Measuring rental inflation: Evidence from Ireland

Janez Kren & Rachel Slaymaker

ESRI Working Paper No. 821

May 2026

Measuring rental inflation: Evidence from Ireland

Janez Kren[1,2] and Rachel Slaymaker*[1,2]

[1]Economic and Social Research Institute
[2]Trinity College Dublin

Abstract

Accurate measurement of rental inflation is essential for understanding housing market dynamics, informing policy, and ensuring reliable inflation statistics. This paper constructs a repeat rent index (RRI) using newly available administrative data from the Residential Tenancies Board (RTB), with near population-level coverage of both new and ongoing tenancies from Q2 2022-Q1 2025. Applying both hedonic and RRI methodologies to the same dataset isolates methodological effects from data scope and coverage effects and enables a formal decomposition of hedonic inflation estimates. Nationally, annual rental inflation measured by the RRI is consistently 2.1–2.8 percentage points lower than hedonic estimates. Decomposition shows that within-property rent growth accounts for 53% of hedonic inflation in this period, while entry and exit effects contribute roughly one-third, reflecting higher standardised rents for entrants and lower rents for exits. These findings highlight the role of market churn in shaping hedonic estimates and illustrate the importance of index choice for policy evaluation. Hedonic indices are best suited for capturing market-level dynamics, while RRIs provide a more accurate measure of affordability trends and compliance with rent controls. Our results demonstrate the value of comprehensive administrative data for improving rental market measurement and informing housing policy in regulated markets.

Keywords: Rental Inflation; Repeat Rent Index; Hedonic Index

JEL codes: R31, R21, E31

---

*Corresponding author: rachel.slaymaker@esri.ie.
This work was funded under a programme of research between the Economic and Social Research Institute and the Residential Tenancies Board. We thank officials at the RTB for access to their microdata. The views expressed in this paper are solely those of the authors.

Non-Technical Summary

Accurate measurement of rental inflation is essential for understanding rental market dynamics, for assessing affordability, and informing broader housing market and economic policies. This study uses newly available administrative data from the Residential Tenancies Board (RTB), which provides coverage of both new and ongoing tenancies in the private rental sector following a 2022 legislative reform requiring annual registration of all tenancies. The availability of these data allows for the construction of a Repeat Rent Index (RRI) for the first time in Ireland and enables a direct comparison with the hedonic rent index method currently used.

The key distinction between these two methodologies lies in what they measure. The RRI tracks the same properties over time, capturing rent changes within individual properties. This makes it particularly suited to measuring affordability trends for households (and properties) in the sector and for assessing the effectiveness of rent caps in limiting within-property rent rises. In contrast, the hedonic index does not track individual properties; instead, it uses all observed rents in each period and adjusts for differences in property characteristics. This means hedonic measures capture changes in the average rent paid and also reflect changes in the composition of the rental market over time, including the entry of new properties and the exit of other properties.

Our analysis for the period Q2 2022-Q1 2025 shows that annual rental inflation measured by the RRI was consistently 2.1–2.8 percentage points lower than hedonic estimates. A formal decomposition reveals that 53 per cent of hedonic inflation in this period reflects within-property rent growth, while roughly one-third is driven by market churn, the entry and exit of properties with systematically different rent levels. Properties entering the sector over this period had rents around 11 per cent higher than those exiting, highlighting the role of property turnover in shaping hedonic estimates.

Local authority level analysis reinforces these findings. In rent-controlled urban areas, RRI estimates of rental inflation were low, while hedonic estimates often exceeded rent cap thresholds. Importantly, two local areas can display similar hedonic inflation rates, yet the underlying drivers differ markedly. This distinction matters because it signals different affordability pressures: churn-driven inflation reflects structural market changes, while within-tenancy growth directly impacts sitting tenants more prominently.

These differences have important implications for interpretation. Hedonic indices are crucial for understanding broader market pressures, driven by factors such as population growth, supply constraints, elevated construction costs, and high interest rates, all of which influence the mix of properties and overall rent levels. However, they can overstate the inflation experienced by households in the sector and lead to misinterpretation of the effectiveness of rent regulation in limiting rent increases. Conversely, the RRI provides a timely and clearer picture of affordability developments for households, but does not capture the impact of market churn on overall rent levels in the sector.

Using both indices together provides a more comprehensive and timely understanding of rental market developments. This research demonstrates the value of leveraging administrative data to improve the accuracy and transparency of rental market indicators.

1 Introduction

Housing costs account for a substantial share of household expenditure. Accurate measurement of rental inflation is critical for understanding housing market dynamics, assessing affordability, and informing policy interventions. Rising house prices and declining homeownership rates across many advanced economies have increased reliance on private rental sectors for housing provision, while affordability pressures have prompted renewed interest in regulatory tools such as rent controls (Kholodilin, 2024). Against this backdrop, the construction of reliable rent indices has become central not only to housing policy evaluation, but also to ensuring the accuracy of headline inflation measures such as the Consumer Price Index (CPI) and resulting broader macroeconomic analysis.

Recent empirical evidence highlights significant discrepancies between commonly used US rent index measures (Adams et al., 2024) and the material bias this can introduce into the CPI (Ambrose et al., 2023). By extension, inaccurate rent measurement can distort macroeconomic policy assessments, leading to misguided interpretations of monetary policy effectiveness (Ambrose et al., 2018), while improved measurement enhances understanding of its distributional impacts (Abramson et al., 2025). Beyond macroeconomic considerations, inaccurate rent measures risk misinterpreting the effectiveness of housing policies designed to address affordability concerns. These issues are particularly salient in regulated rental markets, where published rent indices often underpin assessments of compliance with rent caps.

Whereas repeat-index methods dominate US housing price measurement, European measures typically rely on hedonic models (Spiegel, 2025). Despite the importance of accurate rental inflation measurement, decisions around the method and coverage of rent index indicators are often restricted by data availability. Compared to house price data, there has traditionally been a lack of high-quality rental price data (Ambrose et al., 2015). Existing approaches frequently rely on survey-based indices, which suffer from non-response bias and under-reporting of rent changes during tenant turnover (Crone et al., 2010), or on listing-based measures that exhibit selective coverage and bias toward high-turnover or premium market segments (Adams et al., 2024). A series of recent studies have sought to improve the measurement of rental inflation (Abramson et al., 2025; Adams & Verbrugge, 2025; Adams et al., 2024; Ambrose et al., 2015, 2023; Boesel et al., 2021; Clark, 2022; Kim, 2024), yet applications in regulated European rental markets remain scarce, largely due to the absence of comprehensive administrative data.

In this paper we address these gaps by exploiting newly available administrative data covering the universe of Irish private rental tenancies to construct a repeat rent index (RRI) and compare it to the hedonic approach currently used. Applying both methods to the same dataset isolates methodological effects from data source differences and enables a formal decomposition of hedonic inflation estimates into within-property growth and entry/exit effects. Our analysis draws on a dataset created by a 2022 legislative change requiring the annual registration of all private tenancies with Ireland's rental regulator, the Residential Tenancies Board (RTB), providing near population-level coverage of both new and ongoing tenancies over the period Q2 2022–Q1 2025. This dataset avoids the selective coverage of commercial datasets that can exhibit bias toward certain market segments or property types (Adams et al., 2024) and overcomes the non-response bias common in survey data (Crone et al., 2010). Moreover, these data permit the construction of granular, local-level rent index estimates, a notable challenge for survey-based indices outside major urban markets (Boesel et al., 2021).

A number of papers are closely related to our research, but our work is most similar to Adams et al. (2024), who examine the divergence between different rent index estimates in the US and identify the factors driving these differences. While their analysis focuses on scope and dataset variation using RRI methods, we instead isolate methodological effects within a single, comprehensive administrative dataset. Specifically, we introduce a formal decomposition of hedonic inflation estimates into within-property growth, entry and exit effects, and residual components, providing a clearer understanding of why hedonic and repeat rent indices diverge. Our use of near population-level administrative microdata enables this decomposition both nationally and across local authorities, providing granular insights into rental price dynamics that have not previously been possible. Our contribution is threefold: (i) we exploit a new administrative dataset covering the universe of Irish tenancies; (ii) we construct and compare hedonic and repeat rent indices within a single data source; and (iii) we formally decompose hedonic inflation into within-property growth and entry/exit effects, providing new insights into rental market dynamics.

