How Much Do Text-Derived Quality Measures Matter?

LLM-based variables and the robustness of residential property price indices

Tom Gillespie
University of Galway

Ronan C. Lyons
Trinity College Dublin, CEPH & Central Bank of Ireland

International Conference on Real Estate Statistics
Tokyo, February 2026

Central research question:

Do LLM-extracted listing variables materially alter measured RPPI growth over time?

Key takeaway:

LLM-based quality measures substantially improve model fit, but have only modest effects on aggregate RPPI growth.

Roadmap

Motivation & Context

Data & Measurement

Empirical Strategy

Results

Implications and Conclusions

Why this matters for RPPIs

- **RPPI objective:** measure *pure price change* for a constant-quality housing bundle.
- Hedonic models control for observable characteristics to isolate time effects.
- But some key quality dimensions — especially internal condition — are typically unobserved.

Core concern

Bias arises only if the mix of unobserved quality changes systematically over time.

If internal condition varies within areas but is stable over time, the RPPI may remain robust even if condition is omitted.

The measurement challenge: cross-section vs time

- Internal condition, refurbishment, amenities, and views are strongly capitalised into prices.
- Cross-sectional effects can be large.

But RPPIs are identified from time variation

- Omitted quality matters for the index only if its distribution shifts over time.
- The key empirical question is therefore dynamic, not static.

This distinction – cross-sectional pricing vs. time-series identification – is central to interpreting LLM-based measurement.

Our approach

- Setting: Ireland, 2014–2025.
- Data: 730,000+ residential listings with structured fields and full text.
- We extract multiple dwelling-level variables from descriptions using a constrained LLM schema:
  - Internal condition (1–5)
  - Site size, garden, views
  - Construction period
- We embed these directly into a standard hedonic RPPI framework.

Design principle

Hold sample, geography, and functional form constant. Compare RPPIs with and without LLM-derived variables.

Positioning

- Hedonic RPPIs increasingly use rich microdata.
- Listings text contains additional structured information not captured in standard datasets.
- LLMs make that information extractable at scale.

Our emphasis

Not “can an LLM predict prices?”
But: **does incorporating LLM-measured quality alter official-style index dynamics?**

This is a robustness question about measurement stability, not a forecasting exercise.

Roadmap

- Motivation & Context
- Data & Measurement
- Empirical Strategy
- Results
- Implications and Conclusions

Data: listings and setting (Ireland, 2014–2025)

- **Source:** Daft.ie residential *sale* listings (Ireland's dominant platform; near-universal market coverage).
- **Sample:** 731,543 listings 2014–2025, after cleaning and merging phrase/LLM data.
- **Outcome:** listed price (logged in regressions); matched to transacted price for subsample
- **Structured fields:**
  - property type (simplified categories),
  - bedrooms / bathrooms (combined categories),
  - floor area (sqm; partial coverage),
  - new development flag,
  - geolocation (micro-market assignment; ~400 markets nested within counties/cities).
- **Key advantage:** full description for nearly all listings, enabling scalable measurement of otherwise unobserved variables

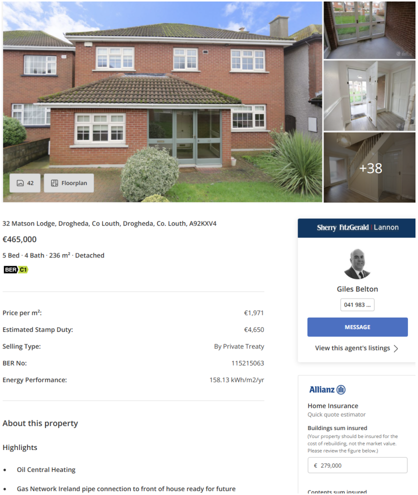

Controls: dwelling attributes, text measures, fixed effects

- **Dwelling attributes (structured):**
  - BER category (energy rating),
  - floor area category (sqm bins),
  - new development indicator,
  - price-change indicator,
  - beds × property-type interactions.
- **Text-derived controls:**
  - Regex phrase indicators (amenities, vintage, views, etc.)
  - LLM-derived variables – as below.
- **Spatio-temporal structure:**
  - Micro-market fixed effects: bespoke set of ~400 real estate markets nested within cities and counties.
  - Month fixed effects (time dummies).

