sales method ([Statistics Canada, 2019](#)).

2.5. The gap

Table 1 summarises what the closest predecessors do and what remains. RTD and HMTS treat each geography separately and obtain stability from windows and splices rather than from an estimated signal-to-noise ratio; the Francke hierarchy and the small-area models estimate the pooling weight but with a repeat-sales observation equation, time-constant hyper-parameters and a single deviation level; reconciliation operates on forecasts of additive series. The combination examined here is (i) a single-sale hedonic observation equation with fine location control, (ii) a four-level administrative hierarchy, (iii) innovation variances calibrated across cells by empirical Bayes so that the pooling weight is both cell-specific and time-varying through the Kalman gain, (iv) explicit reliability accounting, and (v) real-time evaluation with a measured registration lag. We establish by simulation and on the data whether this combination improves on its components; we do not claim that any component is new.

Table 1. Closest predecessors and the remaining space

Approach	What it does	What SHARP-HPI adds or changes
Rolling time dummy (Hill et al., 2022)	Pooled hedonic on a rolling window; only the newest movement is chained; low-volume safeguards	Cell-specific, estimated pooling across a hierarchy; model-based intervals; revisions measured rather than avoided by construction
HMTS (Ishaak et al., 2024)	Imputed average prices by base period smoothed series-by-series with a state-space model, windows spliced	One joint hierarchy with cross-cell calibration of the signal-to-noise ratio; robust innovations; single-sale observation equation
Hierarchical repeat sales (Francke and van de Minne, 2017)	Common trend plus cluster deviations, Student- t errors, repeat-sales pairs	Hedonic single sales (no pair selection), four nested levels, empirical-Bayes calibration of τ^2 across cells, reliability grades
Fay–Herriot / SAE (Fay and Herriot, 1979 ; Rao and Molina, 2015)	Static area-level shrinkage with estimated variance components and MSE	Dynamic (state-space) version across a geographic hierarchy; direct, model-assisted and unidentified cells reported
Forecast reconciliation (Wickramasuriya et al., 2019)	Coherent forecasts of additive hierarchies	Tested as an aggregation diagnostic for log-linear stock-weighted aggregates; reported when it smooths rather than improves

3. Quebec transaction data and institutional setting

3.1. Institutional setting

In Quebec every transfer of immovable property is registered at the land registry (*Registre foncier du Québec*), and the deed records the price. Deeds are published shortly after registration and redistributed by data providers. Municipal property assessment is governed by the *Loi sur la fiscalité municipale*: each municipality (or its MRC) deposits a triennial assessment roll that records, for every assessment unit, its use code (CUBF), number of dwellings, physical link, floor area of storeys, lot area, year of construction, number of storeys and the assessed values of land and building. The roll is open data and is updated with a lag relative to construction. There is no multiple-listing requirement: a substantial share of sales occur without a broker, so registry data cover more of the market than brokered-resale indices do.

3.2. Source, snapshot and provenance

The data are the transaction table of *Vrai Prix* (www.vrai-prix.com), a platform that ingests daily the “Transactions immobilières” feed of the Journal de Québec / Journal de Montréal (api.qub.ca), which republishes land-registry deeds, geocodes them, and matches them to the provincial assessment roll by address. We took one immutable SQLite snapshot of the production database on 21 October 2026 (SHA-256 recorded in the configuration) and never read the live file again. The snapshot contains

769,995 sales dated 4 January 2021 to 7 October 2026 and 3,747,008 roll units, of which 613,153 distinct units are matched to at least one sale. The provider’s structural fields are the roll’s (year built identical in 99.2% of matched sales, floor area in 94.1%); its assessed value is the value in force at the sale date, which we use only for a screening rule. A full data dictionary and a generated provenance report accompany the repository.

Four earlier extractions of the same table (6 and 9 August, 10 and 17 October 2026) measure the publication process (Table 6). In every vintage more than 99.96% of the sales that the final snapshot dates on or before the vintage’s last observed date were already present, no row was deleted and no price revised; the feed publishes with a lag of 10 to 13 days and does not backfill older months. Consequently the only incompleteness in real time is the final two weeks, and the real-time experiments of Section 6 model a ten-day publication delay with a partial last month. The table has one date per sale; deed signing, registration and publication cannot be distinguished, a limitation shared with every registry-based index.

3.3. Coverage

Monthly volumes (Figure 1) start low (2,478 sales in January 2021, a feed start-up), average 10,557 per month over 2022–2025 with a strong June–July peak tied to Quebec’s moving-day calendar, and end with a partial October 2026 (2,722 sales in seven days). Volumes in 2026 run 8–14% below the same months of 2025. The provider collects no sale below \$50,000, a hard floor that truncates the lowest-price rural markets. Sales are geocoded, and every one falls inside a municipal polygon of the provincial administrative boundaries (SDA, version 2026-07); the roll’s municipality code agrees with the spatial join for 99.96% of matched sales. The series are on 2026 boundaries; municipalities constituted or merged during the sample are listed in the repository.

3.4. Classification and population definition

Property classes are built from the roll rather than from the provider label: a residential use code (1000) with one dwelling and an “integrated” physical link is a condominium; one dwelling with a detached, semi-detached or row link is a house; two to five dwellings a plex; six or more a multi-unit building; use codes 1100 and 1211 are cottages and mobile homes. The provider label is used only when the roll is missing or incomplete; conflicts between a complete roll record and the provider label (4.3% of rows) are flagged and the roll class kept. The headline population is houses, condominiums and plexes with roll characteristics. Table 2 lists the exclusion rules in priority order. Sales dated before the roll’s year of construction are lots or unfinished structures (the roll is a current vintage, so a lot sold in 2021 and built on in 2022 now carries the house’s attributes); vacant land, non-residential use codes, parcels sold more than ten times in six years (developer




Figure 1. Registered sales per month in the snapshot, January 2021 to October 2026 (all property types, before exclusions).

or shared parcels), same-day sales of the same parcel at different prices (undivided co-ownership shares and condominium projects), multi-parcel deeds (same day, price and street or 1-km cell), exact duplicates and prices above \$10 million are excluded. The price screen is relative: a sale is excluded when its log price-to-assessment ratio deviates by more than $\log 4$ from the median of its year, class and region, using the assessment in force at the sale. This is not the only outlier treatment; every estimator below also uses a robust likelihood. The headline sample is 632,006 sales (82.1% of the raw table), 422,658 houses, 122,989 condominiums and 86,359 plexes (Table 3).

Table 2. Population definition: exclusion ledger of the Vrai Prix snapshot (769,995 registered sales, 2021-01-04 to 2026-10-07)

Rule (first match, in priority order)	Rows	% of raw
Retained: headline sample	632,006	82.1
Class not in headline (cottage, mobile, 6+ units, unclassified)	38,280	5.0
Price/assessment ratio outside $\pm \log 4$ of year–class–region median	23,653	3.1
Vacant land / road / forest (roll)	18,798	2.4
Non-residential use code (roll)	15,520	2.0
Same parcel, same day, different prices	14,709	1.9
No assessment-roll characteristics	11,533	1.5
Parcel with more than 10 sales in 2021–2026	6,981	0.9
Roll building year after the sale year	4,996	0.6
Multi-parcel deed (same day, price, street or 1-km cell)	2,663	0.3
Exact duplicate row	824	0.1
Price above \$10 million	32	0.0

Rules are applied in the order listed; a row is counted once, under the first rule it triggers. Independent flag counts are in results/cleaning/flag_counts.csv.

Table 3. Headline estimation sample by year and property class

Year	Houses	Condominiums	Plexes	Total
2021	86,922	25,326	19,570	131,818
2022	72,580	20,972	15,566	109,118
2023	60,326	17,679	11,148	89,153
2024	68,426	20,758	13,199	102,383
2025	77,465	22,871	15,681	116,017
2026	56,939	15,383	11,195	83,517
Total	422,658	122,989	86,359	632,006

2026 covers January 1 to October 7 (partial).

3.5. Characteristics, missingness and what is not observed

The available quality variables are floor area of storeys, lot area, year of construction, number of storeys, building genre (bungalow, two-storey, split-level, mansard, one-module), physical link and, for plexes, the number of dwellings. Within the headline sample their completeness is high and stable over time (Table 4; the raw-row version in Table 33 shows lower shares in 2021 because land and unclassified rows are more frequent then), which is the condition for a time-dummy model: a variable whose missingness drifts creates artificial index steps. One provider field, the building style, was backfilled over the sample (12% present in January 2021, 87% in January 2025) and is banned from every model. Bedrooms, bathrooms, parking, basement finish, construction quality, renovations, waterfront, occupancy, financing and arm’s-length status are not observed, and no image or text features are used. The roll’s neighbourhood code turned out to be a unit-level identifier and is unusable; fine location is instead a nested metric grid (Section 5). Characteristics are a single current roll vintage: renovations after the sale are invisible, and new constructions of 2025–2026 may still be coded as land, which lowers new-build coverage at the end of the sample.

3.6. Geography

The hierarchy is province, 17 administrative regions, 103 MRCs and 1,125 municipalities with at least one sale (Table 5); metropolitan communities (CMM for Montréal, CMQ for Québec) and boroughs are carried for aggregation and diagnostics. Statistics Canada census metropolitan area boundaries could not be obtained from the agency’s site (an automated bot check blocked every download), so the metropolitan communities stand in for CMAs in external comparisons; the difference in coverage is documented. Within municipalities, location is controlled by the finest cell of a 500-m / 1-km / 2-km / 5-km grid in the NAD83 Quebec Lambert projection that contains at least fifteen sales of the class, so that no cell is a singleton and changing sale locations cannot masquerade as appreciation.

Table 4. Completeness of the two main characteristics and share of new constructions in the headline sample, by year and class

Headline sample (% of sales)	2021	2022	2023	2024	2025	2026
Houses: floor area	99.1	99.2	99.2	99.2	99.0	99.1
Houses: year built	99.0	99.1	99.1	99.1	98.9	99.0
Houses: new construction (						
Condominiums: floor area	98.2	98.4	98.4	98.6	98.4	98.5
Condominiums: year built	99.9	99.8	99.8	99.8	99.8	99.8
Condominiums: new construction (						
Plexes: floor area	97.7	97.7	97.7	97.8	97.8	97.9
Plexes: year built	97.9	97.7	97.9	98.0	98.0	98.0
Plexes: new construction (						

Lot area, storeys, physical link and building genre are complete for more than 90% of headline sales in every year (results/profile/completeness_by_year.csv). The provider's building style (12% present in January 2021, 87% in January 2025) is excluded from every model. New constructions (sold in the roll's construction year) fall to zero in 2025–2026 because the roll vintage lags construction.

Table 5. Municipalities by number of roll-matched sales, 2021–2026

Sales 2021–2026	Municipalities	Sales
0-9	25	152
10-24	72	1,328
25-49	141	5,303
50-99	223	16,760
100-249	304	49,238
250-499	150	52,304
500-999	75	53,475
1000-4999	99	195,777
5000+	20	355,639

Table 6. Earlier extractions of the transaction table compared with the final snapshot (publication lag audit)

Extraction	Rows	Last date	Lag (days)	Final rows $\leq$ last date	Share present	Deleted	Price revised
2026-08-06	745,119	2026-07-27	10	745,227	99.99		
2026-08-09	745,119	2026-07-27	13	745,227	99.99		
2026-10-10	766,596	2026-09-29	11	766,682	99.99		
2026-10-17	769,333	2026-10-06	11	769,678	99.96		

4. Estimand and measurement challenges

4.1. Definition

Let g index a geography in the hierarchy (province, region, MRC, municipality), $k \in \{\text{house, condominium, plex}\}$ a property class and t a month. Let $\mathcal{X}$ be the space of observable quality and location attributes and $F_{g,k}$ a fixed reference distribution of attributes for the pair (g, k) . The object of measurement is the quality-constant log price level

$$L_{g,k,t} = \int_{\mathcal{X}} \mathbb{E}[\log P \mid x, g, k, t] dF_{g,k}(x), \quad (1)$$

and the index is $I_{g,k,t} = 100 \exp(L_{g,k,t} - \bar{L}_{g,k,2021})$, where $\bar{L}_{g,k,2021}$ is the average over the months of 2021 (the primary normalisation); January 2021 = 100 is offered as a secondary normalisation. The base does not enter any growth rate: monthly and annual log changes $\Delta L_{g,k,t}$ and $\Delta_{12} L_{g,k,t}$ are invariant to it, which is verified in the code tests. Equation (1) is a price index for a *constant basket*, not for the stock or for the transacted properties; the stock enters only through aggregation weights (Section 9), and representative dollar prices are published separately as the exponential of the mean log price of the cell's sample at base-period prices, carried forward by the index.

4.2. Identification in a hedonic model

With a log-linear hedonic specification $\mathbb{E}[\log P \mid x, g, k, t] = f_k(x; \theta_t) + \alpha_{\ell(x)} + L_{g,k,t}$, the level $L_{g,k,t}$ is identified only up to the location of the reference distribution: a change in which dwellings sell changes the sample mean of $f_k(x; \theta_t) + \alpha_{\ell}$ and must not be read as a price change. Three design choices follow. First, location enters through fixed effects at a scale fine enough that the composition of sales within a cell is approximately constant in the attributes that are not observed; the coarser the cell, the more unobserved location quality leaks into the index. Second, the quality prices θ_t are allowed to drift smoothly: a floor-area elasticity that changes from 0.42 to 0.30 over five years (Table 14) would otherwise bias the index for cells whose sold dwellings are larger or smaller than the provincial reference. Third, the reference distribution is the cell's own sample over the whole period, so that the index is a within-cell comparison; common-support rules record the share of a cell's sales inside the central 98% of the class's floor-area and age distributions.