Our analysis shows that annual rental inflation measured by the RRI is consistently 2.1–2.8 percentage points lower than hedonic estimates. Our formal decomposition reveals that approximately one-third of the inflation captured by the hedonic index stems from differences in standardised average rent levels between properties entering and exiting the market, underscoring the role of property turnover in shaping hedonic estimates. In contrast, the within-property component, captured by the RRI, accounts for 53%. This indicates that market churn is the primary driver of divergence between the two indices. In rent-controlled, urban markets, RRI-based estimates of rental inflation were typically low compared to currently published hedonic figures often in excess of permitted rent cap thresholds. Conversely, in non-urban, non-rent-controlled markets, the within-property growth rates, captured by RRI, were much higher, while the effects of change in composition were smaller. This highlights the need to distinguish churn-driven inflation from rising affordability pressures for sitting tenants.

These findings have important implications for both measurement and policy. The hedonic index is best suited to capturing market-level dynamics and the influence of churn on average rent levels, reflecting broader pressures such as population growth, supply shocks, elevated construction costs, and high interest rates. In contrast, the RRI better represents household affordability trends and compliance with rent caps by isolating within-property changes. This distinction is critical for policymakers assessing market dynamics, the effectiveness of rent regulation, identifying areas where caps may require recalibration, and evaluating affordability trends over time. More broadly, the choice of index should be guided by the specific policy question at hand and whether it concerns aggregate market dynamics or the lived experience of households in the rental sector.

The remainder of the paper is structured as follows. Section 2 reviews relevant literature. Section 3 outlines the institutional context of the Irish rental sector and introduces the dataset. Section 4 describes the hedonic and repeat rent index methodologies. Section 5 presents rental inflation estimates from both approaches and provides a decomposition explaining their divergence. Finally, section 6 concludes

2 Related literature

This paper contributes to three strands of the real estate economics literature. First, recent studies have sought to improve the measurement of housing and rental price dynamics through the use of novel data sources (Abramitzky et al., 2025). Examples include Gindelsky et al. (2021), who leverage granular Zillow data to better estimate housing services and costs, and Bricongne et al. (2023), who employ high-frequency web-scraped data to track housing market developments during the COVID-19 pandemic. These contributions highlight the growing importance of granular, timely data for understanding housing market developments. Our work builds on this by exploiting comprehensive administrative microdata covering the near-universe of Irish private rental tenancies, overcoming the non-response bias limitations of survey data (Crone et al., 2010) and the selective coverage of listings-based datasets (Adams et al., 2024).

Second, our work relates to the methodological literature on housing price index construction, particularly debates around the merits and drawbacks of hedonic versus repeat index approaches (De Haan, 2013; De Haan & Diewert, 2013; Hill, 2013; Hill & Trojanek, 2022; Malpezzi et al., 2003). Hedonic models utilise all available observations and adjust for property characteristics, allowing timely incorporation of new entrants, but are sensitive to specification and prone to omitted variable bias (De Haan & Diewert, 2013; Malpezzi et al., 2003). In contrast, repeat-sales methods mitigate quality-related bias by tracking the same unit across periods (De Haan, 2013), yet suffer from data attrition, representativeness concerns, and lagged responsiveness to market changes. Recent methodological contributions have sought to improve repeat-sales methods in contexts where data are sparse or volatile, such as granular local markets. For example, Contat and Larson (2024) propose a new algorithm to estimate city-level repeat-sales price indices, by aggregating heterogeneous sub-markets or populations, permitting granular analysis while mitigating sampling bias in markets with limited transactions. While these papers typically focus on house prices, the underlying issues of measurement accuracy, data coverage, sample representativeness and spatio-temporal trade-off are also relevant to rental markets.

Third, our work is most closely related to recent advances in rent index construction, specifically through the use of repeat rent index (RRI) models (Adams & Verbrugge, 2025; Adams et al., 2024; Ambrose et al., 2015, 2023; Boesel et al., 2021; Clark, 2022; Kim, 2024). Ambrose et al., 2015 introduced the RRI, applying the repeat-sales methodology (Bailey et al., 1963) to the rental market, while Ambrose et al. (2023) highlight the material bias inaccurate rental inflation measurement can have on the CPI. Boesel et al. (2021) and Adams and Verbrugge (2025) both highlight how property heterogeneity and location-specific factors influence rent inflation estimates. Boesel et al. (2021) show that shifts in housing preferences during the COVID-19 pandemic, i.e. greater demand for single-family detached homes in suburban locations, created compositional effects that distorted rent indices. Similarly, Adams and Verbrugge, 2025 note that inflation rates can differ substantially between unit types within the same local area, highlighting the need for structural, as well as locational representativeness. Our work is most similar to Adams et al. (2024) who examine the divergence between different rent index estimates in the US and find both scope (new versus all tenancies) and data coverage to be the key factors driving these differences.

Our study builds on these contributions by applying repeat rent index methods to a regulated European rental market using near-population level administrative data. Specifically, we compare rental inflation estimates from both RRI and hedonic methods within a single dataset, allowing us to isolate methodological effects from data source differences. We also formally decompose hedonic inflation into within-property growth and entry/exit effects, providing new insights into the drivers of rental inflation.

3 Institutional background and data

Historically considered a transitory tenure, the private rental sector (PRS) in Ireland has expanded significantly over the past three decades. According to the 2022 Census, nearly 20% of Irish households now reside in privately rented accommodation, up from just 8% in 1991. This shift reflects broader trends in declining homeownership and rising affordability pressures, which have been well documented in the Irish context (Corrigan et al., 2019; Disch & Slaymaker, 2023; Slaymaker et al., 2022).

Figure 1 presents the long-run trend in rental inflation for new tenancies, based on the ESRI/RTB hedonic rent index. It shows a sharp acceleration in rents beginning in 2013, coinciding with Ireland's post-global financial crisis economic recovery. In response to these pressures, the Irish government introduced a system of rent controls in late 2016 known as Rent Pressure Zones (RPZs). Annual rent increase caps of 4% applied in these areas, with both within-tenancy and between-tenancy rent adjustments regulated. These caps were effective in moderating rental inflation in designated areas (O'Toole et al., 2021), with national rent growth falling from 11% in early 2017 to 4.1% prior to the COVID-19 pandemic.

Despite a further tightening of RPZ rules in 2021, linking allowable rent increases to the lower of either Harmonised Index of Consumer Prices (HICP) inflation or a 2% ceiling, rental inflation has re-accelerated in the post-COVID-19 pandemic period. By Q2 2023, new tenancy rents were rising at an annual rate of 11%. A newly developed hedonic index for existing tenancies i.e., sitting tenants, first introduced in 2022 after the change in data collection, also indicates sustained inflation of 5-6% throughout late 2023 and the first half of 2024 (Figure 1). These figures have been widely interpreted as evidence of increased demand-side pressures and widespread non-compliance with rent control regulations. However, using matched property-level data from the Residential Tenancies Board (RTB), covering over 180,000 properties, Slaymaker et al. (2024) find that the average rent increase for individual properties was just 2.6% over the same period, notably lower than both published index estimates. This divergence in the post-COVID-19 period motivates the need to explore alternative rent index methodologies. In particular, it raises questions about how different index constructions and underlying data samples capture market churn versus household affordability developments and, where applicable, the effectiveness of policies designed to limit rental inflation. These issues are central to better understanding rental price trends not only in Ireland, but across countries facing similar affordability pressures.