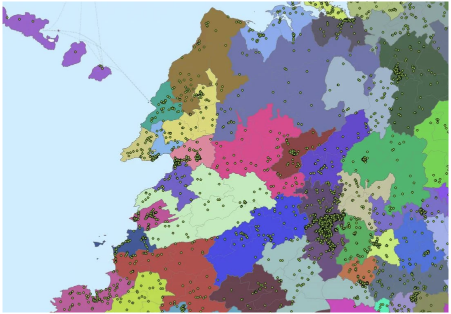

LLM extraction: from descriptions to structured variables (prompted JSON)

- **Input:** full listing description text.
- **Output:** a **single strict JSON object** per listing, following a fixed schema (no free text).
  - Variables are **explicitly defined** and **bounded** (limited categories / units)
  - Missingness is standardised (`0` or `"na"`)
  - Extraction is **auditable** (prompt + schema + examples are fixed)

Extracted fields

1. **Condition category** (more shortly)
2. **Construction year** (4-digit year, or 0)
3. **Site size** and units (ha/ac/sqm/sqft, or 0/"na")
4. **Construction period** (e.g. "Georgian", "Interwar", "Period")
5. **Garden indicators:** `has_garden`, `front_garden`, `back_garden`
6. **Garden size** `garden_size` + `garden_size_units` (sqm/sqft/length units, or 0/"na")
7. **Named style (rare):** `construction_style` (eg. Strain/Stringer)
8. **Views:** `sea_view`, `green_view`

LLM extraction: from descriptions to structured variables (prompted JSON)

- **Input:** full listing description text.
- **Output:** a **single strict JSON object** per listing, following a fixed schema (no free text).
  - Variables are **explicitly defined** and **bounded** (limited categories / units)
  - Missingness is standardised (`0` or **`"na"`**)
  - Extraction is **auditable** (prompt + schema + examples are fixed)

Extracted fields

1. **Condition category** (more shortly)
2. **Construction year** (4-digit year, or 0)
3. **Site size** and units (ha/ac/sqm/sqft, or 0/"na")
4. **Construction period** (e.g. "Georgian", "Interwar", "Period")
5. **Garden indicators:** `has_garden`, `front_garden`, `back_garden`
6. **Garden size** `garden_size` + `garden_size_units` (sqm/sqft/length units, or 0/"na")
7. **Named style (rare):** `construction_style` (eg. Strain/Stringer)
8. **Views:** `sea_view`, `green_view`

LLM extraction: from descriptions to structured variables (prompted JSON)

- **Input:** full listing description text.
- **Output:** a **single strict JSON object** per listing, following a fixed schema (no free text).
  - Variables are **explicitly defined** and **bounded** (limited categories / units)
  - Missingness is standardised (`0` or `"na"`)
  - Extraction is **auditable** (prompt + schema + examples are fixed)

Extracted fields

1. **Condition category** (more shortly)
2. **Construction year** (4-digit year, or 0)
3. **Site size** and units (ha/ac/sqm/sqft, or 0/"na")
4. **Construction period** (e.g. “Georgian”, “Interwar”, “Period”)
5. **Garden indicators:** `has_garden`, `front_garden`, `back_garden`
6. **Garden size** `garden_size` + `garden_size_units` (sqm/sqft/length units, or 0/"na")
7. **Named style (rare):** `construction_style` (eg. Strain/Stringer)
8. **Views:** `sea_view`, `green_view`

LLM extraction: from descriptions to structured variables (prompted JSON)