4.3. Measurement challenges specific to thin markets

A municipality with four sales a month yields a monthly mean residual with a standard error of roughly $0.19/\sqrt{4} \approx 0.10$ log points for houses, an order of magnitude larger than any plausible monthly price change. Any estimator that reports a monthly change from such a cell without pooling

is reporting noise, and any estimator that pools without saying how much is hiding the problem. The framework therefore reports, for every cell-month, the actual number of sales, the effective sample size implied by the posterior variance, the share of the estimate that comes from the local data (the Kalman gain), a 95% interval, and a grade. A model estimate for a month with zero local sales is a model-assisted interpolation and is labelled as such; cells whose effective sample size falls below five are labelled insufficiently identified and suppressed in the headline tables while remaining in the panel for transparency.

Two further challenges are specific to the data. The publication delay is short and backfill is absent, so real-time incompleteness is confined to the last month; but the roll's single vintage means that the characteristics are post-sale, and the \$50,000 floor truncates the price distribution in the cheapest markets. Both are discussed where they matter.

5. SHARP-HPI: framework and identification

5.1. Decomposition

For a sale i of class k in month $t(i)$, located in municipality $u(i)$, MRC $m(i)$ and region $r(i)$, and in fine-location cell $\ell(i)$, the model is

$$\log P_i = f_k(x_i; \theta_{k,t(i)}) + \alpha_{k,\ell(i)} + \mu_{k,t(i)} + \delta_{r(i),k,t(i)}^R + \delta_{m(i),k,t(i)}^M + \delta_{u(i),k,t(i)}^U + \varepsilon_i. \quad (2)$$

The function f_k is a flexible hedonic surface (cubic in centred log floor area, quadratic in centred log lot area, age in nine bins, storeys, building genre, physical-link subtype for houses and number of dwellings for plexes, with time-stable missingness indicators) whose floor-area and lot-area coefficients are allowed to move linearly in a centred time index; $\alpha_{k,\ell}$ is a location fixed effect at the finest grid cell with at least fifteen sales of the class; $\mu_{k,t}$ is the provincial path (a month effect); and the three δ terms are nested deviations of the region from the province, of the MRC from its region and of the municipality from its MRC. The quality-constant log level of municipality u is $L_{u,k,t} = \mu_{k,t} + \delta_{r(u),k,t}^R + \delta_{m(u),k,t}^M + \delta_{u,k,t}^U$; the level of an MRC omits the last term and that of a region the last two. The cell-level spatial residual $u_{\text{spatial},t}$ of the protocol is absorbed by $\alpha_{k,\ell}$ in the cross-section and by the nesting in time; its remaining extent is measured by Moran's I of the residuals (Section 9).

5.2. Temporal evolution of the deviations

Each deviation follows a local-level process,

$$\delta_{g,t} = \delta_{g,t-1} + \eta_{g,t}, \quad \eta_{g,t} \sim N(0, \tau_g^2), \quad (3)$$

and is observed through the monthly mean of the stage-one residuals of the cell net of its parent's deviation,

$$\bar{y}_{g,t} = \delta_{g,t} + e_{g,t}, \quad e_{g,t} \sim N(0, \sigma_g^2/n_{g,t}), \quad (4)$$

with no observation in months where $n_{g,t} = 0$. A local-linear-trend alternative adds a stochastic slope to (3); it is estimated as an ablation and did not improve on the local level (Section 7). The observation variance σ_g^2 is the cell's robust residual variance shrunk toward the class variance with weight $N_g/(N_g + 100)$ for cells with few sales.

5.3. Identification and constraints

Three constraints identify (2). (i) The month effects $\mu_{k,t}$ are normalised at zero in January 2021 and the location effects are absorbed, so that μ is the provincial path for a dwelling of fixed attributes in a fixed cell. (ii) The deviations are deviations from the parent: because the location fixed effects absorb each cell's mean residual over the sample, every $\delta_{g,\cdot}$ has zero sample mean by construction and a diffuse initial state; a region's level is therefore the province's plus a zero-mean relative movement, which is what makes the hierarchy additive and the indices coherent in logs. (iii) The innovation variances τ_g^2 are identified from the autocovariance of $\bar{y}_{g,t}$ only when the cell has several observed months; for the many cells where this is not the case the variance is calibrated across cells (next subsection). The real-time property follows from (3)–(4): the filtered estimate of $\delta_{g,t}$ uses $\bar{y}_{g,1}, \dots, \bar{y}_{g,t}$ only.

5.4. Adaptive pooling through empirical-Bayes calibration

The Kalman gain of (3)–(4) is $\kappa_{g,t} = p_{g,t|t-1}/(p_{g,t|t-1} + \sigma_g^2/n_{g,t})$, where $p_{g,t|t-1}$ is the predictive variance of the state. It is the weight of the month's local data in the updated estimate and it depends on the cell through τ_g^2 and σ_g^2 and on the month through $n_{g,t}$: a liquid month in a volatile cell gets a gain near one; an empty month gets zero and the estimate is the parent-consistent prediction. The adaptive element is the calibration of τ_g^2 . Estimating it by maximum likelihood cell by cell is unstable for thin cells (the profile likelihood is flat and often maximised at the boundary, which produces either a frozen deviation or one that chases noise). We instead place a normal prior on $\log \tau_g^2$ shared by all cells of a level, $\log \tau_g^2 \sim N(m_\ell, v_\ell)$, estimate (m_ℓ, v_ℓ) by a grid EM algorithm that maximises the marginal likelihood $\sum_g \log \int L_g(\tau^2) dN(\log \tau^2; m_\ell, v_\ell)$ over the cells of the level, and set each cell's τ_g^2 to its posterior mode. A liquid cell is driven by its own likelihood; a thin cell inherits the level's typical signal-to-noise ratio. This is the empirical-Bayes logic of [Efron and Morris \(1973\)](#) and [Fay and Herriot \(1979\)](#) applied to a dynamic variance component; it is not new, and its value is established by the ablation that replaces it with per-cell maximum likelihood.

5.5. Robust innovations

Transaction-level outliers are handled in stage one by Huber-type iteratively reweighted least squares ($k = 1.345$ on MAD-standardised residuals). Monthly mean residuals can still be contaminated when a cell has few sales, so the state-space step down-weights months whose standardised innovation exceeds 1.345 by inflating their observation variance (two passes), a pragmatic approximation to a Student- t observation density in the spirit of [Francke and van de Minne \(2017\)](#). The Gaussian variant is an ablation.

5.6. Reliability bundle

For every cell-month the panel records $n_{g,t}$; the effective sample size $\text{ESS}_{g,t} = \sigma_g^2 / \text{Var}(\hat{\delta}_{g,t})$, the number of transactions whose simple mean would have the posterior variance of the estimate; the local-data weight $\kappa_{g,t}$; the posterior standard error of the log level, which sums the state variances along the hierarchy and the stage-one variance of $\mu_{k,t}$ (treating levels as independent, an approximation examined by bootstrap); the covariate overlap; the lag-one autocorrelation of the observed deviations; and the one-step-ahead standardised innovation, whose root mean square should be one if the model is calibrated out of time. Grades A to E are thresholds on ESS and the posterior standard error, anchored on the Monte Carlo coverage results (Table 17). The local-data weight is a Kalman gain, not a probability or a quality score, and the panel documents it as such.

6. Estimation, inference, linking and publication vintages

6.1. Algorithm

Estimation is sequential and fast enough to be rerun monthly on a workstation (Table 19, last column). Stage one is a weighted least-squares regression with absorbed location fixed effects and explicit month dummies, iterated with Huber weights; it uses the full history available at the estimation date and returns the provincial path with its standard errors and the residuals. Stage two proceeds level by level: residual means and counts are computed per cell and month; the profile likelihood of each cell on a 48-point grid of $\log \tau^2$ between 10^{-8} and 2×10^{-2} is evaluated by the Kalman filter; the level's prior is estimated by grid EM; each cell is filtered and smoothed with its posterior-mode τ_g^2 and two robustness passes; its deviation (smoothed in the retrospective mode, filtered in the real-time mode) is subtracted from the residuals of its children before the next level. The whole panel for one class (province, 16 regions with sales, up to 103 MRCs and more than a thousand municipalities, 69 months) takes between ten and forty seconds.

6.2. Inference

Intervals are model-based: the variance of a municipal log level is the sum of the smoothed (or filtered) state variances of the municipality, its MRC and its region, plus the squared standard error of the provincial month effect. This ignores the uncertainty in the hedonic coefficients beyond μ , the estimation of σ_g^2 and τ_g^2 , and the covariance between levels. Two checks are reported: the empirical coverage of the intervals in the Monte Carlo experiments, where the truth is known, and a cluster bootstrap on the condominium class that resamples fine-location cells with replacement and re-estimates both stages (Section 9). Intervals are frequentist in spirit (the prior on $\log \tau^2$ is estimated, not elicited), and we call them confidence intervals; they are not posterior credible intervals for the index. Where many municipal changes are compared, we report the share of cells whose 12-month change excludes zero rather than individual significance, and the placebo experiment gives the false-positive rate of that rule.

6.3. Normalisation and linking

The retrospective series are not spliced: they are re-estimated on the full sample and normalised to the 2021 average. Real-time series are produced by re-estimating both stages at each monthly publication date on the data available then. Three publication policies are compared on the resulting vintage database: (a) full re-estimation, in which every published value may be revised; (b) a movement splice, in which only the newest monthly change is appended to a frozen history, as in RTD; (c) a revision-limited policy, in which a month is frozen after its third release, as in HMTS. The vintage database records every release of every cell-month, so the revision statistics of Section 9 are measured, not assumed.

6.4. Real-time vintages

Publication date v_m for month m is the last day of m plus ten days, the measured publication delay; the data available at v_m are the sales dated on or before $v_m - 10$ days, so the newest month is always partial. Releases are simulated monthly from July 2022 (eighteen months of history) to October 2026 for the province, the regions and the municipalities with at least 1,500 sales; the location partition and the empirical-Bayes priors are rebuilt from the available data at each vintage, so that no information from the future enters any real-time value. The October 2026 observation is partial through 7 October in every mode and is flagged as such in the panel.

6.5. Benchmarks

Eleven benchmark estimators are implemented on the same samples: (A) median and mean prices; (B) a stratified median with fixed weights over floor-area terciles $\times$ age terciles $\times$ subtype; (C)

pooled time dummies, ordinary and Huber; (D) rolling time dummies with 7-, 13- and 25-month windows and window, movement and mean splices; (E) chained and direct double-imputation Laspeyres, Paasche and Fisher indices from per-month hedonic regressions with coarse location effects; (F) rolling-window GEKS on bilateral Fisher links with a movement splice; (G) HMTS as described by [Ishaak et al. \(2024\)](#) and the Statistics Netherlands implementation, with a smooth-trend state-space step on each base-period series and windows of 4 and 13 months; (H) Bailey–Muth–Nourse and Case–Shiller repeat sales on pairs of the same assessment unit held at least six months; (I) static empirical-Bayes partial pooling of cell deviations, the mixed-effects benchmark, which is the special case of the state-space model with $\tau^2 = 0$; (J) a gradient-boosting residual index trained strictly out of time on data before each month, reported as a diagnostic because predictive accuracy is not index accuracy; and (K) a re-implementation of the production QWHPI v2.1 architecture (13-month rolling Huber time dummies with mean splice for the province, direct local time dummies for cells with a median of at least 40 sales per month, and a local-level Kalman deviation with per-cell maximum-likelihood τ^2 for thin cells).

7. Simulation design and Monte Carlo results

7.1. Design

Synthetic markets reproduce the features of the Quebec data that matter for index estimation (Appendix D): 60 municipalities nested in 12 MRCs and 4 regions, expected monthly sales drawn from a log-normal distribution with a median of 3.3 sales (a quarter of the municipalities sell less than once a month, a few sell more than 50 times), latent quality-constant log price paths with nested random-walk deviations, a hedonic surface with a nonlinear floor-area effect and an optional drifting elasticity, spatially correlated cell quality, Student- t noise with contamination, June–July seasonality in volumes, resales, missing floor areas and registration lags. Nine scenarios (Table 7) vary the data-generating process: a baseline; a turning point at month 36; a flat market in which the floor area of sold homes trends upward (the median must rise spuriously); a rising market in which sold-home quality declines; drifting hedonic prices with increasing missingness; heavy tails (2.5 degrees of freedom, 4% contamination); a sparse market with 35% of the sales; and two deliberate counter-examples, a market without local dynamics (where pooling should dominate) and a market with strong idiosyncratic municipal dynamics (where pooling should hurt and a direct estimator should win). Each scenario is replicated 80 times with new seeds; every estimator of Section 9 that can be computed on a municipality is run on every cell, and SHARP-HPI is run in its full form and in five ablations (no empirical-Bayes pooling, no time-varying prices, local linear trend, Gaussian innovations, municipal-level location effects), in smoothed and filtered modes. The truth for a municipality is its latent path; for a region or the province it is the stock-weighted mean of its

municipalities’ paths. Reported metrics are averaged over cells and replications with Monte Carlo standard errors over replications (Tables 8–13, Figures 2–3).