Figure 1: Evolution of standardised average rents and annual inflation rates - New and existing tenancies

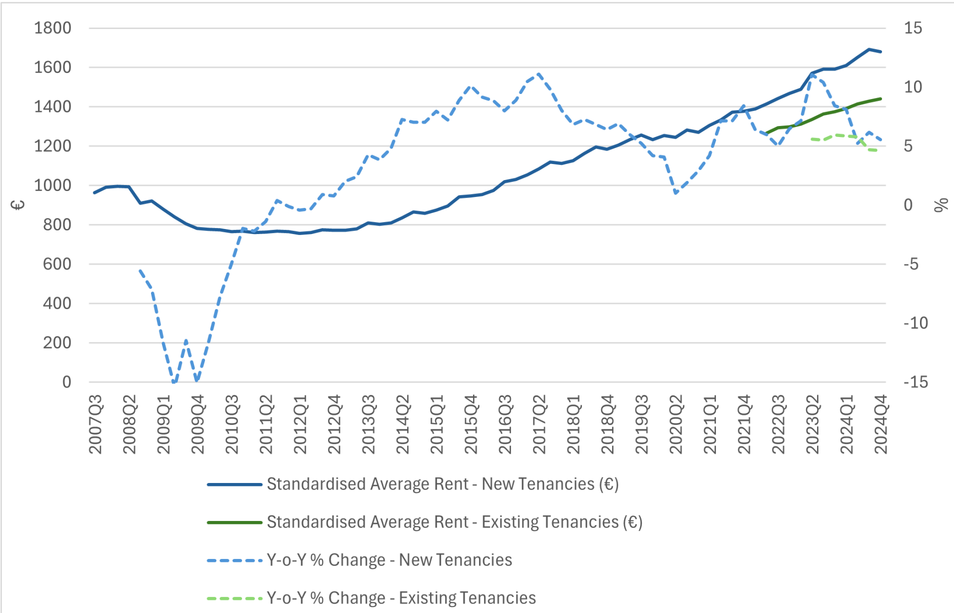

Source: ESRI/RTB Rent Price Index.

Our analysis draws on administrative data from Ireland's rental regulator, the Residential Tenancies Board (RTB), covering the universe of rental tenancy registrations from Q2 2022-Q1 2025. A legislative reform enacted in April 2022 made it mandatory for landlords to register ongoing tenancies annually, in addition to the longstanding requirement to register new tenancies at commencement[1]. While RTB data on new tenancies have long been used to construct hedonic rent indices (Coffey et al., 2022), this legislative reform marks a significant expansion in data coverage. Post reform, this dataset now provides near-population level coverage of Ireland's formal rental sector on an annual basis. Annual, property-level coverage nationwide, capturing both new and ongoing tenancies, permits the construction of a repeat rent index for the first time.

The dataset includes detailed information on contracted rent, payment frequency, tenancy start date, registration type (new or ongoing), full address (including Eircode) and property characteristics (e.g., dwelling and property type, number of bedrooms and occupants). It excludes student-specific and non-market price accommodation such as homes by local authorities and informal family arrangements without a formal tenancy agreement. Importantly, rents reflect full market prices, exclusive of subsidies. Unique identifiers for properties, tenancies and landlords allow for the precise tracking of tenancy status (new or ongoing) and rent adjustments over time.

Relative to existing data sources used in the literature, such as survey-based indices (e.g., BLS CPI in the US) or listing-based measures, the RTB dataset offers several key advantages. First, it avoids the non-response bias associated with tenant turnover (Crone et al., 2010). New tenancies and therefore rent increases that occur when new tenants move in are often under-reported in survey data, due to greater rates of sample attrition, leading to downward bias in rent index estimates. In this dataset, as registration is completed by landlords regardless of whether tenants change or remain, this potential bias is removed. Second, it enables granular, local-level analysis, overcoming the sample size limitations that affect survey data outside major urban centres (Abramson et al., 2025; Boesel et al., 2021).

Finally, unlike listing-based indices that only capture asking rents for new tenancies and/or other commercial datasets that are often biased toward high-turnover or premium market segments (Adams et al., 2024), the RTB data provide comprehensive coverage of actual contracted rents across the entire rental market. New tenancy rent indices derived from online listings platforms are particularly susceptible to sample selection bias and changing coverage over time. Several structural factors contribute to this: listing platforms typically charge fees, which incentivise selective use by landlords, especially with the rise of free alternatives such as social media. Moreover, in tight rental markets, properties may be let before

[1] Since December 2016, tenants in Ireland have been entitled to remain in a property for up to six years with no new lease agreement, regardless of any specified fixed term. This meant that prior to the 2022 reform, the rents for each property were typically only observed sporadically, unless tenants changed on a regular basis.

being publicly advertised. For example, Ireland's largest property listings platform recorded 1.5 listings per RTB registration in 2010, but this ratio fell to just 0.6 by 2024 (Gillespie et al., 2025), indicating a substantial erosion in coverage relative to the actual rental sector.

4 From a Hedonic to a Repeat Rent Index

Our aim is to examine how both the underlying data samples and the choice of index modelling method affect estimates of rental inflation. While Adams et al. (2024) examine various repeat style indices and focus primarily on differences in the scope of the underlying datasets (new or all tenancies, coverage etc.), we instead focus on the impact of the choice between hedonic and repeat index methods using data from only one source. We first create an all tenancies[2] hedonic rent index based on the approach outlined in Lawless et al. (2018) which takes the following form:

$$lnP_{it} = \beta_0 + \beta X_i + \gamma_1 T_1 + ... + \gamma_N T_N + \epsilon_{it} \quad (1)$$

where P_{it} is the rental price of property i at time t, X_i is a vector of property characteristics[3] and β is the vector of coefficients associated with each property characteristic. $T_1 - T_N$ are binary time dummy variables which equal 1 if a property is registered and its rent observed in that quarter and 0 otherwise and ϵ_{it} is an error term. In a hedonic framework, a property's rent level is determined by the property's attributes. The model adjusts and controls for the quality and characteristics of rental units to capture the mix of properties across time periods, while the coefficients on the time dummy variables ($\gamma_1 - \gamma_N$) capture changes in the price or rent of a constant quality representative dwelling i.e. the rent index.

An advantage of hedonic models is they utilise all available data i.e. no tracking of individual properties over time is required. The inclusion of all properties observed each period means the impact of new entrants on market prices is accounted for in a timely manner. The main limitation of the hedonic approach is around the selection of property and locational characteristics and concerns around omitted variable bias. In practice, decisions are largely influenced by data availability, but may make a material difference to estimates (Ambrose et al., 2015; Hill, 2013), though Malpezzi et al. (2003) argues this is less of an issue where the model is used to produce an index as the bias in one coefficient is likely offset by another. In our case, this bias may stem from poorly filled in energy efficiency ratings and the absence of property age indicators. In addition, many factors are difficult to measure and are therefore

[2] Currently separate ESRI/RTB rent index indicators for new and existing tenancies are published quarterly.

[3] Specifically we include dummy variables for the number of bedrooms, property type, local electoral area (LEA), as well as the number of tenants relative to the number of bedrooms.

typically unobserved in datasets. Examples include how well the property has been maintained, the quality of fixtures and fittings[4], sunlight, damp and noise.

One way to mitigate against omitted variable bias concerns in a hedonic framework is to estimate a repeat rent index model instead. Taking the hedonic model shown above in equation 1, the price of rental unit i in period t is given as follows:

$$lnP_{it} = \beta_0 + \beta X_i + \gamma_t T_t + \epsilon_{it} \tag{2}$$

while the price of rental unit i in a second period s is:

$$lnP_{is} = \beta_0 + \beta X_i + \gamma_s T_s + \epsilon_{is} \tag{3}$$

Subtracting 3 from 2 gives:

$$lnP_{it} - lnP_{is} = \gamma_t T_t - \gamma_s T_s + \epsilon_{it} - \epsilon_{is} \tag{4}$$

The change in rental price of unit i between periods s and t is purely a function of the two time periods and the change in the error term. In other words, the repeat rent index approach instead measures rental price growth by controlling for the time-invariant, observed and unobserved aspects of housing quality by tracking rental price developments for the same rental units over time, focusing on pairs of rent observations for the same unit. By comparing identical units over time, this approach minimises biases from changes in the composition and quality of the rental market as a whole. Long used for house prices (Case & Shiller, 1989), this approach has more recently been applied to rents (Ambrose et al., 2015; Boesel et al., 2021).