- **Input:** full listing description text.
- **Output:** a **single strict JSON object** per listing, following a fixed schema (no free text).
  - Variables are **explicitly defined** and **bounded** (limited categories / units)
  - Missingness is standardised (`0` or `"na"`)
  - Extraction is **auditable** (prompt + schema + examples are fixed)

Extracted fields

1. **Condition category** (more shortly)
2. **Construction year** (4-digit year, or 0)
3. **Site size** and units (ha/ac/sqm/sqft, or 0/"na")
4. **Construction period** (e.g. "Georgian", "Interwar", "Period")
5. **Garden indicators:** `has_garden`, `front_garden`, `back_garden`
6. **Garden size** `garden_size` + `garden_size_units` (sqm/sqft/length units, or 0/"na")
7. **Named style (rare):** `construction_style` (eg. Strain/Stringer)
8. **Views:** `sea_view`, `green_view`

LLM extraction: from descriptions to structured variables (prompted JSON)

- **Input:** full listing description text.
- **Output:** a **single strict JSON object** per listing, following a fixed schema (no free text).
  - Variables are **explicitly defined** and **bounded** (limited categories / units)
  - Missingness is standardised (`0` or `"na"`)
  - Extraction is **auditable** (prompt + schema + examples are fixed)

Extracted fields

1. **Condition category** (more shortly)
2. **Construction year** (4-digit year, or 0)
3. **Site size** and units (ha/ac/sqm/sqft, or 0/"na")
4. **Construction period** (e.g. “Georgian”, “Interwar”, “Period”)
5. **Garden indicators:** `has_garden`, `front_garden`, `back_garden`
6. **Garden size** `garden_size` + `garden_size_units` (sqm/sqft/length units, or 0/"na")
7. **Named style (rare):** `construction_style` (eg. Strain/Stringer)
8. **Views:** `sea_view`, `green_view`

LLM extraction: from descriptions to structured variables (prompted JSON)

- **Input:** full listing description text.
- **Output:** a **single strict JSON object** per listing, following a fixed schema (no free text).
    - Variables are **explicitly defined** and **bounded** (limited categories / units)
    - Missingness is standardised (`0` or `"na"`)
    - Extraction is **auditable** (prompt + schema + examples are fixed)

Extracted fields

1. **Condition category** (more shortly)
2. **Construction year** (4-digit year, or 0)
3. **Site size** and units (ha/ac/sqm/sqft, or 0/"na")
4. **Construction period** (e.g. "Georgian", "Interwar", "Period")
5. **Garden indicators:** `has_garden`, `front_garden`, `back_garden`
6. **Garden size** `garden_size` + `garden_size_units` (sqm/sqft/length units, or 0/"na")
7. **Named style (rare):** `construction_style` (eg. Strain/Stringer)
8. **Views:** `sea_view`, `green_view`

LLM extraction: from descriptions to structured variables (prompted JSON)

- **Input:** full listing description text.
- **Output:** a **single strict JSON object** per listing, following a fixed schema (no free text).
  - Variables are **explicitly defined** and **bounded** (limited categories / units)
  - Missingness is standardised (`0` or `"na"`)
  - Extraction is **auditable** (prompt + schema + examples are fixed)

Extracted fields

1. **Condition category** (more shortly)
2. **Construction year** (4-digit year, or 0)
3. **Site size** and units (ha/ac/sqm/sqft, or 0/"na")
4. **Construction period** (e.g. “Georgian”, “Interwar”, “Period”)
5. **Garden indicators:** `has_garden`, `front_garden`, `back_garden`
6. **Garden size** `garden_size` + `garden_size_units` (sqm/sqft/length units, or 0/"na")
7. **Named style (rare):** `construction_style` (eg. Strain/Stringer)
8. **Views:** `sea_view`, `green_view`

LLM extraction: from descriptions to structured variables (prompted JSON)

- **Input:** full listing description text.
- **Output:** a **single strict JSON object** per listing, following a fixed schema (no free text).
    - Variables are **explicitly defined** and **bounded** (limited categories / units)
    - Missingness is standardised (`0` or `"na"`)
    - Extraction is **auditable** (prompt + schema + examples are fixed)