Table 7. Monte Carlo scenarios (60 months, 4 regions $\times$ 3 MRCs $\times$ 5 municipalities, log-normal municipality sizes)

Scenario	Drift (					
baseline		0.4	1.2/0.6/0.4/0.60	+0.0	+0.00	4.0 / 1
turning point	0.6 $\rightarrow$ -0.8 at m36		1.2/0.6/0.4/0.60	+0.0	+0.00	4.0 / 1
flat quality up		0.0	0.6/0.6/0.4/0.60	+0.4	+0.00	4.0 / 1
rising quality down		0.8	1.2/0.6/0.4/0.60	-0.4	+0.00	4.0 / 1
drifting hedonics		0.4	1.2/0.6/0.4/0.60	+0.0	+0.15	4.0 / 1
heavy tails		0.4	1.2/0.6/0.4/0.60	+0.0	+0.00	2.5 / 4
sparse		0.4	1.2/0.6/0.4/0.60	+0.0	+0.00	4.0 / 1
no local dynamics		0.4	1.2/0.1/0.1/0.05	+0.0	+0.00	4.0 / 1
idiosyncratic towns		0.4	1.2/0.2/0.4/2.00	+0.0	+0.00	4.0 / 1

Quality trend = monthly drift of the mean log floor area of sold homes (composition); β_{fa} drift = change of the floor-area elasticity over the sample; sales scale multiplies expected monthly sales (baseline median $\approx$ 3.3 per municipality).

7.2. Accuracy of monthly changes

At the municipal level the ranking is unambiguous and stable across scenarios (Table 8, Figure 2). In the baseline the RMSE of the monthly log change is 1.58 log points for SHARP-HPI, 2.91 for the QHPI architecture (per-cell maximum-likelihood deviations), 12.52 for a cell-level pooled time dummy, 6.01 for the stratified median and 24.28 for the median; repeat sales are estimable in only 30% of the municipal cell-months and are the noisiest where they exist. The composition scenarios show what quality adjustment buys: in the flat market with rising sold quality the median’s monthly changes carry a bias of 0.211 log points per month (about 3% per year of spurious appreciation), the hedonic estimators’ 0.052. The real-time (filtered) SHARP-HPI has an RMSE of 3.02, twice the smoothed value: half of the final accuracy at the municipal level comes from information that arrives after the month, which is the revision cost documented on the real data.

7.3. What each component contributes

The ablations identify the mechanisms. Empirical-Bayes calibration of the innovation variance improves the RMSE only slightly in the baseline (1.62 without it) but matters for calibration: the 95% intervals cover 98% of the true levels with it and 97% without, and in the heavy-tailed scenario 91% against 92%, because per-cell maximum likelihood collapses to the boundary in thin cells. The local-linear-trend variant is never better than the local level and doubles the revision between filtered and smoothed estimates (4.57 against 2.36 log points). Gaussian innovations cost 2.25

Table 8. Monte Carlo: RMSE of monthly log changes ($\times 100$) at the municipality level, by scenario (Monte Carlo standard error in parentheses)

Method	baseline	turning point	flat quality up	rising quality down	drifting hedonics	heavy tails	sparse	no local dynamics	idiosyncratic towns
Median	24.28 (0.16)	24.28 (0.16)	24.71 (0.16)	23.95 (0.16)	25.89 (0.17)	27.37 (0.22)	29.43 (0.13)	24.28 (0.16)	24.38 (0.15)
Stratified median	6.01 (0.02)	6.02 (0.02)	6.00 (0.02)	6.00 (0.02)	6.25 (0.02)	7.60 (0.03)	5.57 (0.03)	5.97 (0.02)	6.16 (0.02)
Pooled TD (Huber)	12.52 (0.10)	12.52 (0.10)	12.55 (0.10)	12.54 (0.10)	13.01 (0.11)	15.15 (0.16)	15.24 (0.14)	12.52 (0.10)	12.57 (0.11)
Rolling TD 13 m	7.48 (0.10)	7.48 (0.10)	7.50 (0.10)	7.48 (0.10)	7.75 (0.10)	8.14 (0.13)	7.84 (0.19)	7.48 (0.10)	7.48 (0.10)
Repeat sales	39.26 (0.45)	39.23 (0.46)	39.08 (0.44)	39.80 (0.47)	39.49 (0.45)	51.70 (0.99)	42.97 (0.84)	39.25 (0.45)	39.17 (0.42)
QHPI v2.1 arch.	2.91 (0.05)	2.91 (0.05)	2.90 (0.05)	2.91 (0.05)	2.93 (0.05)	7.89 (0.07)	3.57 (0.04)	2.84 (0.05)	3.30 (0.04)
SHARP-HPI (smoothed)	1.58 (0.02)	1.58 (0.02)	1.58 (0.02)	1.58 (0.02)	1.62 (0.02)	4.62 (0.08)	2.40 (0.04)	1.35 (0.02)	2.52 (0.02)
SHARP-HPI (filtered)	3.02 (0.04)	3.02 (0.04)	3.02 (0.04)	3.02 (0.04)	3.06 (0.04)	8.23 (0.10)	4.17 (0.06)	2.73 (0.04)	4.24 (0.03)
SHARP: no EB pooling of τ^2	1.62 (0.02)	1.62 (0.02)	1.62 (0.02)	1.62 (0.02)	1.66 (0.02)	4.51 (0.08)	2.45 (0.04)	1.40 (0.02)	2.55 (0.02)
SHARP: fixed quality prices	1.58 (0.02)	1.58 (0.02)	1.58 (0.02)	1.58 (0.02)	1.61 (0.02)	4.62 (0.08)	2.40 (0.04)	1.35 (0.02)	2.52 (0.02)
SHARP: local linear trend	1.61 (0.02)	1.61 (0.02)	1.61 (0.02)	1.61 (0.02)	1.65 (0.02)	4.95 (0.08)	2.45 (0.04)	1.40 (0.02)	2.51 (0.02)
SHARP: Gaussian	2.25 (0.04)	2.25 (0.04)	2.24 (0.04)	2.23 (0.04)	2.29 (0.04)	6.77 (0.09)	3.21 (0.05)	2.13 (0.04)	2.84 (0.03)
SHARP: municipality-level location FE	2.05 (0.04)	2.05 (0.04)	2.05 (0.04)	2.05 (0.04)	2.08 (0.04)	2.84 (0.06)	3.30 (0.08)	1.88 (0.04)	2.81 (0.03)

Averages over cells and replications; see Table 7 for the scenarios. Lower is better; the final two scenarios are the deliberate counter-




Figure 2. Monte Carlo: RMSE of monthly log changes (log scale) at the municipality and region levels by scenario and method; whiskers are Monte Carlo standard errors.

against 1.58 in the baseline and 6.77 against 4.62 with heavy tails. Fine location effects matter in every scenario but one: replacing them by municipal effects raises the RMSE to 2.05 in the baseline. The exception is the heavy-tailed scenario, the first counter-example, where the municipal-location variant is the best estimator (2.84 against 4.62): with 4% contamination and 2.5 degrees of freedom, fine cells of fifteen sales absorb outliers into their fixed effects and transmit them to the residual means, and a coarser partition is more robust. Time-varying quality prices change nothing in the baseline and help marginally when prices drift.

7.4. Counter-examples and the regional level

The second counter-example, strong idiosyncratic municipal dynamics, did not produce the reversal it was designed for at the municipal level: the smoothed SHARP-HPI remains the most accurate (2.52 against 3.30 for per-cell estimation) because the empirical-Bayes prior is wide enough to let volatile cells keep large innovation variances. The cost appears in real time: the filtered estimate lags the truth by +0.48 months by cross-correlation, against +0.11 in the baseline, and its coverage falls to 95%. The regional level (Table 9) is where SHARP-HPI's design costs accuracy: with 60 to 1,000 sales a month per region, the QHPI architecture, which estimates each liquid region with its own hedonic coefficients and smooths lightly, has an RMSE of 1.17 against 1.40 for SHARP-HPI, whose regional path is the pooled provincial path plus a smoothed deviation; the Monte Carlo confirms the repeat-sales evidence of Section 9 that region-specific estimation is preferable where the data allow it, and the panel's regional series should be read with that caveat. At the provincial level all hedonic estimators agree within 0.1 log point and the medians are three times noisier (Table 10); the stock-weighted bottom-up aggregate of municipal SHARP indices is as accurate as the direct provincial estimate.

Table 9. Monte Carlo: RMSE of monthly log changes ($\times 100$) at the region level, by scenario (Monte Carlo standard error in parentheses)

Method	baseline	turning point	flat quality up	rising quality down	drifting hedonics	heavy tails	sparse	no local dynamics	idiosyncratic towns
Median	7.24 (0.16)	7.24 (0.16)	7.34 (0.16)	7.19 (0.16)	7.56 (0.16)	7.34 (0.17)	12.37 (0.26)	7.20 (0.15)	7.49 (0.14)
Stratified median	6.03 (0.13)	6.03 (0.13)	6.12 (0.13)	6.09 (0.14)	6.18 (0.14)	6.35 (0.14)	9.12 (0.16)	6.01 (0.13)	6.47 (0.13)
Pooled TD (Huber)	2.63 (0.04)	2.63 (0.04)	2.64 (0.04)	2.63 (0.04)	2.73 (0.04)	2.56 (0.04)	4.64 (0.07)	2.61 (0.04)	2.89 (0.05)
Rolling TD 13 m	2.99 (0.05)	2.99 (0.05)	3.00 (0.05)	2.99 (0.05)	3.10 (0.05)	2.99 (0.05)	5.45 (0.09)	2.99 (0.05)	3.11 (0.05)
Repeat sales	12.03 (0.21)	12.03 (0.21)	12.05 (0.21)	12.04 (0.20)	12.18 (0.21)	15.80 (0.30)	22.22 (0.39)	12.00 (0.21)	12.59 (0.22)
QHPI v2.1 arch.	1.17 (0.01)	1.23 (0.01)	0.84 (0.01)	1.25 (0.01)	1.19 (0.01)	1.21 (0.01)	2.14 (0.06)	1.08 (0.01)	1.27 (0.01)
SHARP-HPI (smoothed)	1.40 (0.02)	1.40 (0.02)	1.41 (0.02)	1.40 (0.02)	1.44 (0.02)	2.32 (0.06)	2.24 (0.04)	1.23 (0.02)	1.58 (0.03)
SHARP-HPI (filtered)	1.58 (0.02)	1.58 (0.02)	1.58 (0.02)	1.58 (0.02)	1.62 (0.02)	2.94 (0.07)	2.44 (0.04)	1.30 (0.02)	1.76 (0.03)
SHARP: no EB pooling of τ^2	1.39 (0.02)	1.39 (0.02)	1.40 (0.02)	1.39 (0.02)	1.43 (0.02)	2.16 (0.06)	2.24 (0.04)	1.24 (0.02)	1.58 (0.03)
SHARP: fixed quality prices	1.40 (0.02)	1.40 (0.02)	1.41 (0.02)	1.40 (0.02)	1.44 (0.02)	2.32 (0.06)	2.23 (0.04)	1.23 (0.02)	1.58 (0.03)
SHARP: local linear trend	1.39 (0.02)	1.39 (0.02)	1.40 (0.02)	1.39 (0.02)	1.43 (0.02)	2.34 (0.06)	2.23 (0.04)	1.23 (0.02)	1.57 (0.03)
SHARP: Gaussian	1.70 (0.03)	1.70 (0.03)	1.70 (0.03)	1.70 (0.03)	1.74 (0.03)	2.72 (0.07)	2.70 (0.05)	1.58 (0.03)	1.84 (0.03)
SHARP: municipality-level location FE	1.97 (0.04)	1.97 (0.04)	1.97 (0.04)	1.97 (0.04)	2.01 (0.04)	2.18 (0.06)	3.24 (0.08)	1.87 (0.04)	2.11 (0.04)

Averages over cells and replications; see Table 7 for the scenarios. Lower is better; the final two scenarios are the deliberate counter-

Table 10. Monte Carlo: RMSE of monthly log changes ($\times 100$) at the province level, by scenario (Monte Carlo standard error in parentheses)