Based on the seminal work of Bailey et al. (1963), our simple repeat index takes the following form:

$$lnP_{it} - lnP_{is} = \gamma_1 D_{i1} + ... + \gamma_N D_{iN} + v_{it} \tag{5}$$

where $D_{ij} = -1$ if the first observation in a pair took place in period j, $D_{ij} = 1$ if the second observation in a pair took place in period j and $D_{ij} = 0$ otherwise. The γ parameters approximate percentage differences[5]. This method treats separate pairs of observations for the same property as independent from one another i.e. a pairwise fixed effect rather than a property fixed effect. Only comparing a property with its most recent registration aims to ensure we are comparing like with like and remove the effects of renovations or a deterioration in maintenance which may

4 In Ireland it is the norm for rental units to be let furnished.

5 This method estimates the geometric mean price growth rate. Boesel et al. (2021) instead estimate the arithmetic mean price growth rate by substituting the -1s for minus the initial price and the 1s for the property pair's second observed price.

occur over time. As our dataset records the near-population of rental contracts we do not weight our baseline rent index estimates. However, we do apply Census weights as a robustness check in section 5.4.

Case and Shiller (1989) argued the error term in equation 5 is heteroskedastic due to rising prices and irregular observation intervals. To address this, they and Ambrose et al. (2015) applied a three-step correction: OLS estimation of (5), regression of squared residuals to obtain fitted values, and weighted least squares using the inverse of these values. This gives greater weight to closely timed property pairs. Clark (2022), Adams et al. (2024) and Abramson et al. (2025) found this correction had minimal impact on index estimates. Given the annual frequency of our data, the impact should be minimal. Nevertheless, we include this adjustment as a robustness check in section 5.4.

The attraction of the repeat rent index method lies in its simplicity, requiring only information on price, period and a dwelling identifier. By default this approach controls for key factors that do not change over time such as location, as well as other unobservable property characteristics De Haan (2013). It does, however, have several limitations. First, it only includes properties observed in at least two periods, thereby discarding large volumes of data and raising concerns around sample selection bias. Unlike hedonic indices that incorporate new market entrants immediately, repeat indices are lagged measures that exclude new properties until a second observation occurs. Moreover, as they track rent changes rather than levels, they ignore the broader, market-wide changes in standardised average rental prices over time brought about by new entrants and property exits that are captured in hedonic indices (Malpezzi et al., 2003). We explore these issues empirically in the next section.

5 Results: comparing rent index estimates

This section compares rental inflation estimates derived from hedonic and repeat rent index (RRI) methodologies, and decomposes the observed differences to identify the underlying drivers.

5.1 Data sample

Our analysis draws on administrative data from the Residential Tenancies Board (RTB), the Irish rental regulator. The dataset comprises legally mandated registrations submitted by landlords, covering the universe of both new and ongoing private rental tenancy registrations between Q2 2022 and Q1 2025. We exclude late registrations, those referring to tenancies initiated prior to Q2 2022 but submitted after that date. To construct our hedonic sample, we winsorize monthly rent values at the 1st and 99th percentiles within each Local Electoral Area (LEA)[6] and quarter.[7] This results in a hedonic sample of 724,494 observations covering both new and ongoing tenancies.

For the repeat rent index (RRI) we construct a quarterly property-level panel dataset by removing duplicate registrations for properties observed more than once per quarter, those with inconsistently recorded property characteristics, as well as properties observed only once during the sample period. Properties observed only once, so called singletons, cannot be used to measure rent changes over time, and duplicates within a quarter may reflect administrative anomalies or multi-unit registrations. Finally we perform a second outlier removal for extreme values in the top and bottom 1% of rent growth pairs. These steps result in a final RRI sample of 590,017 registrations, or 373,858 matched property pairs (growth rates).

As the RRI methodology can only include properties observed at least twice, this raises concerns about sample representativeness and whether the subset of properties included in the RRI differs systematically from the broader sample used in the hedonic model, leading to potential selection bias. The issue of data attrition in repeat-sales methodologies is well documented. For example, Hill and Trojanek (2022) note that it is common to discard two-thirds or more of the available data when constructing repeat-sales indices. Moreover, they highlight concerns such as lemons bias, where lower quality units are more likely to transact, and fast transactions bias, where frequently traded units differ systematically from others. While these concerns are highly relevant in a housing sales context, they are arguably less applicable to our rental context. This is due to Ireland's legal requirement for landlords to register rental tenancies annually, which provides regular, consistent coverage of Ireland's rental stock.

To assess the representativeness of our RRI sample, Table 1 compares key property characteristics between our full hedonic and RRI samples. We conduct t-tests to evaluate whether differences in sample means are statistically significant differences. As is typical with large datasets, some differences are statistically significant; however, the magnitudes are small. For example, the raw monthly rents only differ by only €7. In terms of geographic distribution, Dublin accounts for a slightly larger share of the RRI sample (45.0 vs 43.9%), while the proportions for other urban areas are nearly identical. Similarly, differences in property size and type are negligible. That the hedonic and RRI samples are so similar is reassuring and provides confidence that our RRI estimates will provide representative estimates of rental inflation.

6 Ireland has 166 LEAs; a sub-county administrative division used for statistical and electoral purposes.

7 Where fewer than 100 observations exist in an LEA for a given quarter, outliers are identified using the 1st and 99th percentiles of the corresponding county-level rent distribution.

Table 1: Comparison of Hedonic versus RRI samples

	Hedonic	RRI	Diff
Monthly rent (€)	1455	1462	−7.011***
lnRent	7.175	7.184	−0.009***
New tenancy dummy	0.258	0.237	0.021***
Number of tenants (max 6)	1.801	1.801	0.000
Dummy: 1 if more tenants than bedrooms	0.132	0.133	−0.001*
One-bedroom property	0.190	0.192	−0.002**
Two-bedrooms	0.379	0.387	−0.008***
Three-bedrooms	0.297	0.294	0.003***
Four-bedrooms	0.113	0.109	0.004***
Five-or more bedrooms	0.021	0.019	0.002***
Detached property	0.104	0.098	0.007***
Semi-detached	0.217	0.213	0.004***
Terraced	0.144	0.143	0.001
Apartment	0.476	0.489	−0.013***
Other property type	0.059	0.057	0.002***
Dublin	0.439	0.450	−0.011***
Greater Dublin Area (GDA)	0.077	0.076	0.001*
Other cities	0.143	0.143	0.000
Other areas	0.341	0.331	0.010***
Number of observations	724,494	590,017	

Notes: Final column provides a t-test to gauge statistically significant differences in means between hedonic and RRI samples. Significance levels: *** p<0.01, ** p<0.05, * p<0.1 .

5.2 Hedonic versus Repeat Rent Index estimates

Figure 2 compares our newly constructed repeat-rent index (RRI) with its hedonic counterpart. Both indices measure rental growth across all tenancies, encompassing both newly initiated leases and ongoing agreements[8]. As the analysis is based on near population-level data, neither index is weighted; for robustness we revisit this issue in Section 5.4. Both indices have a base quarter of Q2 2022, yet by the end of the sample period in Q1 2025, the hedonic index is notably higher than the RRI (112.4 vs. 107.1). Panel B presents annual inflation rates, where RRI estimates range from 1.7-2.5% depending on the quarter, consistently lower than the hedonic estimates of 3.8-5.1%. Across the sample period, inflation estimates derived from

[8] There is much debate about whether indices should capture all tenancies or solely new tenancies to better capture new market conditions. See Adams et al. (2024) for further discussion and how this choice affects estimates of US rental inflation. Separate hedonic measures for new and existing tenancies are already produced in Ireland (see Figure 1). Our focus is instead on understanding the differences in inflation estimates from hedonic and RRI methodologies.

the RRI are between 2.1 and 2.8 percentage points below those from the hedonic index.