Extracted fields

1. **Condition category** (more shortly)
2. **Construction year** (4-digit year, or 0)
3. **Site size** and units (ha/ac/sqm/sqft, or 0/"na")
4. **Construction period** (e.g. “Georgian”, “Interwar”, “Period”)
5. **Garden indicators:** `has_garden`, `front_garden`, `back_garden`
6. **Garden size** `garden_size` + `garden_size_units` (sqm/sqft/length units, or 0/"na")
7. **Named style (rare):** `construction_style` (eg. Strain/Stringer)
8. **Views:** `sea_view`, `green_view`

LLM extraction: from descriptions to structured variables (prompted JSON)

- **Input:** full listing description text.
- **Output:** a **single strict JSON object** per listing, following a fixed schema (no free text).
  - Variables are **explicitly defined** and **bounded** (limited categories / units)
  - Missingness is standardised (`0` or `"na"`)
  - Extraction is **auditable** (prompt + schema + examples are fixed)

Extracted fields

1. **Condition category** (more shortly)
2. **Construction year** (4-digit year, or 0)
3. **Site size** and units (ha/ac/sqm/sqft, or 0/"na")
4. **Construction period** (e.g. “Georgian”, “Interwar”, “Period”)
5. **Garden indicators:** `has_garden`, `front_garden`, `back_garden`
6. **Garden size** `garden_size` + `garden_size_units` (sqm/sqft/length units, or 0/"na")
7. **Named style (rare):** `construction_style` (eg. Strain/Stringer)
8. **Views:** `sea_view`, `green_view`

How the extraction is constrained (to reduce inference and drift)

- **Strict schema:** model must output *only* the JSON object with the specified keys.
- **No inference rule:** extract *only explicitly mentioned* information; otherwise return `0` / `"na"` / `"false"`.
- **Unit restrictions:** only allowed units are accepted (e.g. site size: ha/ac/sqm/sqft).
- **Prioritisation rules:**
  - if multiple site-size units appear, prioritise: ha $>$ ac $>$ sqm $>$ sqft
  - **do not** use garden size as site size (and vice versa)
  - if only garden dimensions given (e.g. 5m$\times$10m), record the *larger dimension* as a length unit
- **Multiple properties rule:** if more than one property is described, extract for the **first** only.

Is it a black box?

Extraction is a constrained measurement step that can be documented and reproduced.

Measuring internal condition: the 1–5 category definition

Condition categories

- **1: Luxurious / high-end**: e.g. magnificent, exclusive, elegant, period features, architect-designed, branded finishings
- **2: Excellent / turnkey**: e.g. showhouse, walk-in, high-quality finish, "turn-key"
- **3: Good / well-maintained**: e.g. well-maintained; recently modernised/upgraded/renovated
- **4: Needs work**: e.g. potential; requires upgrades; needs cosmetic upgrades
- **5: Derelict / major refurbishment**
- **0: Cannot classify** from explicit text cues

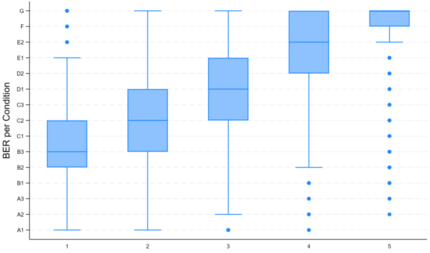

LLM-derived quality measures for RPPI compilation

- **Internal condition (1–5) targets the core unobserved dimension.**
  - Captures finish, refurbishment, and maintenance language strongly capitalised into prices.
  - Separates “nice” from “new”: a 1960s home can be high-condition if refurbished.
  - Ordinal, interpretable categories (luxury → turnkey → average → needs work → derelict).
  - Strict rule-based assignment: only explicit cues are coded; otherwise 0.
- **Additional extracted variables capture composition beyond condition.**
  - Site size (relevant in rural and edge-of-urban markets).
  - Garden presence/size and views (amenities often missing in structured data).
  - Construction period/style and year (period premia when explicitly marketed).
- **Integration into hedonic RPPIs:**
  - Compare specifications with and without LLM (or RegEx phrases) variables, holding sample (and other variables) constant
  - Can also check robustness using transaction prices, matched via Eircode identifier

Relevance for official statistics

The contribution is not a single score, but a general framework: converting free text into bounded, auditable variables that can be incorporated into standard RPPI methodologies.