Method	baseline	turning point	flat quality up	rising quality down	drifting hedonics	heavy tails	sparse	no local dynamics	idiosyncratic towns
Median	3.56 (0.08)	3.56 (0.08)	3.62 (0.08)	3.55 (0.08)	3.69 (0.08)	3.59 (0.09)	6.07 (0.13)	3.56 (0.08)	3.62 (0.08)
Stratified median	3.08 (0.07)	3.08 (0.07)	3.13 (0.08)	3.06 (0.07)	3.14 (0.07)	3.11 (0.08)	5.21 (0.13)	3.07 (0.07)	3.15 (0.09)
Pooled TD (Huber)	1.24 (0.02)	1.24 (0.02)	1.25 (0.02)	1.24 (0.02)	1.29 (0.03)	1.24 (0.02)	2.13 (0.04)	1.22 (0.02)	1.37 (0.03)
Rolling TD 13 m	1.39 (0.03)	1.39 (0.03)	1.40 (0.03)	1.39 (0.03)	1.44 (0.03)	1.41 (0.03)	2.36 (0.05)	1.39 (0.03)	1.44 (0.03)
Repeat sales	5.32 (0.10)	5.32 (0.10)	5.33 (0.10)	5.33 (0.10)	5.37 (0.10)	7.42 (0.17)	9.57 (0.21)	5.28 (0.10)	5.66 (0.13)
QHPI v2.1 arch.	1.39 (0.03)	1.39 (0.03)	1.40 (0.03)	1.39 (0.03)	1.44 (0.03)	1.41 (0.03)	2.36 (0.05)	1.39 (0.03)	1.44 (0.03)
SHARP-HPI (smoothed)	1.25 (0.02)	1.25 (0.02)	1.25 (0.02)	1.25 (0.03)	1.29 (0.03)	1.26 (0.02)	2.14 (0.04)	1.22 (0.02)	1.37 (0.03)
SHARP: no EB pooling of τ^2	1.25 (0.02)	1.25 (0.02)	1.25 (0.02)	1.25 (0.03)	1.29 (0.03)	1.26 (0.02)	2.14 (0.04)	1.22 (0.02)	1.37 (0.03)
SHARP: fixed quality prices	1.25 (0.02)	1.25 (0.02)	1.26 (0.02)	1.25 (0.03)	1.29 (0.03)	1.26 (0.02)	2.13 (0.04)	1.22 (0.02)	1.37 (0.03)
SHARP: local linear trend	1.25 (0.02)	1.25 (0.02)	1.25 (0.02)	1.25 (0.03)	1.29 (0.03)	1.26 (0.02)	2.14 (0.04)	1.22 (0.02)	1.37 (0.03)
SHARP: Gaussian	1.59 (0.03)	1.59 (0.03)	1.59 (0.03)	1.59 (0.03)	1.63 (0.03)	2.25 (0.05)	2.62 (0.05)	1.57 (0.03)	1.68 (0.04)
SHARP: municipality-level location FE	1.88 (0.04)	1.88 (0.04)	1.88 (0.04)	1.88 (0.04)	1.92 (0.05)	1.96 (0.05)	3.18 (0.09)	1.87 (0.04)	1.96 (0.05)
SHARP: stock-weighted bottom-up	1.26 (0.02)	1.26 (0.02)	1.27 (0.02)	1.26 (0.03)	1.31 (0.03)	1.75 (0.04)	2.15 (0.04)	1.23 (0.02)	1.40 (0.03)

Averages over cells and replications; see Table 7 for the scenarios. Lower is better; the final two scenarios are the deliberate counter-

Table 11. Monte Carlo: empirical coverage (%) of nominal 95% intervals for municipal index levels

Method	baseline	turning point	flat quality up	rising quality down	drifting hedonics	heavy tails	sparse	no local dynamics	idiosyncratic towns
Pooled TD (Huber)	97	97	97	97	97	97	96	97	97
QHPI v2.1 arch.	89	89	87	87	89	83	89	90	89
SHARP-HPI (smoothed)	98	98	97	98	98	92	99	99	96
SHARP-HPI (filtered)	97	97	96	96	97	89	98	98	95
SHARP: no EB pooling of τ^2	97	97	96	96	98	91	98	99	93
SHARP: local linear trend	97	97	96	96	97	90	97	98	93
SHARP: Gaussian	98	98	97	97	98	92	98	98	97

Table 12. Monte Carlo: mean bias of monthly log changes ($\times 100$) at the municipality level

Method	baseline	turning point	flat quality up	rising quality down	drifting hedonics	heavy tails	sparse	no local dynamics	idiosyncratic towns
Median	-0.043	-0.043	+0.211	-0.282	-0.040	+0.093	+0.049	-0.043	-0.002
Pooled TD (Huber)	-0.018	-0.018	+0.029	-0.064	-0.019	-0.016	-0.012	-0.018	-0.005
QHPI v2.1 arch.	-0.001	-0.000	+0.046	-0.045	+0.001	-0.005	-0.005	-0.001	-0.010
SHARP-HPI (smoothed)	+0.003	+0.003	+0.052	-0.043	+0.004	-0.005	+0.004	+0.002	-0.005
SHARP-HPI (filtered)	+0.001	+0.001	+0.050	-0.045	+0.002	-0.012	+0.008	+0.001	+0.005

Table 13. Monte Carlo: timing diagnostics at the municipality level

Scenario	Reps	SHARP lag (months)	SHARP revision RMSE	QHPI lag	sec./rep
baseline	80	+0.01	2.36	+0.01	—
turning point	80	+0.00	2.36	+0.02	—
flat quality up	80	+0.01	2.37	+0.03	—
rising quality down	80	+0.01	2.37	+0.02	—
drifting hedonics	80	+0.01	2.39	+0.01	—
heavy tails	80	-0.01	4.80	-0.00	—
sparse	80	+0.02	3.16	+0.01	—
no local dynamics	80	+0.00	2.20	-0.00	—
idiosyncratic towns	80	+0.02	3.10	-0.00	—

Lag = argmax of the cross-correlation between estimated and true 3-month changes (negative = estimate lags).
Revision RMSE = RMSE between filtered (real-time) and smoothed monthly changes ($\times 100$).




Figure 3. Monte Carlo: empirical coverage of nominal 95% intervals for municipal index levels.

8. Quebec monthly indices, 2021–2026

8.1. Stage one: quality prices

Table 14 reports the stage-one regressions. The within-cell fit is 0.53 for houses, 0.63 for condominiums and 0.48 for plexes, with robust residual standard deviations of 0.185, 0.134 and 0.214 log points. Floor area carries the expected convexity (Figure 4), age profiles show a premium for near-new dwellings and depreciation that flattens after sixty years for houses and reverses for old condominiums in central locations, and the time interactions are negative for floor area in all three classes: the elasticity of price with respect to floor area fell during the sample (-0.134 per unit of the centred time index for houses), which is consistent with smaller dwellings appreciating faster than larger ones. A fixed-coefficient model (the “no time-varying prices” ablation) changes municipal monthly changes by 0.48 log points on average, little, but the drift is real and costless to allow.

8.2. Stage two: calibrated signal-to-noise ratios

Table 15 reports the empirical-Bayes priors. The prior median innovation standard deviation of regional deviations is 3.66% per month for houses with a tight prior variance (0.2), that of municipalities 7.32% with a wide prior (1.8): municipalities differ genuinely in how much they move relative to their MRC, and the calibration lets liquid cells keep their own estimate while thin cells inherit the typical ratio. Figure 5 shows the consequence: the mean local-data weight rises with liquidity from near zero below one sale a month to above 0.8 in the largest cities, and the effective sample size tracks the actual count for liquid cells and exceeds it for thin cells, which is the pooling.

Table 14. Stage-1 robust hedonic regressions (selected coefficients; dependent variable log price)

Regressor	Houses	Condominiums	Plexes
log floor area (centred)	0.418	0.697	0.326
log floor area ²	0.039	0.107	-0.015
log floor area ³	0.048	-0.017	-0.025
log lot area (centred)	0.131	0.043	0.144
log lot area ²	-0.006	-0.018	-0.002
Age 2–5	0.182	0.014	0.205
Age 11–20	0.076	-0.127	0.022
Age 21–40	-0.039	-0.219	-0.130
Age 61–80	-0.281	-0.222	-0.342
Age 120+	-0.365	-0.201	-0.429
Two storeys	-0.039	0.017	0.043
Semi-detached	-0.091	—	—
Row	-0.131	—	—
Triplex	—	—	0.066
5 units	—	—	0.207
log floor area $\times$ time	-0.134	-0.091	-0.119
log lot area $\times$ time	0.046	0.123	0.085
Floor area missing	-0.088	-0.095	0.064
Age missing	-0.088	-0.039	-0.088
Observations	422,196	122,914	86,245
Robust residual s.d.	0.185	0.134	0.214
Within R^2	0.533	0.629	0.480

Huber-IRLS ($k = 1.345$) with fine-location fixed effects (nested metric grid, finest cell with at least 15 sales) and month effects; time index tt runs from -0.5 (January 2021) to 0.5 (October 2026). Full coefficient lists in results/panel/stage1_meta.json.




Figure 4. Stage-one quality profiles: age effects (left) and floor-area effects (right), by class.

Table 15. Empirical-Bayes priors for the innovation standard deviation of cell deviations

Class / level	Prior median τ (	
house / region	3.66	0.24
house / mrc	2.07	5.35
house / municipality	7.32	1.79
condo / region	1.89	0.76
condo / mrc	1.89	7.98
condo / municipality	1.18	11.30
plex / region	2.81	0.07
plex / mrc	1.70	7.63
plex / municipality	4.16	3.73

Grid EM over all cells of a level with at least six observed months; τ is the monthly standard deviation of the latent deviation from the parent path.




Figure 5. Municipal cells: mean local-data weight (Kalman gain) and mean effective sample size against median monthly sales.

8.3. The panel

The panel contains 2,654 geography–class cells (2,321 municipal) over 69 months. The provincial house index stands at 154.9 in September 2026 (2021 average = 100), condominiums at 136.4 and plexes at 151.2, with twelve-month changes of 6.1%, 1.3% and 4.8%. Table 16 gives the twelve largest municipalities with intervals and grades, Figure 6 six municipalities spanning the liquidity range, and Figure 7 the map of the municipal house index with the reliability hatching. Across the 55 municipalities graded A to C for houses in September 2026, the index ranges from 133 at the tenth percentile to 174 at the ninetieth (median 153): the dispersion of five-year appreciation across Quebec municipalities is as large as the provincial appreciation itself. Table 17 gives the distribution of grades: most regional and MRC cell-months are graded A to C, while a majority of municipal cell-months are graded D or E and are published for transparency only. The series are not seasonally adjusted; the monthly profile visible in Figure 6 (dips in winter months) is a composition effect in unobserved attributes that the hedonic controls do not remove and that users of monthly changes should treat as seasonal.

Table 16. SHARP-HPI, September 2026 (2021 average = 100), twelve largest municipalities by sales: index [95% interval] grade

Municipality	Houses	Condominiums	Plexes
Montréal	132.5 [129, 137] B	128.5 [126, 131] A	135.7 [130, 141] B
Québec	166.8 [162, 171] A	152.9 [148, 158] B	164.2 [156, 172] B
Laval	142.3 [139, 146] A	132.9 [129, 137] A	146.9 [139, 155] C
Gatineau	142.1 [138, 147] B	127.2 [123, 131] B	145.7 [138, 154] C
Longueuil	137.3 [132, 143] B	133.5 [130, 137] A	133.8 [126, 142] C
Sherbrooke	173.9 [167, 181] B	158.7 [152, 166] B	150.7 [140, 162] D
Saguenay	167.3 [161, 173] B	152.6 [146, 160] B	168.8 [159, 180] C
Lévis	169.6 [164, 176] B	151.2 [145, 158] B	165.7 [155, 177] C
Trois-Rivières	188.1 [180, 197] B	170.8 [162, 181] C	204.6 [189, 221] D
Terrebonne	152.4 [147, 158] B	133.1 [128, 138] B	153.5 [142, 166] D
Saint-Jérôme	154.3 [148, 161] B	145.4 [140, 151] B	155.2 [145, 166] C
Saint-Jean-sur-Richelieu	152.9 [147, 159] B	136.0 [129, 143] C	147.5 [138, 158] D


8.4. Aggregation and coherence

Section 4 defined the headline as a constant-basket index. The provincial path $\mu_{k,t}$ is estimated on all sales and is therefore transaction-weighted across municipalities. Table 18 compares it with bottom-up aggregates of the municipal indices weighted by the housing stock (dwelling units or assessed value, from the roll), by transactions, and equally. The stock-weighted house aggregate differs from the direct provincial path by +1.8 log points over 2021–2026 (root mean squared gap 1.3 per month) and the equal-weighted aggregate by +6.7: small municipalities appreciated

Table 17. Distribution of reliability grades across cell-months (percent of cell-months)

Level / class	A	B	C	D	E	Cell-months
mrc / condo	10	13	23	10	45	5,600
mrc / house	10	47	33	8	2	7,140
mrc / plex	0	7	34	37	22	7,070
municipality / condo	3	12	14	12	59	17,150
municipality / house	0	2	5	10	83	77,630
municipality / plex	0	0	1	4	94	67,690
province / condo	100	0	0	0	0	70
province / house	100	0	0	0	0	70
province / plex	1	99	0	0	0	70
region / condo	32	24	30	1	12	1,120
region / house	57	33	7	3	0	1,190
region / plex	0	43	41	10	6	1,190

A: effective sample size ≥ 150 and posterior s.e. of the log level $< 1.5\%$; B: ≥ 75 and $< 2.5\%$; C: ≥ 30 and $< 3.5\%$; D: ≥ 15 and $< 5\%$; E otherwise (published for transparency, flagged).



Shaded band: 95 % interval. Grey bars (bottom of each panel): monthly sales count.

Figure 6. SHARP-HPI house indices with 95% intervals for six municipalities of decreasing liquidity (smoothed, retrospective mode).

SHARP-HPI municipal house index, September 2026 (hatched: reliability grade D/E; grey: suppressed or no sales)




Figure 7. Municipal house index in September 2026 (2021 = 100) on 2026 municipal boundaries; hatched cells are graded D or E.

more than large ones, so the choice of weights is a choice of estimand, not a reconciliation error. A reconciliation projection that forces the direct and bottom-up series to agree would move the provincial series toward whichever weighting is imposed; we do not impose it and report the gap instead.