Figure 2: Comparison of Hedonic and RRI estimates

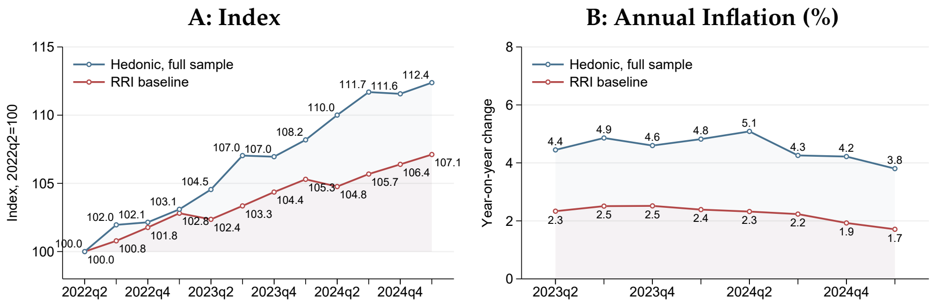

Notes: These are all tenancies indices i.e. include both new and ongoing tenancies.

The hedonic all-tenancies model indicates a broad upward trend until its peak in Q2 2024, before declining in the second half of 2024 and start of 2025. In contrast, annual rental inflation estimates from the RRI model show a consistent downward trend throughout 2024 and into 2025, falling from 2.5% in Q4 2023 to 1.7% by the first quarter of 2025. It is important to note that the hedonic samples change each quarter, incorporating properties that enter or exit the market (or the dataset), whereas the RRI is based on matched pairs and does not capture this market-level turnover. The RRI results in particular appear consistent with HICP inflation falling below 2% from March 2024 onwards, thereby reducing the maximum allowable rent increases in Rent Pressure Zones (RPZ). Given that during our sample period over 80% of tenancies are located within RPZs, this regulatory constraint is likely to have a significant impact on national-level rental inflation.

While the baseline RRI inflation estimates presented in Figure 2 are moderate overall, Figure 3 provides further insight by disaggregating properties into four tenancy-type groups. Focusing on the inflation rates shown in panel B of Figure 3, the majority of properties were ongoing tenancies in both periods, and the inflation rates for these existing-existing pairs closely mirror the baseline RRI estimates. Properties that transitioned from new tenancies in the first period to ongoing tenancies in the second exhibited the lowest rent increases (at or below 1%), indicating that tenants early in their lease are least likely to experience rent adjustments. In contrast, properties with tenant turnover between the two periods saw substantially larger rent increases, particularly those with existing tenants in the first and then new tenants in the second period (in the range of 7-9%). This likely reflects a catch-up effect, where rents had remained unchanged for several years prior to the tenancy change. Properties with new tenants in both periods also experienced elevated rent growth earlier in the sample (4-5%), although these increases converged towards the baseline RRI by the final two quarters in the sample period.

Figure 3: Splitting out RRI estimates by property pair types

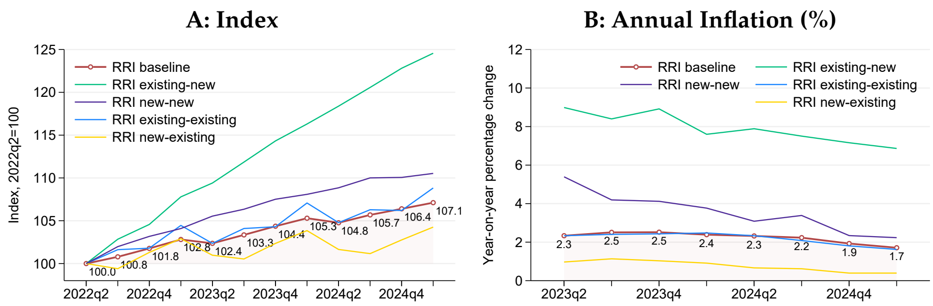

Notes: New-new pairs involve new tenants in both periods. Existing-new pairs reflect a change in tenant in the most recent period, but not in the previous one. New-existing pairs refer to tenants who began their tenancies one year ago and still remain in the property. Existing-existing pairs involve tenants who have remained at the same property for two or more years.

5.3 Explaining the gap

Having established the magnitude of the difference between our hedonic and RRI estimates, we now turn to examining the underlying reasons for this divergence. Two key factors may explain the discrepancy: differences in sample composition, and the impact of property entry and exit from the dataset. We address each in turn.

5.3.1 Sample composition

Our baseline hedonic regression includes the broadest possible set of registrations, excluding only observations identified as outliers based on rent levels. In contrast, three additional cleaning steps are required to construct the quarterly property-level panel dataset required for the RRI index, as detailed in Section 5.1. Although Table 1 showed only minute differences in observable characteristics between the two samples, our objective here is to assess whether, and to what extent, the differing properties in each sample influence the quarterly estimates of rental inflation. To do so, we take each step in turn.

First are observations dropped during the panel cleaning process, including multiple registrations submitted in the same quarter and where matched property pairs have inconsistent property characteristics[9]. To assess whether these observations affect the rental growth rates, we re-estimate the hedonic model by excluding them from the sample. As shown by the yellow line in Figure 4, their removal leads

[9] We removed 51,763 (7%) observations at this stage. The majority were multiple registrations assigned to the same unit in the same quarter. This can occur when a single property ID includes multiple units i.e. rooms, each with separate rental agreements, where leases only last a short time, or where there is a poor match.

to slightly higher growth rates, thus actually increasing the overall gap between the hedonic and RRI estimates.

Figure 4: Comparing Hedonic regression and RRI on a consistent sample

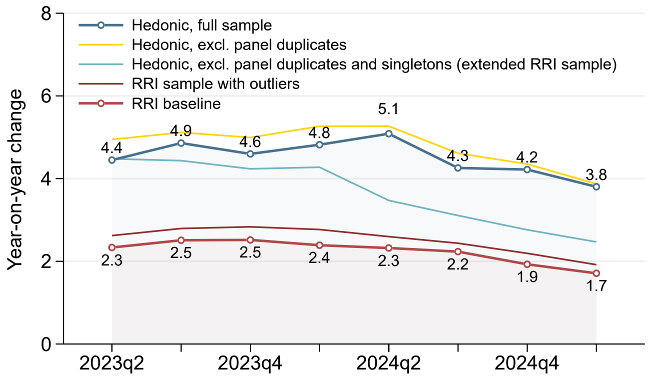

Notes: The underlying sample in the RRI with outliers (dark red) is the same as the extended RRI sample hedonic model (light blue). Thus this represents the gap in hedonic and RRI estimates for an identical sample.

The second group of observations in the hedonic regression, that are excluded from the RRI sample, are those properties only seen once, which we call *singletons*[10]. The combined effect of removing both panel-duplicates and singletons is illustrated by the light blue line in Figure 4. For the first four quarters (2023q2-2024q1), the hedonic estimates on the RRI sample are fairly close to the baseline hedonic model. This suggests the observed gap between the hedonic and RRI indices during this period is primarily attributable to methodological differences rather than sample composition. Notably, these four `'core'` quarters provide sufficient periods both before and after the observation window to reliably track property entries and exits. The role of entries and exits will be explored in greater detail in the next subsection.

In contrast, during the final year of the sample, the hedonic estimates based on the RRI sample move somewhat closer to the RRI baseline estimates. This shift can be explained by the nature of singletons in the last year of the sample; predominantly market entrants that will likely get second observation in the future. Adams et al. (2024) note repeat transaction indices are especially prone to revisions because newly available observations update the estimates for earlier periods. In particular, the most recent period(s) (year in our case, given the annual registration requirement) is particularly volatile, as the sample size for estimating the index is

10 We identify 72,994 singletons, 10% of the total sample. Singletons include both properties registered in the early periods of our sample which have since exited (either the market, or just the sample in the case of missed or late registrations), and properties towards the end of our sample period that have entered the market (or just sample if previously unregistered).

smallest at time t and only increases for later periods as additional repeat pairs become available.