Roadmap

Motivation & Context

Data & Measurement

Empirical Strategy

Results

Implications and Conclusions

Empirical strategy

- **Goal:** test whether adding LLM-derived variables changes the RPPI path.
- **Design principle:**
  - Same sample (2014–2025).
  - Same geography and time structure.
  - Same baseline hedonic controls.
- Compare four nested specifications:
  - (1) Structured fields only.
  - (2) + Phrase indicators.
  - (3) + LLM variables.
  - (4) + Phrases and LLM combined.

Key comparison

Do LLM variables materially change (i) model fit and (ii) the implied RPPI growth path?

Hedonic specification

$$\ln P_{it} = \underbrace{\alpha_t}_{\text{Time effects (RPPI)}} + \mathbf{X}_{it}^{\text{struct}}\beta + \mathbf{X}_{it}^{\text{text}}\theta + \varepsilon_{it}$$

- P_{it}: listed (or transaction) price.
- α_t: month fixed effects (the implied RPPI); key question is whether text variables shift α_t.
- $\mathbf{X}^{\text{struct}}$: BER, size bin, new development, beds $\times$ type, etc.
- $\mathbf{X}^{\text{text}}$: phrase indicators and/or LLM-derived variables.

All-in-one (vs. rolling windows)

Main: pooled "all-in-one" regressions, transparent RMSE comparisons across specifications Robustness: "rolling windows" regressions, allows coefficients to evolve over the cycle

Roadmap

Motivation & Context

Data & Measurement

Empirical Strategy

Results

Implications and Conclusions

Model fit and internal condition effects (pooled regressions)

	(1) Basic	(2) Phrase	(3) LLM	(4) All
Internal condition (base = Category 4)				
Category 1			0.296	0.259
			(28)	(26)
Category 2			0.119	0.106
			(26)	(25)
Category 3			0.056	0.056
			(15)	(16)
Category 5			-0.149	-0.135
			(-11)	(-9.7)
RegEx phrases	N	Y	N	Y
LLM variables	N	N	Y	Y
Micro-market FE	Y	Y	Y	Y
Other controls	Y	Y	Y	Y
Observations	731,543	731,543	731,543	731,543
RMSE	0.277	0.264	**0.250**	**0.244**
R^2	0.811	0.829	0.846	0.853

Large, precisely estimated condition effects and substantial RMSE reduction when LLM variables are included

- **Internal condition is strongly priced.**
  - Category 1 (luxury): +26-30% relative to base.
  - Category 5 (derelict): -13-15%.
  - All highly statistically significant.
- **LLM variables materially improve fit.**
  - RMSE falls from 0.277 $\rightarrow$ 0.250 ($\simeq$ 10% reduction).
  - Further falls to 0.244 when combined with phrase indicators.
  - Gains larger than phrases alone.
- **Interpretation for RPPI work:**
  - LLM variables capture economically meaningful quality variation.
  - They substantially reduce residual cross-sectional dispersion.
  - The key question becomes: are these quality differences sufficiently time-varying to alter the index path?

Takeaway

LLM-based internal condition is economically large and statistically precise – but its importance for aggregate price measurement depends on time variation, not cross-sectional pricing alone.