Table 18. Provincial aggregation: stock-weighted, transaction-weighted and equal-weighted bottom-up aggregates of municipal indices versus the direct provincial path (log $\times 100$, January 2021 to September 2026)

Class	Weights	Municipalities	Bottom-up total change	Direct total change	RMSE vs direct	Max g
house	Housing stock (units)	1093	58.8	56.9	1.32	
house	Housing stock (assessed value)	1093	54.9	56.9	1.70	
house	Transactions 2021–2026	1109	59.5	56.9	1.41	
house	Equal municipality weights	1109	63.6	56.9	2.29	
condo	Housing stock (units)	244	41.8	44.1	0.89	
condo	Housing stock (assessed value)	244	40.1	44.1	1.51	
condo	Transactions 2021–2026	245	43.4	44.1	0.58	
condo	Equal municipality weights	245	49.7	44.1	3.20	
plex	Housing stock (units)	962	49.1	52.8	1.64	
plex	Housing stock (assessed value)	962	45.6	52.8	3.29	
plex	Transactions 2021–2026	967	51.3	52.8	0.83	
plex	Equal municipality weights	967	54.5	52.8	1.84	

9. Benchmarks, validation, ablation and robustness

9.1. Benchmarks at the provincial and regional levels

Table 19 reports every benchmark at the provincial level: the cumulative log change from January 2021 to September 2026 and the standard deviation of monthly log changes, with wall-clock times.

Figure 8 plots the series. Four facts stand out. First, the quality-adjusted methods agree on the magnitude of the house appreciation: pooled time dummies, rolling time dummies with any window and splice, hedonic imputation, GEKS, HMTS, the QHPI architecture and SHARP-HPI all lie between 55.0 and 60.1 log points, against 46.4 for the median and 67.6 for Case–Shiller repeat sales. The median understates because the sold houses became smaller and older over the sample; repeat sales overstate because resold houses are renovated. Second, the chained imputation indices drift relative to their direct counterparts for condominiums and plexes (Laspeyres 45.4 against Paasche 36.8 for condominiums), the classical chain-drift of monthly chaining with per-period hedonic functions estimated on a few thousand sales. Third, methods differ far more in month-to-month volatility than in trend: the standard deviation of monthly changes is about 1.9 log points for the time-dummy family and SHARP, 2.8 for imputation and GEKS, 3–4 for the median and repeat sales, and 1.3–1.9 for HMTS, whose state-space smoothing of each base-period series is visible. Fourth, the gradient-boosting residual index, which forecasts each month’s prices from models trained on earlier months, rises by only 46.9 log points: a predictive model that learns location and time jointly absorbs part of the appreciation into its fitted surface, which is why predictive accuracy is not index accuracy. SHARP-HPI with and without time-varying quality prices differ by less than one log point in total at the provincial level; the drift matters below the province.




Figure 8. Provincial indices by method (2021 average = 100), January 2021 to September 2026.

Table 20 moves to the 16 administrative regions, where sample sizes range from 60 to 1,200 sales a month and the methods that estimate a separate hedonic function per region become unstable. The regional comparison has no truth; we report the volatility of each series and its root mean squared distance to the regional Case–Shiller index over 12-month changes. Three results are relevant to the design. The median and stratified median are two to three times as volatile as the hedonic methods and the farthest from repeat sales. Among hedonic methods, the QHPI architecture, which fits each liquid region with its own coefficients and smooths lightly, is the least volatile (2.33 for houses)

Table 19. Province-level benchmark indices, January 2021 to September 2026: cumulative log change ($\times 100$) and standard deviation of monthly log changes ($\times 100$)

Method	H: total	H: s.d. Δ	C: total	C: s.d. Δ	P: total	P: s.d. Δ	sec.
Median price	46.4	3.97	36.0	2.09	50.9	5.74	0
Mean price	47.8	3.79	36.3	2.47	43.9	4.16	0
Stratified median	54.0	3.35	37.3	1.88	49.6	4.61	0
Pooled time dummy (OLS)	58.5	2.58	40.9	1.35	48.7	2.42	1
Pooled time dummy (Huber)	56.5	1.86	41.0	1.06	51.4	1.94	4
Rolling TD, 7 m, mean splice	56.1	1.89	37.3	1.01	51.0	2.06	12
Rolling TD, 13 m, mean splice	55.6	1.87	37.2	1.02	50.6	1.91	18
Rolling TD, 13 m, movement splice	55.0	1.86	37.1	1.00	51.8	1.91	18
Rolling TD, 13 m, window splice	56.5	1.91	38.7	1.13	52.0	2.04	18
Rolling TD, 25 m, mean splice	56.4	1.86	39.4	1.03	51.1	1.87	32
Hedonic imputation Laspeyres (chained)	56.4	2.76	45.4	1.50	55.6	2.44	5
Hedonic imputation Paasche (chained)	56.5	2.76	36.8	1.45	57.8	2.53	5
Hedonic imputation Fisher (chained)	56.5	2.75	41.1	1.45	56.7	2.44	4
Hedonic imputation Fisher (direct)	59.0	2.76	41.5	1.69	50.1	2.30	4
GEKS-Fisher, 13 m	58.8	2.78	41.0	1.47	52.3	2.48	16
GEKS-Fisher, 25 m	59.0	2.78	41.1	1.47	51.8	2.45	27
HMTS, 4 m	59.2	1.93	42.4	1.06	53.2	1.56	—
HMTS, 13 m	60.1	1.92	42.2	0.90	54.9	1.29	—
Repeat sales (BMN)	68.6	3.02	38.6	1.71	57.9	4.58	0
Repeat sales (Case–Shiller)	67.6	2.97	37.6	1.73	58.5	4.61	0
GBM out-of-time residual	46.9	2.54	36.0	1.46	39.3	2.32	191
QHPI v2.1 architecture	55.6	1.87	37.2	1.02	50.6	1.91	18
SHARP stage 1 without time-varying prices	56.7	1.91	41.0	1.06	51.3	1.98	21
SHARP-HPI	56.9	1.91	41.5	1.04	50.9	1.97	21

H = houses, C = condominiums, P = plexes. Rolling TD = rolling time dummy; GEKS on bilateral double-imputation Fisher links with movement splice; HMTS = hedonic multilateral time series (Ishaak et al., 2024); QHPI = re-implementation of the QWHPI v2.1 architecture; last column = wall-clock seconds for the house sample on one Apple M-series core.

and the closest to repeat sales (9.9); SHARP-HPI, whose regional path is the pooled provincial path plus a smoothed deviation, is less volatile than the pooled time dummy (3.75 against 4.41) but not closer to repeat sales than it (12.7 against 12.6). The reason is identifiable: with 400 to 1,200 sales a month the regional Kalman gain is close to one, so the SHARP regional path is essentially the pooled model’s monthly regional mean, and the pooled model constrains quality prices to be common across regions. The Monte Carlo experiments of Section 7, where the truth is known, show that this costs little when quality prices are common and that regional estimates are then close to the truth; on the real data the repeat-sales arbiter, biased though it is, suggests that region-specific quality prices carry information at the regional level. Figure 9 shows six regions.

Table 20. Region-level benchmarks (16 administrative regions): mean standard deviation of monthly log changes ($\times 100$) and RMSE of 12-month log changes against the Case–Shiller repeat-sales index ($\times 100$)

Method	H: s.d. Δ	H: RMSE vs RS	C: s.d. Δ	C: RMSE vs RS	P: s.d. Δ	P: RMSE vs RS
Median price	7.91	13.7	9.17	13.0	13.27	27.5
Stratified median	6.16	13.2	6.21	12.2	10.00	26.0
Pooled time dummy (Huber)	4.41	12.6	4.65	11.1	7.11	24.6
Rolling TD, 13 m, mean splice	4.87	12.8	5.11	12.0	8.91	24.7
Hedonic imputation Fisher (chained)	5.42	12.0	3.79	7.6	7.23	17.8
GEKS-Fisher, 25 m	5.19	9.7	3.69	6.5	6.30	13.7
HMTS, 13 m	1.59	10.1	0.95	6.2	1.09	14.2
QHPI v2.1 architecture	2.33	9.9	1.48	6.6	2.20	14.5
SHARP-HPI (real-time filtered)	4.42	12.8	2.28	11.2	3.61	24.7
SHARP-HPI	3.75	12.7	1.60	11.0	2.56	24.8

Repeat sales are an independent (pair-based) arbiter, not the truth: they carry renovation and selection biases. Lower s.d. is not better per se; the Monte Carlo (Section 7) separates noise from signal.

9.2. Agreement with repeat sales

Table 21 compares SHARP-HPI with Case–Shiller repeat-sales indices estimated on the same cells where at least 300 pairs exist. At the provincial level the two house series have a correlation of 0.99 in levels; across regions the median correlation is 0.95 and across the 20 municipalities with enough pairs 0.81. The repeat-sales total change exceeds the hedonic one by a median of -0.1 log points for houses at the regional level and -22.9 for condominiums. The sign and the class pattern are the classical ones: repeat sales of houses embed renovations and the selection of resold dwellings (Clapp et al., 1991; Gatzlaff and Haurin, 1997; Nowak and Smith, 2020), which bias them upward relative to a constant-quality index, while condominiums are renovated less. Repeat sales are an independent arbiter of the movement, not the truth, and the Monte Carlo evidence is the only place where the truth is known.




Figure 9. Regional house indices by method for six regions of decreasing liquidity (2021 = 100).

Table 21. Agreement between SHARP-HPI and Case-Shiller repeat-sales indices estimated on the same cells (cells with at least 300 pairs)

Level / class	Cells	Pairs	Corr. levels (median)	Corr. 12-m changes (median)	RMSE levels (median, $\times 100$)	RS
province / house	1	45,354	0.99	0.90	3.0	
province / condo	1	18,518	0.99	0.94	1.6	
province / plex	1	9,508	0.98	0.56	3.3	
region / house	16	45,242	0.95	0.41	6.4	
region / condo	8	17,471	0.89	0.68	4.8	
region / plex	10	8,603	0.86	0.35	10.2	
municipality / house	24	18,982	0.81	0.37	12.6	
municipality / condo	7	11,212	0.88	0.65	4.9	
municipality / plex	5	3,702	0.86	0.33	8.9	

Repeat-sales pairs: same assessment unit, holding period of at least six months, both sales in the headline sample.

9.3. Composition shock

Table 22 and Figure 10 report a synthetic quality-composition shock: in four regions, the smallest 40% of houses (by floor area) are removed from the sample from January 2024 onward, which raises the median price mechanically. The median index records a spurious 2023-to-2024 jump of +23.4 log points in the Montréal region and +12.5 in the Capitale-Nationale; the pooled time dummy moves by -0.5 and -0.4, and SHARP-HPI by +0.3 and -0.7. The hedonic controls remove the shock almost entirely; the residual effect is the extrapolation of the floor-area curve to a truncated support.

Table 22. Synthetic composition shock (houses): the smallest 40% of houses are removed from the sample from January 2024; change in the 2023-to-2024 mean log index ($\times 100$)

Region	Method	2023→2024 change, original	Same, shocked sample	Shock effect
Region 03	Median price	+9.5	+22.0	+12.5
Region 03	Pooled time dummy (Huber)	+9.3	+8.9	-0.4
Region 03	SHARP-HPI	+9.3	+8.6	-0.7
Region 06	Median price	+4.0	+27.4	+23.4
Region 06	Pooled time dummy (Huber)	+3.3	+2.8	-0.5
Region 06	SHARP-HPI	+3.2	+3.5	+0.3
Region 07	Median price	+5.7	+18.7	+12.9
Region 07	Pooled time dummy (Huber)	+5.5	+4.6	-0.9
Region 07	SHARP-HPI	+5.6	+5.4	-0.2
Region 16	Median price	+6.1	+19.7	+13.6
Region 16	Pooled time dummy (Huber)	+6.6	+6.1	-0.5
Region 16	SHARP-HPI	+6.9	+6.7	-0.1




Figure 10. Composition shock (houses, four regions): shift of the 2023-to-2024 change after removing the smallest 40% of houses from January 2024.

9.4. Placebo

When sale months are randomly permuted within each municipality (keeping composition, location and the municipality's sales calendar but destroying the time structure), the estimated municipal deviations should be flat. Table 23 shows that they are: the median standard deviation of the placebo deviation path is below 1.6 log points at the municipal level and below 0.3 at the MRC level, and the share of placebo 12-month changes outside their 95% band is 0.7% for houses and 0.4% for condominiums at the municipal level, below the nominal 5%. At the regional level for houses the rate is 11.8%, above nominal: with 50 or more sales a month the permuted means are not independent across months because the permutation preserves within-municipality volumes, and the residual seasonal composition survives. Figure 11 contrasts the placebo and actual distributions of 12-month municipal changes.

Table 23. Placebo: sale months permuted within municipality (three replications); estimated deviations should be flat

Class / level	Cells	s.d. of deviation ($\times 100$)	Median max 12-m change ($\times 100$)	% of 12-m changes outside 95	
condo / mrc	240	0.22	0.32	0.5	0.3
condo / municipality	735	0.17	0.26	0.4	0.4
condo / region	48	0.05	0.07	0.0	0.1
house / mrc	306	0.25	0.35	1.6	0.2
house / municipality	3327	1.58	2.39	0.7	1.4
house / region	51	0.84	2.70	11.8	0.7
plex / mrc	303	0.28	0.33	0.2	0.3
plex / municipality	2901	0.95	1.32	0.1	1.6
plex / region	51	0.09	0.15	0.3	0.1




Figure 11. Placebo (months permuted within municipality) versus actual absolute 12-month changes of municipal house indices.