The last remaining sample difference between the hedonic and RRI occurs because the baseline RRI model excludes outliers in terms of growth rates. These are observations where the rent *levels* are not outliers and are therefore included in the hedonic model sample, but have a *growth rate* in the top or bottom 1%, and are thus excluded from the baseline RRI sample. To assess their impact, we re-estimate the RRI model with the growth-outliers included. As shown by the dark red line in Figure 4, including growth outliers results in slightly higher rental inflation estimates, which are on average 0.3 percentage points higher. This reflects the skewed distribution of rent changes, with a longer tail on the high-growth side. Due to this asymmetry, including outliers leads to higher mean estimates compared to the RRI baseline.

Taken all together, Figure 4 shows in our baseline models, the average gap between RRI and hedonic estimates is 2.24 percentage points in the four core quarters[11]. After all three sample adjustments which lead to an identical underlying sample for both methods, the gap reduces by 0.65 percentage points to 1.60. Thus, the sample composition can explain some of the gap, but 71% of the gap still remains even where the RRI and hedonic models are estimated on an identical sample.

5.3.2 Entry and exit

In Section 5.3.1, we find that for the core quarters (2023q2–2024q1), hedonic estimates on the RRI sample closely match those from the baseline hedonic model. This indicates that the gap between the hedonic and RRI indices during this period is mainly due to methodological differences rather than sample composition. We now explore how property entry and exit patterns contribute to differences in rental inflation estimates from the two models.

For the purpose of this analysis, an 'entry' is defined as a panel entrant i.e. it has entered into our RRI sample, or when a property is first observed in the registrations dataset (excluding any properties seen in the first year of the dataset) and is then seen again at least once. A panel 'exit' is defined if the property was in the data at least twice, but has not been seen subsequently even though at least four quarters have since passed. The observations seen only once ('singletons') are treated as a separate category. In practice, as we saw in section 5.3.1, many singletons seen in the last year of our sample will likely be entrants in time, but the earliest a matched pair would be seen is 2025q2, so they do not yet meet our definition of a panel entrant and we therefore treat them as a separate category. Similarly, we can only determine what is an actual exit ex-post. In other words, for observations in the last four quarters we do not know yet if they will appear again or not. For this reason we focus solely on the core quarters (2023q2–2024q1).

[11] Note this only refers to the four core quarters as the final year of our sample is affected by the conflation of singletons and entrants.

It is important to clarify here that entries and exits refer to transitions within the dataset, not necessarily the rental market itself. For instance, apparent exits may not be actual exits from the rental market, they could instead represent late registrations, or a temporarily vacant property that is available to rent. Additionally, there could be some cases where there is an issue reliably tracking the same property over time. In which case, the same property would first appear as an exit, and then reappear as an entry under a new identification number. Given the extensive property ID, landlord ID and address matching process, these cases are likely to be minimal.

Table 2 shows the number of observations depending on the observations' position inside the panel. In the four core quarters, there are 38,707 entrants (first time the property is seen in the panel), 25,363 exits (the last observation in the panel that did not have any registrations afterwards), and another 12,572 singleton properties that were seen neither before nor later. In this core period, the entries account for 16% of all observations, and exits for 10%.

Table 2: Number of observations by position in the panel

		Position in the panel					
	Singleton	First (entry)	Middle	Last (exit)	Other	Total	
2022q2	7,201	36,599	0	0	1,638	45,438	
2022q3	9,866	53,736	326	31	2,671	66,630	First
2022q4	7,977	44,392	965	92	2,120	55,546	year
2023q1	8,322	44,040	2,521	243	2,293	57,419	
2023q2	3,593	12,249	30,592	4,783	5,353	56,570	
2023q3	3,051	10,486	45,206	6,911	6,053	71,707	Core
2023q4	2,571	7,961	37,752	6,201	4,637	59,122	quarters
2024q1	3,357	8,011	38,087	7,468	4,957	61,880	
2024q2	7,238	463	3,820	41,695	5,893	59,109	
2024q3	7,863	273	2,796	55,185	6,754	72,871	Last
2024q4	6,086	101	1,050	46,798	4,958	58,993	year
2025q1	5,869	0	0	48,904	4,436	59,209	
Total	72,994	218,311	163,115	218,311	51,763	724,494	
Core total	12,572	38,707	151,637	25,363	21,000	249,279	
Core share	5%	16%	61%	10%	8%	100%	

Notes: Singletons are properties with only one observation in the sample period. Middle observations have both a registration prior and another registration afterwards. Category 'Other' are observations included in the hedonic regressions, but dropped in the panel cleaning process, such as multiple registrations in the same quarter.

Next we want to know if the entrants and exits are different from other observations and each other regarding rent levels. Figure 5 shows that entrants have on average 7.48% higher standardised average rent levels compared to observations that have a registration both before and after (so called middle observations). Meanwhile, the exiting properties have on average 3.56% lower standardised average rents. Combining these two differences, that is 11% higher rents for entrants compared to contemporaneous exits. Rent levels of singletons in the core period are 1.5% lower than middle observations. Note these singletons are quite different from the ones seen in the last year of the data (and therefore only in the last year of the hedonic sample, from 2024q2-2025q1)[12].

Figure 5: Relative rent levels of entrants and exits to RRI panel

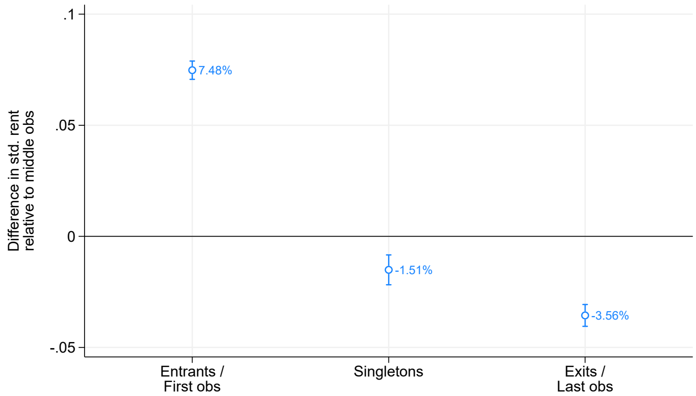

Notes: Observations in core quarters (2023q2-2024q1). Bars show 99% confidence interval. Singletons in the core period see similar rent levels to properties that have a registration both before and after (so called middle observations). Note these singletons are quite different from the ones seen in the last year of the data (and therefore only in the last year of the hedonic sample, from 2024q2-2025q1).

In Figure 6 the comparison between entrants and exits is extended beyond the mean. It shows the kernel density estimates of the standardised rent for both entries and exits in the four core quarters. The entire distribution of entrants is shifted rightwards relative to the exiting properties. Therefore, the difference in the mean seen in Figure 5 is not driven by a small number of outliers, but rather by systematically higher rent levels for entrants compared to those of exiting properties.

12 Singletons in the core period can be thought of as systematic errors. In contrast, singletons in the final year of our sample will be a mix of systematic errors and market entrants included in the hedonic sample, but not yet seen twice, so not yet in the RRI sample.

Figure 6: Comparison of standardised average rent levels of panel entrants and exits

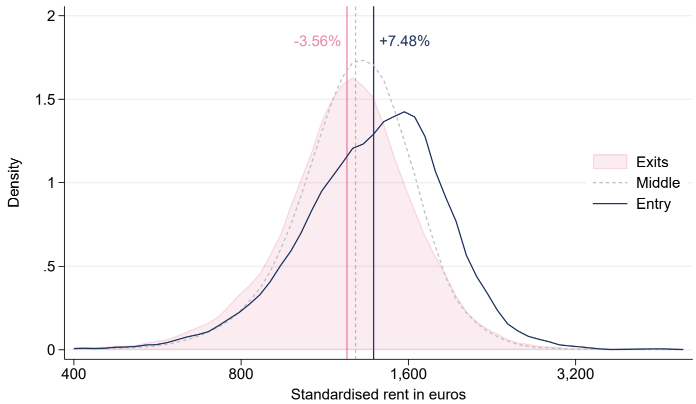

Notes: Observations in core quarters (2023q2-2024q1).