Residential Property Price Indices, with/without LLM variables

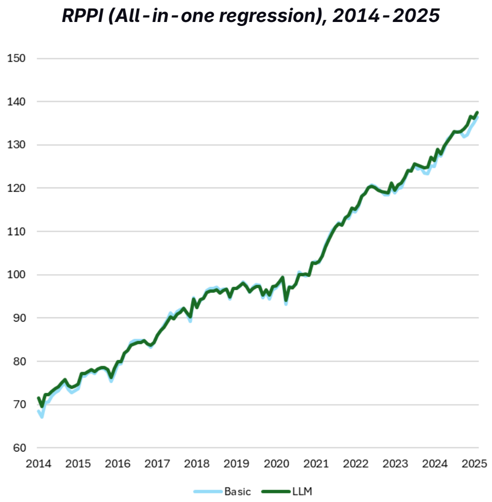

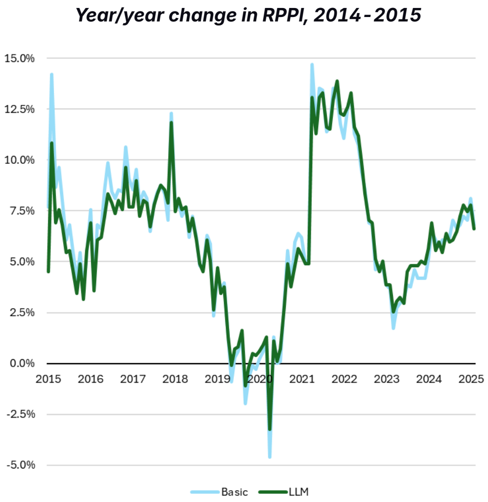

RPPI comparison: Do LLM variables change the index path?

Compare index levels and yr/yr growth 2014-2025

- **Index levels:** LLM and All specs sit slightly above Basic/Phrase early on, but differences are small (typically 1–2 index points).
- **Growth rates:** Year-on-year changes are highly similar across specifications.
- **Turning points unchanged:** 2019 slowdown, 2020 flat, 2021–22 surge, 2023 moderation all identical in timing and magnitude.

LLM variables materially improve cross-sectional fit, but have only **modest effects on aggregate RPPI growth**. Existing hedonic indices appear broadly robust to omitted internal condition.

Jan Index	Basic	Phrase	LLM	All
2014	68.5	69.8	71.5	71.4
2015	73.7	75.5	74.8	75.4
2016	79.3	80.7	79.9	80.5
2017	86.1	87.2	86.1	86.6
2018	92.8	93.4	92.5	93.0
2019	96.9	97.5	96.9	97.2
2020	97.0	97.4	97.4	97.5
2021	103.0	103.4	102.6	103.0
2022	114.5	114.9	115.1	115.3
2023	118.9	119.1	119.6	119.7
2024	125.0	125.0	126.4	126.2
2025	135.1	134.7	136.2	135.9

YoY Growth	Basic	Phrase	LLM	All
2015	7.7%	8.1%	4.5%	5.5%
2016	7.6%	6.8%	6.9%	6.8%
2017	8.5%	8.1%	7.7%	7.6%
2018	7.8%	7.1%	7.5%	7.4%
2019	4.4%	4.4%	4.7%	4.6%
2020	0.2%	-0.1%	0.6%	0.3%
2021	6.2%	6.1%	5.3%	5.7%
2022	11.1%	11.2%	12.2%	11.9%
2023	3.9%	3.7%	3.9%	3.9%
2024	5.1%	4.9%	5.7%	5.4%
2025	8.1%	7.8%	7.8%	7.7%

Why doesn't the index move? Cross-section vs time

- Internal condition has large price effects (Category 1: +26-30%; Category 5: 13-15%)
- But the RPPI depends on **time variation** in quality mix.

Condition distribution is remarkably stable

- Cat2 rises gradually (36% → 42%)
- Cat3 falls modestly (49% → 43%)
- Cat1 and Cat5 remain rare ($\simeq$ 3%)

Implication

LLM condition variables explain **cross-sectional dispersion**, but the mix of condition changes only slowly over time.
$\Rightarrow$ Limited impact on aggregate RPPI growth.