9.5. Residual spatial dependence

Moran's I of the stage-one residuals on ten nearest neighbours, computed in June of each year province-wide and within the three largest cities (Table 24, Figure 12), is small: within cities it ranges between -0.02 and 0.04 and is significant at the 5% level in a minority of the months tested. The 500-m to 1-km location cells absorb most of the spatial structure; what remains is a second-order concern for the index, which averages residuals over cells, but it means that model-based intervals understate the variance in cells where the sold dwellings cluster.

Table 24. Residual spatial autocorrelation of stage-1 residuals (June of each year)

Class / scope	Months tested	Mean Moran's I (10-NN)	Mean z	% months significant
condo / Montréal	6	0.024	2.56	66.7
condo / Québec	6	0.044	2.85	66.7
condo / province	6	0.068	9.64	100.0
house / Gatineau	6	0.019	1.18	16.7
house / Montréal	6	-0.002	-0.28	0.0
house / Québec	6	0.003	0.29	0.0
house / province	6	0.017	5.27	83.3
plex / Montréal	6	0.005	0.33	0.0
plex / Québec	2	-0.002	-0.23	0.0
plex / province	6	0.025	3.05	83.3




Figure 12. Moran's I (10 nearest neighbours) of stage-one house residuals, June of each year.

9.6. Bootstrap versus model-based intervals

A cluster bootstrap that resamples fine-location cells with replacement and re-estimates both stages (40 draws, condominiums) gives the comparison in Table 25. For the province the bootstrap standard error of the log level is about half the model-based one: the model-based μ standard error from the weighted regression does not exploit the within-cell dependence that the bootstrap preserves, and is

conservative. For municipalities with 200 to 600 sales the bootstrap is 1.34 times the model-based standard error (median over the cells tested): the stage-two variances treat the monthly residual means as independent observations with known variance, and the bootstrap adds the variance of the hedonic surface and of the location effects. For the thinnest cells the ratio falls below one, because resampling cells rather than sales leaves a cell with three sales almost unchanged. We report the model-based intervals in the panel and this calibration table alongside them; a user who wants conservative intervals for mid-sized municipalities should inflate them by the tabulated ratio.

Table 25. Cluster bootstrap (40 draws, resampling fine-location cells) versus model-based standard errors, condominiums ($\times 100$)

Cell (condominiums)	Sales	Model s.e. level	Bootstrap s.e. level	Ratio	Model s.e. 12-m	Bootstrap s.e. 12-m	Ratio
province QC	122,989	0.87	0.50	0.58	1.21	0.61	0.50
municipality 78032	61	7.53	11.67	1.55	10.65	14.93	1.40
municipality 71060	204	2.48	3.91	1.58	3.50	4.71	1.34
municipality 61025	199	2.65	3.71	1.40	3.73	4.55	1.22
municipality 60037	17	4.07	3.19	0.78	5.70	3.42	0.60
municipality 68010	3	10.60	6.18	0.58	14.41	5.61	0.39
municipality 63023	1	14.84	12.31	0.83	20.75	10.71	0.52
municipality 77043	199	3.70	6.96	1.88	5.16	9.00	1.74
municipality 59015	348	2.18	5.74	2.63	3.07	7.77	2.53
municipality 45072	621	2.83	4.33	1.53	3.96	5.58	1.41
municipality 78010	12	10.31	9.31	0.90	15.87	10.02	0.63
municipality 32043	5	8.40	6.52	0.78	12.29	6.77	0.55
municipality 53065	3	9.24	4.51	0.49	12.73	3.66	0.29
region 02	754	1.76	2.81	1.60	2.49	3.37	1.36
region 07	4,851	1.38	1.76	1.28	1.92	2.01	1.05
region 04	1,187	2.07	3.44	1.67	2.89	4.33	1.50
region 06	46,957	0.98	0.91	0.93	1.37	1.12	0.81
region 17	608	2.18	3.53	1.62	3.07	4.25	1.38

9.7. Ablations on the Quebec data

Table 26 reports, for each class and level, how far each variant’s municipal monthly changes depart from the full model and the out-of-time calibration of its one-step-ahead innovations. Removing the empirical-Bayes calibration (per-cell maximum likelihood τ^2) changes house municipal monthly changes by 0.71 log points on average and lowers the mean local-data weight, because many thin cells end at the lower boundary of τ^2 and are frozen to their parent. The Gaussian variant and the municipality-level location variant change the series by 3.16 and 3.49 log points, the largest departures: outlying months and within-municipality location shifts both matter. Removing the MRC level changes municipal series by 3.66 log points, which measures how much the administrative hierarchy itself contributes. The out-of-time innovation diagnostics are above one for every variant (1.38 for the full model at the municipal level): the model’s one-step predictive variances are too small by 20–40% for houses, consistent with the bootstrap, and no variant closes the gap; the local-linear-trend variant is slightly worse than the local level.

Table 26. Real-data ablations at the municipality level: RMSE of monthly changes against the full model ($\times 100$) and one-step-ahead standardized innovation diagnostics

Variant	H: Δ vs full	H: OOS z	H:						
Full model	0.00	1.38	86	0.00	1.22	90	0.00	1.33	87
No EB pooling of τ^2	0.71	1.42	85	0.35	1.22	90	0.96	1.33	87
Fixed quality prices	0.48	1.38	86	0.44	1.21	90	0.42	1.32	87
Local linear trend	1.04	1.42	85	0.88	1.19	90	1.24	1.32	87
Gaussian innovations	3.16	1.27	88	1.74	1.08	92	1.92	1.24	88
Municipality-level location FE	3.49	1.35	87	1.55	1.11	91	1.64	1.27	88
No MRC level	3.66	1.34	87	1.47	1.21	90	1.71	1.31	87

OOS z = root mean square of the standardized one-step-ahead prediction error of the cell's monthly mean deviation (1 if calibrated); % in band = share of innovations inside ± 1.96 .

9.8. Sample screens

Table 27 shows that the provincial time-dummy path is insensitive to the sample screens: keeping the assessment-ratio outliers, dropping type conflicts or new constructions, applying a monthly 1–99% price trim, or keeping only roll-classified sales changes the 2021–2026 house path by at most one log point at any month, and the Huber and ordinary least-squares versions differ by less than two log points in total. The robustness comes from the robust likelihood and the location controls, not from the screens.

Table 27. Sensitivity of the provincial time-dummy path to the sample screens (log $\times 100$)

Class	Sample	Sales	Total change (Huber)	Total change (OLS)	RMSE monthly vs headline	Ma
house	Headline sample	422,658	56.6	58.4	0.00	
house	Assessment-ratio outliers kept	437,144	57.2	59.6	0.27	
house	Type conflicts dropped	420,995	56.5	58.5	0.04	
house	New construction dropped	409,697	56.0	56.8	0.17	
house	Monthly 1–99					
house	Roll-classified sales only	419,217	56.5	58.2	0.07	
condo	Headline sample	122,989	43.9	44.1	0.00	
condo	Assessment-ratio outliers kept	124,678	43.6	43.7	0.12	
condo	Type conflicts dropped	122,037	43.4	43.7	0.05	
condo	New construction dropped	122,450	43.8	43.8	0.09	
condo	Monthly 1–99					
condo	Roll-classified sales only	122,982	43.8	44.0	0.01	

9.9. Real-time vintages and revisions

Releases were simulated monthly from July 2022 to October 2026 for the province, the regions and the 64 municipalities with at least 1,500 sales, with the ten-day publication delay measured in Section 3; at each release both stages are re-estimated on the data available then. Two publication

modes are compared. Under *full re-estimation* the published value of a month is the smoothed estimate given the data available at the release, so history is revised; under the *frozen first release* the value of a month is its one-sided filtered estimate and never changes except through stage-one re-estimation, which is the state-space analogue of a movement splice. Table 28 reports, for full re-estimation, the mean absolute revision of the monthly log change between the k -th release and the final retrospective series; Table 29 the same for the frozen policy; Table 30 the provincial policies; Figure 13 the profile and the fan of provincial vintages.

At the provincial level the first-release monthly change is within 0.16 log points of the final value for houses (0.16 for condominiums, 0.26 for plexes), about a tenth of the typical monthly change (1.50), and the third release is no closer: the provincial revision comes entirely from stage-one re-estimation with more data, not from the latest month, so a frozen policy loses nothing at the province. At the regional and municipal levels the two modes differ. For houses the first-release municipal monthly change is revised by 1.72 log points under full re-estimation and the gap to the final series under the frozen policy is 2.31; for condominiums 0.86 against 1.82. The revision declines over the first six releases as the smoother incorporates later months and then stabilises at the stage-one component; twelve-month changes are revised by 2.09 log points at first release for houses. The frozen policy therefore costs about half a log point per month at the municipal level relative to publishing revised values, the price of never revising, and the revision-limited policy (freeze at the third release) captures most of the gain of full re-estimation. These magnitudes are the timeliness–accuracy trade-off that a publisher of municipal indices faces; they are measured on this data and would differ with a slower registry.

Table 28. Real-time vintages: mean absolute revision of the monthly log change ($\times 100$) between the k -th release and the final retrospective estimate

Class / level	Rel. 1	Rel. 2	Rel. 3	Rel. 6	Rel. 12	12-m change, rel. 1	% rev > 1 pt
house / province	0.16	0.17	0.16	0.15	0.15	0.13	0
house / region	1.62	1.56	1.57	1.51	1.59	2.16	44
house / municipality	1.72	1.61	1.59	1.55	1.62	2.09	54
condo / province	0.16	0.16	0.16	0.17	0.16	0.25	0
condo / region	0.48	0.43	0.43	0.43	0.39	1.23	12
condo / municipality	0.86	0.78	0.74	0.73	0.66	2.40	27
plex / province	0.26	0.25	0.26	0.27	0.24	0.35	0
plex / region	0.76	0.64	0.59	0.57	0.61	1.55	26
plex / municipality	1.58	1.42	1.35	1.29	1.27	3.95	48

Releases simulated monthly from July 2022 to October 2026 with a 10-day publication delay; municipalities with at least 1,500 sales. Final = smoothed full-sample panel.

Table 29. Frozen-first-release (one-sided filtered) policy: mean absolute gap of the monthly log change ($\times 100$) between the real-time value and the final smoothed estimate

Class / level	Rel. 1 vs final	Rel. 12 vs final	12-m change, rel. 1	Rel. 1 vs last real-time
house / province	0.16	0.15	0.13	0.16
house / region	1.93	1.94	2.24	1.78
house / municipality	2.31	2.43	2.42	1.61
condo / province	0.16	0.16	0.25	0.16
condo / region	0.98	0.98	1.93	0.61
condo / municipality	1.82	1.87	3.59	1.07
plex / province	0.26	0.24	0.35	0.26
plex / region	1.37	1.41	2.05	0.86
plex / municipality	3.03	3.00	5.26	1.99

Under this policy a month's value changes only through stage-1 re-estimation; the gap to the final series is the price of never revising (filter versus smoother), and it does not decline with later releases.

Table 30. Provincial monthly change ($\times 100$): first and third release versus final

Class	Months	MAE first release	Bias first release	MAE third release	RMSE first	RMSE third	Mean $ \Delta $
condo	52	0.16	+0.013	0.15	0.21	0.20	0.77
house	52	0.16	-0.013	0.16	0.20	0.20	1.50
plex	52	0.26	-0.014	0.26	0.35	0.35	1.53

Full re-estimation policy (smoothed at each vintage). A revision-limited policy freezes the value at the third release; the frozen-first-release (movement-splice) policy is in Table 29.




Figure 13. Left: mean absolute revision of the monthly change by release number, houses (solid: full re-estimation; dashed: frozen first release). Right: every monthly vintage of the provincial house index, July 2022 to October 2026 (the darkest line is the latest).

9.10. External comparison

Table 31 compares the provincial and metropolitan SHARP series with the Teranet–National Bank and CREA indices, rebased to the 2021 average and compared on annual means to avoid the thin January 2021. In the Québec City metropolitan area the agreement is close: the SHARP house aggregate rises by 47.9 log points from 2021 to 2026 against 45.5 for the CREA single-family benchmark and 43.7 for Teranet (all types), and the condominium aggregate by 43.9 against 47.1 for CREA apartments. In Montréal the SHARP house aggregate rises by 35.6 against 25.0 for CREA and 26.1 for Teranet, a gap of about ten log points, while condominiums are within five points of CREA apartments. Table 32 shows that the gap is in the data rather than in the estimator: repeat sales on our own Montréal house sample rise by 36.1 log points, the same as the SHARP aggregate, whereas repeat sales on all three classes together rise by 30.1, closer to Teranet’s all-type 26.1. Two explanations remain for the house gap relative to CREA: the CREA benchmark uses bedroom, bathroom and other listing attributes that absorb quality upgrades our roll variables cannot, and brokered resales are a different population from registry deeds. The correlation of 12-month changes between SHARP and the external series is between 0.95 and 0.95 in Montréal; in Québec City it is lower because the SHARP metropolitan aggregate is noisier month to month than the smoothed external series.