5.3.3 Decomposition

Having established that panel entrants and exits represent a notable share of total observations and that their rent levels are significantly different than other observations, the final step is to quantify how these two facts explain the gap between the hedonic and RRI model estimates of rental inflation. This is achieved by decomposing the change in the hedonic model into a component representing a RRI-like change and other components that explain the difference (entry and exit, outliers, and sample composition). Again, we focus on the four core quarters in the middle of our sample (2023q2-2024q1) for which we can more reliably track entrants and exits. This decomposition is a modified and extended version of Melitz and Polanec (2015) where it was used to decompose firms' aggregate productivity.

Estimates from the hedonic model (γ_t) for a period t can be expressed as a geometric mean of all standardised rents $\hat{P}_i$ in that period. For any two period pair, the sum of log-rents can then be split into two groups: properties that are overlapping in both periods (θ) and properties that are either new entrants (W_1) or exits (X_0). The total number of observations in the base period is $T_0 = \theta_0 + X_0$ and in the second period $T_1 = \theta_1 + W_1$, where by definition $\theta_0 = \theta_1$.

$$\gamma_1 = \frac{1}{T_1} \sum_{i}^{T_1} \ln \hat{P}_{i1} = \frac{\theta_1}{T_1} \frac{1}{\theta_1} \sum_{i}^{\theta_1} \ln \hat{P}_{i1} + \frac{W_1}{T_1} \frac{1}{W_1} \sum_{i}^{W_1} \ln \hat{P}_{i1} \tag{6}$$

$$\gamma_0 = \frac{1}{T_1} \sum_{i}^{T_1} \ln \hat{P}_{i0} = \frac{\theta_0}{T_0} \frac{1}{\theta_0} \sum_{i}^{O_0} \ln \hat{P}_{i0} + \frac{X_0}{T_0} \frac{1}{X_0} \sum_{i}^{X_0} \ln \hat{P}_{i0} \tag{7}$$

The two expressions can be rewritten as the weighted mean of the geometric mean rent in the overlapping sample ($\overline{\ln P_\theta}$) and the geometric mean of either entry or exit observations ($\overline{\ln P_W}$ and $\overline{\ln P_X}$ respectively):

$$\gamma_1 = s_{\theta 1}\overline{\ln P}_{\theta 1} + s_{W1}\overline{\ln P}_{W1} = \left(1 - s_{W1}\right)\overline{\ln P}_{\theta 1} + s_{W1}\overline{\ln P}_{W1} \tag{8}$$

$$\gamma_0 = s_{\theta 0}\overline{\ln P}_{\theta 0} + s_{X0}\overline{\ln P}_{X0} = \left(1 - s_{X0}\right)\overline{\ln P}_{\theta 0} + s_{X0}\overline{\ln P}_{X0} \tag{9}$$

After some further algebraic simplification, taking the difference between the two equations leads to the following decomposition:

$$\gamma_1 = \overline{\ln P}_{\theta 1} + s_{W1}\left(\overline{\ln P}_{W1} - \overline{\ln P}_{\theta 1}\right) \tag{10}$$

$$\gamma_0 = \overline{\ln P}_{\theta 0} + s_{X0}\left(\overline{\ln P}_{X0} - \overline{\ln P}_{\theta 0}\right) \tag{11}$$

$$\underbrace{\gamma_1 - \gamma_0}_{\text{Change in hedonic}} = \underbrace{\overline{\Delta \ln P_\theta}}_{\substack{\text{Av. growth}\\ \sim\text{RRI}}} + \underbrace{s_{W1}\left(\overline{\ln P}_{W1} - \overline{\ln P}_{\theta 1}\right)}_{\text{Entry}} - \underbrace{s_{X0}\left(\overline{\ln P}_{X0} - \overline{\ln P}_{\theta 0}\right)}_{\text{Exit}} \tag{12}$$

Therefore, the change in hedonic estimates can be decomposed into[13]:

(1) Average within-property growth rate in the matched property pair (or overlapping) sample, which approximately corresponds to the RRI estimates[14]

(2) Entry, calculated as the share of properties entering and their average rent relative to the rent of the overlapping sample in the same period

(3) Exit, calculated as a share of properties exiting and their average rent in previous period relative to the rent of the overlapping sample in the previous period

(4) Not shown in the derivation, but the decomposition can be extended further to separate outliers and other observations not included in the hedonic regression.

13 There are a couple of caveats before we can bring this decomposition to the data. Typically, when working with quarterly data one would look at the year-on-year change between γ_t and γ_{t-4}. Even though the majority of observations are observed in the same quarter every year, many new tenancies start in different quarter than the previous tenancy in the same property. As turnover tenancies have higher growth rates, excluding them would lead to downward bias. When we report the average growth rate, it will be that of the closest quarter. The second issue is that currently there are only four quarters with reliable entry and exit data. Ideally, the decomposition requires entry at time t and exit information from $t-4$. This cannot be done with the current data. As an approximation we will use contemporaneous exits, both in terms of share of exits and the average rent level.

14 Because our data is reported annually these two are very similar, but this does not hold for the general case.

Table 3: Decomposition of differences between hedonic-model growth rate and average property-level growth rate (~RRI)

		Decomposition of difference				Entry		Exit	
Period	Hedonic	Overlap	Net entry	Outliers	Resid.	Share	Rel. rent	Share	Rel. rent
2023q2	4.45	2.31	1.53	0.17	0.44	0.22	5.50	0.08	-3.99
2023q3	4.86	2.51	1.73	0.21	0.41	0.15	9.21	0.10	-3.99
2023q4	4.60	2.54	1.68	0.24	0.14	0.13	9.20	0.10	-4.24
2024q1	4.82	2.52	1.35	0.26	0.69	0.13	7.95	0.12	-2.63
Total	4.68	2.48	1.52	0.22	0.46	0.16	7.48	0.10	-3.56
%Total	100%	53%	33%	5%	10%				

Notes: Estimates for the 'core' quarters from 2023q2 to 2024q1 only. Overlap refers to matched property pair growth rates i.e RRI. Net entry is the combined entry and exit effect. The bottom row shows the percentage of the hedonic mean growth rate that can be explained by each component.

Table 3 shows these decomposition results for the total and for each of the four quarters separately. All four quarters give similar results so we focus only on the total. Across the four core quarters (2023q2-2024q1) the overall rental inflation rate based on the hedonic model was 4.68%. In the final row, the decomposition is expressed as a percentage of the total hedonic growth rate. It shows that 53% of the hedonic rental growth can be attributed to within-property growth observed in the overlapping sample i.e. the RRI, 33% can be attributed to net entry (25% entry, 8% exit[15]), and the remaining 15% is due to combined outliers, sample composition and approximation errors. This decomposition explains why our hedonic estimates are higher than our RRI estimates of rental inflation, with the differing rent levels of properties entering and exiting the sample responsible for the majority of the gap. It is important to note this finding relates to the four specific quarters in question and may vary across other time periods.

One of the advantages of this new dataset is its ability to capture local variation in rental prices outside of major urban centres, a notable challenge with survey data. Table 4 presents the same decomposition for each of Ireland's local authority areas. Taking two contrasting markets, the urban, rent-controlled South Dublin LA and the rural, non-rent-controlled Tipperary, illustrates the value of having both hedonic and RRI estimates and the insights provided by the decomposition. While overall hedonic inflation estimates were similar in both areas (5.94 vs 5.83%), the decomposition shows fundamentally different drivers. In South Dublin, only 1.8 percentage points reflected within-property rent growth, consistent with the maximum permitted annual rent increases of 2% under the rent controls. Most of the inflation was instead driven by net entry (3.3%), consistent with the large-scale

15 16% of observations are panel entrants with an average rent 7.48% above those in the overlapping sample. These entrants explain 1.16 percentage points of the hedonic growth rate. 11% are panel exits with an average rent 3.56% below the overlapping sample. These exits explain 0.36 percentage points of the hedonic growth rate.

new build-to-rent developments in the area. In contrast, in Tipperary, nearly all inflation (4.7 percentage points) came from within-property growth, indicating worse affordability developments, but limited churn and fewer supply shocks. Importantly, the impact of entry and exit on the hedonic estimates is not confined to new-build properties; it reflects broader market churn. The final and third to last columns in Table 4 show that across almost all local authorities, entrants consistently have higher rents than other properties, while exiting properties have lower rents. Having both indices provides deeper insights into local trends and drivers of the inflation, which is crucial for policymakers when designing and assessing policy interventions.