Share of listings by internal condition score, 2015-2025

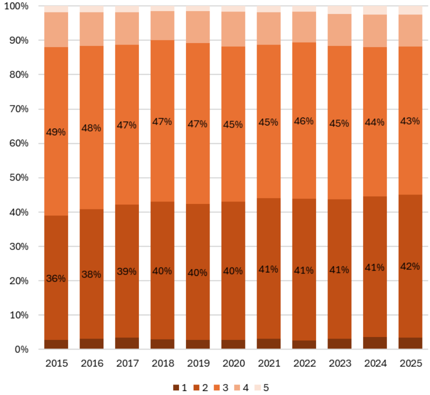

Might LLM variables matter more with missing variables?

- Two important but non-mandatory structured variables that may be correlated with internal condition: Energy rating (BER / EPC), Floor area (sqm)
- Coverage improves dramatically over time, esp as BER made mandatory in mid-2020:
- 2014: 38% missing EPC, 49% missing sqm
- 2024–25: ~2% missing EPC, ~16% missing sqm

Share of listings with missing BER and floor area (sqm), 2014-2025

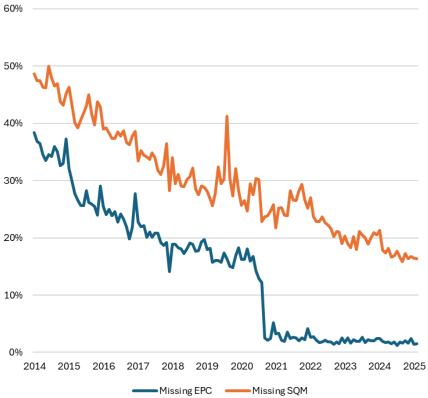

Interpretation

Nonetheless, very little difference – e.g. in quarterly changes – with/without LLM variables before/after mid-2020 (or at any cutoff).

Might LLM variables matter more with missing variables?

- Two important but non-mandatory structured variables that may be correlated with internal condition: Energy rating (BER / EPC), Floor area (sqm)
- Coverage improves dramatically over time, esp as BER made mandatory in mid-2020:
- 2014: 38% missing EPC, 49% missing sqm
- 2024–25: ~2% missing EPC, ~16% missing sqm

Share of listings with missing BER and floor area (sqm), 2014-2025

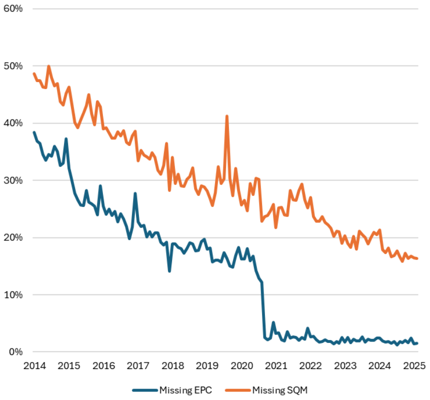

Interpretation

Nonetheless, very little difference – e.g. in quarterly changes – with/without LLM variables before/after mid-2020 (or at any cutoff).

Robustness checks

- **Alternative outcomes:** listed vs transaction prices (matched sample).
- **Alternative index approaches (optional):** rolling window vs all-in-one.
- **Alternative FE structure:** finer/coarser geography; market-time definitions.

Roadmap

- Motivation & Context
- Data & Measurement
- Empirical Strategy
- Results
- Implications and Conclusions

Practical considerations for official statistics

- **1. LLM variables are feasible at scale**
  - 700k+ listings processed with near-complete coverage, structured JSON output enables reproducibility.
  - Integration into existing hedonic pipelines is straightforward.
- **2. Gains are primarily cross-sectional**
  - Substantial RMSE reduction.
  - Large and precisely estimated condition effects.
  - Limited impact on aggregate index growth.
- **3. Governance and version control matter**
  - Documentation essential; historical backfills must be reproducible.
  - Changes in model versions could introduce revisions if not managed carefully.
- **4. Where impact may be larger**
  - Markets with weaker structured data coverage.
  - Periods with sharp shifts in quality mix.
  - Countries without mandatory energy disclosure.