Table 31. Comparison with external indices: change of the annual mean log index from 2021 to 2026 (January–September), log $\times 100$, and correlations of 12-month and 1-month log changes

SHARP series	External series	Geography	SHARP 2026 vs 2021	External 2026 vs 2021	Gap	Corr. 12-m	Corr. 1-m
condo (CMM/CMQ)	CREA MLS HPI, Composite NSA	Montréal CMA	28.4	23.5	+4.9	0.88	0.54
condo (CMM/CMQ)	CREA MLS HPI, Apartment NSA	Montréal CMA	28.4	23.7	+4.7	0.92	0.52
condo (CMM/CMQ)	CREA MLS HPI, Single Family NSA	Montréal CMA	28.4	25.0	+3.4	0.86	0.48
condo (CMM/CMQ)	Teranet-NB, NSA (repeat sales)	Montréal CMA	28.4	26.1	+2.3	0.95	0.63
house (CMM/CMQ)	CREA MLS HPI, Single Family NSA	Montréal CMA	35.6	25.0	+10.6	0.91	0.33
house (CMM/CMQ)	CREA MLS HPI, Composite NSA	Montréal CMA	35.6	23.5	+12.2	0.92	0.33
house (CMM/CMQ)	Teranet-NB, NSA (repeat sales)	Montréal CMA	35.6	26.1	+9.6	0.95	0.70
house (CMM/CMQ)	CREA MLS HPI, Apartment NSA	Montréal CMA	35.6	23.7	+12.0	0.90	0.24
condo province	CREA MLS HPI, Apartment NSA	Quebec (province)	31.6	24.7	+6.9	0.88	0.58
condo province	CREA MLS HPI, Composite NSA	Quebec (province)	31.6	26.5	+5.2	0.84	0.58
condo province	CREA MLS HPI, Single Family NSA	Quebec (province)	31.6	28.1	+3.5	0.83	0.52
house province	CREA MLS HPI, Composite NSA	Quebec (province)	42.3	26.5	+15.9	0.91	0.37
house province	CREA MLS HPI, Single Family NSA	Quebec (province)	42.3	28.1	+14.3	0.91	0.36
house province	CREA MLS HPI, Apartment NSA	Quebec (province)	42.3	24.7	+17.6	0.92	0.30
condo (CMM/CMQ)	Teranet-NB, NSA (repeat sales)	Québec CMA	43.9	43.7	+0.2	0.71	0.13
condo (CMM/CMQ)	CREA MLS HPI, Single Family NSA	Québec CMA	43.9	45.5	-1.6	0.60	0.06
condo (CMM/CMQ)	CREA MLS HPI, Composite NSA	Québec CMA	43.9	44.3	-0.4	0.67	0.11
condo (CMM/CMQ)	CREA MLS HPI, Apartment NSA	Québec CMA	43.9	47.1	-3.1	0.67	0.19
house (CMM/CMQ)	CREA MLS HPI, Apartment NSA	Québec CMA	47.9	47.1	+0.9	0.54	-0.01
house (CMM/CMQ)	CREA MLS HPI, Composite NSA	Québec CMA	47.9	44.3	+3.6	0.59	-0.08
house (CMM/CMQ)	CREA MLS HPI, Single Family NSA	Québec CMA	47.9	45.5	+2.4	0.56	-0.09
house (CMM/CMQ)	Teranet-NB, NSA (repeat sales)	Québec CMA	47.9	43.7	+4.3	0.63	0.54

Teranet–National Bank: repeat sales on all registry transfers, CMA, all residential types. CREA MLS HPI: benchmark hedonic on b resales, CMA / provincial aggregate. SHARP metropolitan series = stock-weighted aggregates of municipal cells inside the Montréal and Québec (CMQ) metropolitan communities; all series rebased to 2021 average = 100 before comparison.

Table 32. Repeat-sales and median-price changes computed on the Vrai Prix data for the geographies of the external indices (annual means, log $\times 100$)

Sample (this data)	Pairs	Repeat sales 2026 vs 2021	Median price 2026 vs 2021
Montréal (CMM) houses	11,722	36.1	31.9
Montréal (CMM) condos	11,605	22.1	21.6
Montréal (CMM) all three classes	26,579	30.1	31.0
Island of Montréal (region 06) all	8,978	19.7	23.5
Québec (CMQ) all three classes	8,794	46.8	45.2
Province houses	45,354	47.7	41.9
Province all three classes	73,384	43.3	38.7

Case–Shiller weighted repeat sales, holding period at least six months. Teranet Montréal CMA (all types) rises by 26.1 and Québec CMA by 43.7 on the same basis.

10. Economic implications and applications

Three uses follow from a monthly municipal panel with reliability accounting. First, assessment: Quebec’s triennial rolls are indexed to a reference date, and the ratio of sale prices to assessed values drifts between rolls; a municipal index with an interval gives assessors and the provincial agency a defensible time-adjustment factor for sales comparison, and the grade tells them when the factor should be the MRC’s or the region’s rather than the municipality’s. Second, monitoring: the dispersion documented in Section 8 means that provincial aggregates hide most of the variation that matters for local housing policy; the panel makes the variation visible while labelling the cells where it is not identified. Third, research: a constant-quality index at the municipal level with a measured revision profile is the input that studies of local demand shocks, migration, short-term rentals or zoning changes need, and the real-time vintages allow such studies to use the information that was available at the time.

Two cautions accompany these uses. The index measures the price of a constant basket in the observable attributes; renovations and unobserved quality improvements are in the index, not out of it, and the comparison with the CREA benchmark index in Section 9 suggests that this matters for houses in Montréal. And the intervals are model-based and, for mid-sized municipalities, narrower than a cluster bootstrap by a factor of about 1.34 (Table 25); users who need conservative inference should widen them accordingly.

11. Limitations and conclusion

11.1. Limitations

The data limit what any estimator can do. Characteristics come from a single current roll vintage, so renovations after the sale are invisible and recent new constructions may be excluded; the

provider collects no sale below \$50,000; the roll's neighbourhood code is unusable, and location is a metric grid rather than an economic neighbourhood; arm's-length status, financing and listing information are absent; and the single date per sale conflates signing and registration. The model is sequential: stage two conditions on stage-one residuals and the intervals omit hedonic-parameter and hyper-parameter uncertainty, which the bootstrap shows to matter for mid-sized municipalities. The hierarchy is administrative; a municipality is shrunk toward its MRC whether or not the MRC is its housing market, and the ablation that removes the MRC level shows that this choice changes municipal monthly changes by a few log points. Seasonality is not modelled. The Monte Carlo design, although rich, is ours, and the two counter-examples were chosen by us. Statistics Canada boundaries could not be obtained, so external comparisons use metropolitan communities.

11.2. Conclusion

We set out to measure constant-quality house price movements where the data are thin, and to say how far they can be trusted. The framework integrates known components; the evidence shows that their combination produces municipal monthly indices with errors a fraction of those of medians and cell-level time dummies in simulation, intervals whose coverage is close to nominal when the innovation variance is calibrated across cells, a revision profile that is measured rather than designed away, and aggregates that agree with external indices where the scopes match. It also shows where the approach should not be used: at the regional level, where sales are plentiful, region-specific hedonic coefficients are preferable to pooled coefficients plus a deviation; with heavy-tailed contamination a coarser location partition is more robust than fine cells of fifteen sales; in markets with strong idiosyncratic local dynamics the real-time estimate lags; and for cells with fewer than fifteen effective observations no estimator reports more than a model-assisted interpolation. The panel, the vintage database and the code are released so that both claims can be checked.

References

- Anenberg, Elliot and Steven Laufer**, “A More Timely House Price Index,” *The Review of Economics and Statistics*, 2017, 99 (4), 722–734.
- Athanasopoulos, George, Rob J. Hyndman, Nikolaos Kourentzes, and Anastasios Panagiotelis**, “Forecast reconciliation: A review,” *International Journal of Forecasting*, 2024, 40 (2), 430–456.
- Bailey, Martin J., Richard F. Muth, and Hugh O. Nourse**, “A Regression Method for Real Estate Price Index Construction,” *Journal of the American Statistical Association*, 1963, 58 (304), 933–942.
- Bogin, Alexander, William Doerner, and William Larson**, “Local House Price Dynamics: New Indices and Stylized Facts,” *Real Estate Economics*, 2019, 47 (2), 365–398.
- Bourassa, Steven C., Eva Cantoni, and Martin Hoesli**, “Robust Repeat Sales Indexes,” *Real Estate Economics*, 2013, 41 (3), 517–541.
- , —, and —, “Robust hedonic price indexes,” *International Journal of Housing Markets and Analysis*, 2016, 9 (1), 47–65.
- Canadian Real Estate Association**, “MLS® Home Price Index (HPI) Methodology,” https://www.crea.ca/files/HPI_Methodology-July-2023-rev-ENG.pdf 2023. July 2023 revision.
- Case, Bradford and John M. Quigley**, “The Dynamics of Real Estate Prices,” *The Review of Economics and Statistics*, 1991, 73 (1), 50–58.
- Case, Karl E. and Robert J. Shiller**, “Prices of Single-Family Homes Since 1970: New Indexes for Four Cities,” *New England Economic Review*, September 1987, pp. 45–56.
- and —, “The Efficiency of the Market for Single-Family Homes,” *American Economic Review*, 1989, 79 (1), 125–137.
- Clapham, Eric, Peter Englund, John M. Quigley, and Christian L. Redfearn**, “Revisiting the Past and Settling the Score: Index Revision for House Price Derivatives,” *Real Estate Economics*, 2006, 34 (2), 275–302.
- Clapp, John M. and Carmelo Giaccotto**, “Revisions in Repeat-Sales Price Indexes: Here Today, Gone Tomorrow?,” *Real Estate Economics*, 1999, 27 (1), 79–104.
- , —, and **Dogan Tirtiroglu**, “Housing Price Indices Based on All Transactions Compared to Repeat Subsamples,” *Real Estate Economics*, 1991, 19 (3), 270–285.

- Contat, Justin and William D. Larson**, “A flexible method of housing price index construction using repeat-sales aggregates,” *Real Estate Economics*, 2024, 52 (6), 1551–1583.
- Court, Andrew T.**, “Hedonic Price Indexes with Automotive Examples,” in “The Dynamics of Automobile Demand,” General Motors Corporation, 1939, pp. 99–117.
- Croushore, Dean**, “Frontiers of Real-Time Data Analysis,” *Journal of Economic Literature*, 2011, 49 (1), 72–100.
- Datta, Gauri Sankar and Partha Lahiri**, “A Unified Measure of Uncertainty of Estimated Best Linear Unbiased Predictors in Small Area Estimation Problems,” *Statistica Sinica*, 2000, 10 (2), 613–627.
- , **J. N. K. Rao, and David Daniel Smith**, “On measuring the variability of small area estimators under a basic area level model,” *Biometrika*, 2005, 92 (1), 183–196.
- de Haan, Jan**, “Hedonic Price Indexes: A Comparison of Imputation, Time Dummy and ‘Re-Pricing’ Methods,” *Jahrbücher für Nationalökonomie und Statistik*, 2010, 230 (6), 772–791.
- **and Frances Krsinich**, “Scanner Data and the Treatment of Quality Change in Nonrevisable Price Indexes,” *Journal of Business & Economic Statistics*, 2014, 32 (3), 341–358.
- Diewert, W. Erwin and Kevin J. Fox**, “Substitution Bias in Multilateral Methods for CPI Construction,” *Journal of Business & Economic Statistics*, 2022, 40 (1), 355–369.
- , **Saeed Heravi, and Mick Silver**, “Hedonic Imputation versus Time Dummy Hedonic Indexes,” in “Price Index Concepts and Measurement,” University of Chicago Press, 2009, pp. 161–196.
- Dorsey, Robert E., Haixin Hu, Walter J. Mayer, and Hui chen Wang**, “Hedonic versus repeat-sales housing price indexes for measuring the recent boom-bust cycle,” *Journal of Housing Economics*, 2010, 19 (2), 75–93.
- Durbin, James and Siem Jan Koopman**, *Time Series Analysis by State Space Methods*, 2 ed., Oxford University Press, 2012.
- Efron, Bradley and Carl Morris**, “Stein’s Estimation Rule and its Competitors—An Empirical Bayes Approach,” *Journal of the American Statistical Association*, 1973, 68 (341), 117–130.
- **and —**, “Data Analysis Using Stein’s Estimator and its Generalizations,” *Journal of the American Statistical Association*, 1975, 70 (350), 311–319.
- Eurostat, ILO, IMF, OECD, UNECE, and World Bank**, *Handbook on Residential Property Price Indices (RPPIs)*, OECD Publishing / Eurostat, 2013.