Overall

LLM-derived variables can strengthen measurement without destabilising the RPPI.

Practical considerations for official statistics

- **1. LLM variables are feasible at scale**
  - 700k+ listings processed with near-complete coverage, structured JSON output enables reproducibility.
  - Integration into existing hedonic pipelines is straightforward.
- **2. Gains are primarily cross-sectional**
  - Substantial RMSE reduction.
  - Large and precisely estimated condition effects.
  - Limited impact on aggregate index growth.
- **3. Governance and version control matter**
  - Documentation essential; historical backfills must be reproducible.
  - Changes in model versions could introduce revisions if not managed carefully.
- **4. Where impact may be larger**
  - Markets with weaker structured data coverage.
  - Periods with sharp shifts in quality mix.
  - Countries without mandatory energy disclosure.

Overall

LLM-derived variables can strengthen measurement without destabilising the RPPI.

Practical considerations for official statistics

- **1. LLM variables are feasible at scale**
  - 700k+ listings processed with near-complete coverage, structured JSON output enables reproducibility.
  - Integration into existing hedonic pipelines is straightforward.
- **2. Gains are primarily cross-sectional**
  - Substantial RMSE reduction.
  - Large and precisely estimated condition effects.
  - Limited impact on aggregate index growth.
- **3. Governance and version control matter**
  - Documentation essential; historical backfills must be reproducible.
  - Changes in model versions could introduce revisions if not managed carefully.
- **4. Where impact may be larger**
  - Markets with weaker structured data coverage.
  - Periods with sharp shifts in quality mix.
  - Countries without mandatory energy disclosure.

Overall

LLM-derived variables can strengthen measurement without destabilising the RPPI.

Practical considerations for official statistics

- **1. LLM variables are feasible at scale**
  - 700k+ listings processed with near-complete coverage, structured JSON output enables reproducibility.
  - Integration into existing hedonic pipelines is straightforward.
- **2. Gains are primarily cross-sectional**
  - Substantial RMSE reduction.
  - Large and precisely estimated condition effects.
  - Limited impact on aggregate index growth.
- **3. Governance and version control matter**
  - Documentation essential; historical backfills must be reproducible.
  - Changes in model versions could introduce revisions if not managed carefully.
- **4. Where impact may be larger**
  - Markets with weaker structured data coverage.
  - Periods with sharp shifts in quality mix.
  - Countries without mandatory energy disclosure.

Overall

LLM-derived variables can strengthen measurement without destabilising the RPPI.

Concluding thoughts

- **Internal condition is economically large and statistically precise.**
  - +30% for top category and -15% for derelict, other things being equal.
- **LLM variables materially improve model fit.**
  - Reduction in RMSE by about one tenth relative to baseline.
  - Stronger contribution than phrase-based NLP alone.
- **But aggregate RPPI dynamics are highly robust.**
  - Index levels differ only slightly.
  - Growth rates and turning points unchanged.
  - Condition mix varies slowly over time.
- **LLM-based measurement improves precision.** It refines, rather than overturns, official price indices.

Bottom line

Modern text-based measurement can enhance RPPIs, while preserving their stability and credibility.

Thank you.

Concluding thoughts

- **Internal condition is economically large and statistically precise.**
  - +30% for top category and -15% for derelict, other things being equal.
- **LLM variables materially improve model fit.**
  - Reduction in RMSE by about one tenth relative to baseline.
  - Stronger contribution than phrase-based NLP alone.
- **But aggregate RPPI dynamics are highly robust.**
  - Index levels differ only slightly.
  - Growth rates and turning points unchanged.
  - Condition mix varies slowly over time.
- **LLM-based measurement improves precision.** It refines, rather than overturns, official price indices.

Bottom line

Modern text-based measurement can enhance RPPIs, while preserving their stability and credibility.

Thank you.