- Fay, Robert E. and Roger A. Herriot**, “Estimates of Income for Small Places: An Application of James-Stein Procedures to Census Data,” *Journal of the American Statistical Association*, 1979, 74 (366), 269–277.
- Francke, Marc K.**, “Repeat Sales Index for Thin Markets,” *The Journal of Real Estate Finance and Economics*, 2010, 41 (1), 24–52.
- **and Alex van de Minne**, “The Hierarchical Repeat Sales Model for Granular Commercial Real Estate and Residential Price Indices,” *The Journal of Real Estate Finance and Economics*, 2017, 55 (4), 511–532.
- Francke, Marc, Lyndsey Rolheiser, and Alex van de Minne**, “Estimating Census Tract House Price Indexes: A New Spatial Dynamic Factor Approach,” *The Journal of Real Estate Finance and Economics*, 2025, 70 (3), 483–514.
- Gatzlaff, Dean H. and Donald R. Haurin**, “Sample Selection Bias and Repeat-Sales Index Estimates,” *The Journal of Real Estate Finance and Economics*, 1997, 14 (1-2), 33–50.
- **and —**, “Sample Selection and Biases in Local House Value Indices,” *Journal of Urban Economics*, 1998, 43 (2), 199–222.
- Gelman, Andrew**, “Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper),” *Bayesian Analysis*, 2006, 1 (3), 515–534.
- Griliches, Zvi**, “Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change,” in “The Price Statistics of the Federal Government,” National Bureau of Economic Research, 1961, pp. 173–196. Staff Paper 3; reprinted in *Price Indexes and Quality Change*, Harvard University Press, 1971.
- Harvey, Andrew C.**, *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press, 1989.
- Hill, Robert J.**, “Hedonic Price Indexes for Residential Housing: A Survey, Evaluation and Taxonomy,” *Journal of Economic Surveys*, 2013, 27 (5), 879–914.
- **and Daniel Melser**, “Hedonic Imputation and the Price Index Problem: An Application to Housing,” *Economic Inquiry*, 2008, 46 (4), 593–609.
- **and Michael Scholz**, “Can Geospatial Data Improve House Price Indexes? A Hedonic Imputation Approach with Splines,” *Review of Income and Wealth*, 2018, 64 (4), 737–756.

- **and Miriam Steurer**, “Real Estate Price Measurement Since the Global Financial Crisis: Methods, Data and Policy Developments,” *Review of Income and Wealth*, 2026, 72 (1). Article e70055.
- **, Michael Scholz, Chihiro Shimizu, and Miriam Steurer**, “Rolling-Time-Dummy House Price Indexes: Window Length, Linking and Options for Dealing with Low Transaction Volume,” *Journal of Official Statistics*, 2022, 38 (1), 127–151.
- Hwang, Min and John M. Quigley**, “Selectivity, Quality Adjustment and Mean Reversion in the Measurement of House Values,” *The Journal of Real Estate Finance and Economics*, 2004, 28 (2-3), 161–178.
- Hyndman, Rob J., Roman A. Ahmed, George Athanasopoulos, and Han Lin Shang**, “Optimal combination forecasts for hierarchical time series,” *Computational Statistics & Data Analysis*, 2011, 55 (9), 2579–2589.
- Ishaak, Farley, Pim Ouwehand, and Hilde Remøy**, “Constructing Limited-Revisable and Stable CPPIs for Small Domains,” *Journal of Official Statistics*, 2024, 40 (3), 380–408.
- Ivancic, Lorraine, W. Erwin Diewert, and Kevin J. Fox**, “Scanner data, time aggregation and the construction of price indexes,” *Journal of Econometrics*, 2011, 161 (1), 24–35.
- Kalman, R. E.**, “A New Approach to Linear Filtering and Prediction Problems,” *Journal of Basic Engineering*, 1960, 82 (1), 35–45.
- Lopez, Esteban and Geoffrey J.D. Hewings**, “Housing price indices for small spatial units,” *Regional Science and Urban Economics*, 2018, 70, 57–71.
- Nowak, Adam D. and Patrick S. Smith**, “Quality-Adjusted House Price Indexes,” *American Economic Review: Insights*, 2020, 2 (3), 339–356.
- Pfeffermann, Danny**, “New Important Developments in Small Area Estimation,” *Statistical Science*, 2013, 28 (1), 40–68.
- Quigley, John M.**, “A Simple Hybrid Model for Estimating Real Estate Price Indexes,” *Journal of Housing Economics*, 1995, 4 (1), 1–12.
- Rao, J.N.K. and Isabel Molina**, *Small Area Estimation*, 2 ed., John Wiley & Sons, 2015.
- Ren, You, Emily B. Fox, and Andrew Bruce**, “Clustering correlated, sparse data streams to estimate a localized housing price index,” *The Annals of Applied Statistics*, 2017, 11 (2), 808–839.

- Rosen, Sherwin**, “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition,” *Journal of Political Economy*, 1974, 82 (1), 34–55.
- Sayag, Doron, Dano Ben-hur, and Danny Pfeffermann**, “Reducing revisions in hedonic house price indices by the use of nowcasts,” *International Journal of Forecasting*, 2022, 38 (1), 253–266.
- Schwann, Gregory M.**, “A Real Estate Price Index for Thin Markets,” *The Journal of Real Estate Finance and Economics*, 1998, 16 (3), 269–287.
- Shiller, Robert J.**, “Arithmetic repeat sales price estimators,” *Journal of Housing Economics*, 1991, 1 (1), 110–126.
- Silver, Mick and Saeed Heravi**, “The Difference Between Hedonic Imputation Indexes and Time Dummy Hedonic Indexes,” *Journal of Business & Economic Statistics*, 2007, 25 (2), 239–246.
- Statistics Canada**, “Methodology of the Residential Property Price Index (RPPI),” Statistics Canada, Prices Analytical Series, Catalogue no. 62F0014M2019006 2019. Released 14 November 2019; archived.
- Teranet and National Bank of Canada**, “Teranet–National Bank House Price Index: Methodology Overview,” <https://housepriceindex.ca/wp-content/uploads/2017/08/Teranet-National-Bank-House-Price-Index-Methodology-Overview.pdf> 2017. Accessed 21 October 2026.
- van de Minne, Alex, Marc Francke, David Geltner, and Robert White**, “Using Revisions as a Measure of Price Index Quality in Repeat-Sales Models,” *The Journal of Real Estate Finance and Economics*, 2020, 60 (4), 514–553.
- Wallace, Nancy E. and Richard A. Meese**, “The Construction of Residential Housing Price Indices: A Comparison of Repeat-Sales, Hedonic-Regression, and Hybrid Approaches,” *The Journal of Real Estate Finance and Economics*, 1997, 14 (1-2), 51–73.
- Wickramasuriya, Shanika L., George Athanasopoulos, and Rob J. Hyndman**, “Optimal Forecast Reconciliation for Hierarchical and Grouped Time Series Through Trace Minimization,” *Journal of the American Statistical Association*, 2019, 114 (526), 804–819.

Appendices

A. Algorithms

Stage one. For class k : (1) build features (conditional-median imputation of floor area, lot area and age within class $\times$ municipality, then class $\times$ region, then class; missingness indicators); (2) assign the location partition: the finest cell among 500 m, 1 km, 2 km, 5 km, municipality and MRC with at least 15 sales of the class; (3) fit $\log P$ on the hedonic terms, their time interactions and month dummies with absorbed location effects by weighted least squares; (4) compute Huber weights $w_i = \min(1, 1.345s/|e_i|)$ with s the MAD-based scale and refit twice; (5) return $\mu_{k,t}$, its standard errors, the residuals and the robust scale.

Stage two. For each level in (region, MRC, municipality): (1) subtract the parent’s deviation path from each residual; (2) aggregate residual means $\bar{y}_{g,t}$ and counts $n_{g,t}$; (3) set $\sigma_g^2 = w s_g^2 + (1 - w) \sigma_k^2$ with $w = N_g/(N_g + 100)$ and s_g the cell’s MAD-based scale; (4) evaluate the Kalman-filter log-likelihood of each cell with at least six observed months on the grid $\log \tau^2 \in [\log 10^{-8}, \log 2 \cdot 10^{-2}]$ (48 points); (5) run the grid EM for (m_ℓ, v_ℓ) with $v_\ell \geq 0.05$; (6) for each cell choose τ_g^2 at the mode of $\log L_g(\tau^2) - (\log \tau^2 - m_\ell)^2/(2v_\ell)$, run the filter, down-weight months with standardised innovations above 1.345 by dividing their observation variance by the Huber weight, re-choose τ_g^2 and re-run (two passes), and apply the Rauch–Tung–Striebel smoother; (7) pass the smoothed (retrospective) or filtered (real-time) deviation to the children.

Kalman recursions. With state δ_t , $\delta_{t|t-1} = \delta_{t-1|t-1}$, $p_{t|t-1} = p_{t-1|t-1} + \tau^2$; if $n_t > 0$: $F_t = p_{t|t-1} + \sigma^2/n_t$, $v_t = \bar{y}_t - \delta_{t|t-1}$, $\kappa_t = p_{t|t-1}/F_t$, $\delta_{t|t} = \delta_{t|t-1} + \kappa_t v_t$, $p_{t|t} = p_{t|t-1}(1 - \kappa_t)$, $\log L += -\frac{1}{2}(\log 2\pi F_t + v_t^2/F_t)$; otherwise $\delta_{t|t} = \delta_{t|t-1}$, $p_{t|t} = p_{t|t-1}$, $\kappa_t = 0$. The initial state is diffuse ($p_0 = 0.25$). The smoother runs backward with gain $c_t = p_{t|t}/p_{t+1|t}$.

B. Data cleaning rules

The rules of Table 2 are implemented in `src/sharp_hpi/cleaning.py`. Classification precedence: (R1) roll building year after the sale year; (R2) roll type “terrain” or use code beginning with 9 or 4; (R3) use code outside {1000, 1100, 1211}; (R4) class from the roll (1000 with 1 dwelling and “Intégr  ” link: condominium; 1000 with 1 dwelling and detached/semi/row link: house; 1000 with 2–5 dwellings: plex; 6 or more: multi-unit; 1100: cottage; 1211: mobile home); (R5) provider label when the roll is incomplete. Flags: exact duplicate (parcel, date, price), same parcel same day different prices, shared parcel (> 10 sales), multi-parcel deed (same date,

price, municipality and street name, or same date, price and 0.01° cell), assessment-ratio outlier ($|\log(P/A) - \text{median}_{\text{year,class,region}}| > \log 4$, with A the assessment in force at the sale), price above \$10 million. Independent flag counts and the provider-versus-roll agreement are in the repository.

C. Geographic concordance

Boundaries: MRNF *Découpages administratifs* 1:20 000, version V2026-07, layers `munic_s` (municipality, MRC and region codes), `comet_s` (metropolitan communities), `arron_s` (boroughs); CRS NAD83 geographic, metric operations in NAD83 / Quebec Lambert (EPSG:32198). Every sale falls inside a municipal polygon. The MAMH municipal register supplies population, designation and constitution dates; municipalities constituted or merged after January 2021 (Amos 2025, Courcelles–Saint-Évariste 2024, Hébertville 2026, La Morandière-Rochebaucourt 2023, La Pocatière 2025, Lac-des-Aigles 2024, Notre-Dame-de-la-Salette 2022, Oujé-Bougoumou 2023, Plessisville 2024) are coded on their 2026 boundaries throughout.

D. Simulation design

Table 7 lists the scenarios. Each replication draws 60 municipalities ($4 \text{ regions} \times 3 \text{ MRCs} \times 5 \text{ municipalities}$) with expected monthly sales $\lambda_u \sim \text{log-normal}(1.2, 1.3^2)$ (median 3.3; the “sparse” scenario multiplies by 0.35), housing stock proportional to λ_u times a $\text{uniform}(80, 160)$ factor, latent paths μ_t (random walk with drift, optional break), δ^R , δ^M , δ^U (random walks with the stated innovation standard deviations; half of the municipalities have $\tau_U/4$), sales $\sim \text{Poisson}(\lambda_u s_t)$ with seasonal factor $s_t = 1 + 0.6 \cos(2\pi(t \bmod 12 - 5)/12)$, log floor area $\sim N(\log 120 + 0.15m_u + qt, 0.35^2)$ with municipality level $m_u \sim N(0, 0.25^2)$ and quality trend q , gamma ages, log lot areas, storeys, a hedonic surface with quadratic floor-area and age terms and an elasticity that drifts by β_{fa} over the sample, spatially correlated cell quality $0.08 \sin(x/1500) \cos(y/1500) + 0.06 \mathbb{I}(\text{radius} < 1.5 \text{ km})$, Student- t noise with the stated degrees of freedom and contamination, resales (12% of sales are of a randomly chosen earlier parcel of the same municipality at least six months earlier), floor area missing at the stated rate, and exponential registration lags with mean 10 days. The truth for a municipality is $L_{u,t}$; for a region, MRC or the province it is the stock-weighted mean of its municipalities’ $L_{u,t}$.

E. Supplementary tables

Table 33. Completeness of characteristics by year of sale, all raw rows (including land, non-residential and unmatched sales)

Variable (% non-missing)	2021	2022	2023	2024	2025	2026
Floor area	79.6	81.6	86.5	88.8	88.2	85.2
Year built	69.8	77.0	86.4	88.9	88.4	85.3
Lot area	99.6	99.7	99.8	99.8	99.8	98.1
Storeys	90.8	90.1	91.2	91.6	92.0	92.9
Physical link	89.9	89.4	90.5	90.9	91.2	92.2
Building genre	86.3	85.6	87.0	87.4	87.8	88.8
Provider style (banned)	12.8	51.4	82.5	85.2	84.8	81.8
Roll identifier	98.8	98.8	99.1	99.4	99.5	99.3

The provider's building style is backfilled over time and is excluded from every model (time-correlated missingness). Lower shares in 2021 reflect the larger share of land and unclassified rows, not missing data in the headline sample (Table 4).