Table 4: Decomposition of differences between hedonic-model growth rate and average property-level growth rate - by Local Authority

			Decomposition of difference			Entry		Exit	
Local authority	Hedonic	Overlap	Net entry	Outliers	Resid.	Share	Rel. rent	Share	Rel. rent
Carlow	3.73	3.18	0.71	0.15	-0.31	0.12	5.43	0.11	-0.40
Cavan	9.96	5.38	1.77	1.64	1.16	0.19	8.79	0.10	-1.32
Clare	6.45	4.68	0.93	0.72	0.12	0.16	3.64	0.10	-3.24
Cork Co.	3.98	2.25	0.91	0.31	0.53	0.16	5.85	0.11	0.08
Cork City	3.54	1.94	1.17	-0.08	0.51	0.13	5.86	0.10	-3.92
Donegal	9.58	5.47	1.59	1.55	0.96	0.17	5.08	0.10	-7.33
Galway Co.	6.56	3.86	1.38	0.57	0.75	0.16	6.83	0.12	-2.32
Galway City	3.86	2.17	1.61	-0.03	0.11	0.12	10.37	0.10	-3.04
Kerry	6.12	4.59	0.74	0.67	0.12	0.15	-1.05	0.09	-9.49
Kildare	4.61	1.94	1.61	-0.13	1.20	0.16	9.61	0.12	-0.87
Kilkenny	4.55	2.27	1.12	0.06	1.10	0.15	6.30	0.11	-1.53
Laois	4.95	2.63	1.62	0.28	0.42	0.15	5.68	0.14	-5.23
Leitrim	6.93	5.19	0.16	1.27	0.31	0.16	2.51	0.13	1.93
Limerick Co.	7.07	3.56	2.02	0.87	0.62	0.17	10.14	0.12	-2.71
Limerick City	4.69	2.06	2.05	0.18	0.40	0.14	11.88	0.12	-3.57
Longford	8.87	5.83	0.81	1.82	0.42	0.14	0.69	0.12	-5.98
Louth	4.47	2.01	1.52	0.15	0.80	0.16	8.95	0.11	-0.75
Mayo	7.10	5.24	0.49	0.71	0.66	0.16	2.34	0.10	-1.00
Meath	3.71	1.96	1.46	0.01	0.29	0.16	5.88	0.12	-4.49
Monaghan	6.05	4.19	0.97	0.44	0.46	0.17	4.37	0.10	-2.20
Offaly	4.94	3.40	1.54	0.40	-0.41	0.16	7.55	0.14	-2.39
Roscommon	7.61	5.24	0.74	0.95	0.68	0.16	3.06	0.12	-2.14
Sligo	5.79	3.58	1.01	0.65	0.54	0.14	5.22	0.11	-2.77
Tipperary	5.83	4.69	1.19	0.29	-0.34	0.16	4.71	0.11	-3.99
Waterford Co.	5.75	3.64	2.03	0.32	-0.25	0.18	8.74	0.12	-4.25
Waterford City	4.36	2.66	0.94	0.46	0.30	0.13	5.79	0.11	-1.52
Westmeath	5.10	2.55	1.87	0.50	0.18	0.15	7.89	0.12	-5.43
Wexford	7.57	4.44	1.40	0.74	1.00	0.16	6.52	0.11	-3.65
Wicklow	3.97	1.87	1.44	0.06	0.60	0.16	9.37	0.13	0.25
Fingal	3.33	1.77	1.53	-0.03	0.05	0.14	6.13	0.10	-6.69
Dún Lao.-Rath.	4.29	1.63	2.19	-0.03	0.49	0.18	10.81	0.08	-3.70
South Dublin	5.94	1.77	3.29	-0.01	0.89	0.19	13.95	0.08	-8.58
Dublin City	3.52	1.72	1.37	0.02	0.41	0.16	6.33	0.09	-4.18
Total	4.68	2.48	1.52	0.22	0.46	0.16	7.48	0.10	-3.56

Notes: Overlap refers to matched property pair growth rates i.e RRI. Net entry is the combined entry and exit effect. Estimates for the 'core' quarters from 2023q2 to 2024q1 only.

5.4 Robustness

In this section we perform a number of robustness checks. As the legally mandated annual registration of private rental tenancies in Ireland should provide near-complete coverage of the PRS in our dataset, we estimate our baseline hedonic and RRI specifications without weights. However, in practice, incomplete compliance with these registration requirements could introduce geographic or structural biases, potentially affecting the representativeness of our estimates. To assess robustness, we therefore re-weight our baseline hedonic index using Local Electoral Area (LEA) weights from the 2022 Census, which reflect the spatial distribution of the rental stock. We would not expect the Census and RTB samples to align perfectly, as the Census captures a single point in time and also includes some informal rental arrangements, such as family agreements, that are not required to be registered with the regulator. Moreover, as Hill (2013) notes, it is not clear whether re-weighting a hedonic index to reflect the entire rental stock, rather than observed transactions each period, is conceptually desirable. Nonetheless, this exercise is a useful test of our dataset's representativeness. Appendix Figure 7 panel A shows that re-weighting does not make a material difference to our estimates. The adjustment fractionally raises the index upwards as some smaller, less urban rental markets are slightly under-represented, but the difference is negligible.

A potential limitation of the RRI approach, as highlighted in section 4, concerns sample representativeness, since only properties observed more than once are retained. While Table 1 indicates that, on average, the hedonic and RRI samples are very similar, differences may arise in specific quarters. To address this, we re-weight the RRI sample to match the full hedonic sample of all tenancies. Specifically, for each quarter, we estimate a logit model where the dependent variable equals one if a property is included in the RRI sample, and regress this on the hedonic characteristics. We then use the inverse predicted probabilities to construct weights for re-estimating the RRI index. Panel B of Figure 7 shows that this adjustment has no material impact on the estimates[16].

As discussed in section 4, another concern with the RRI model we estimate is that the assumption of constant variance may be violated (Case & Shiller, 1989). In other words, there is a possibility that observations with larger time gaps gaps between them will have larger variance. As each property should be observed annually in our dataset, or more frequently if turnover occurs, we would not expect this to be a major concern. Nevertheless, as a robustness check, we perform a three stage feasible generalised least squares (FGLS) to account for this. Like Clark (2022), Adams et al. (2024) and Abramson et al. (2025), we do not find this heteroskedasticity correction to have a meaningful effect on our RRI estimates (Figure 8).

16 Note the differences between our hedonic and RRI estimates attributed to sample composition in section 5.3.1 are due to differences in rent levels between the singletons included in the hedonic, but not yet the RRI sample, not due to differences in the hedonic characteristics we use here to re-weight.

In section 5.3.2 we showed that churn and the differing rent levels of properties entering and exiting the sample is a key driver of the divergence between our hedonic index and RRI estimates. One possible explanation for this, could be that entrants' high rents are due to their higher quality, which due to omitted variable bias is not captured in the hedonic regression. One way to mitigate this is to include dummies for Building Energy Rating (BER) score. Though legally mandatory for most rental units, the BER is not a required field in the registration form and 48.6% of the registrations do not provide the information on energy efficiency.[17] Because new-build properties in Ireland are required to have an A1 or A2 rating, the BER rating would capture not just energy efficiency but also act as a proxy of property quality and age.

To assess the robustness of our findings, we therefore re-estimate the hedonic model, the RRI, and the decomposition using only the subsample of 372,852 registrations (198,560 pairs) that have a non-missing BER rating. Results in Table 5 in the Appendix show that in this subsample with BER controls, the overall growth rate was 4.14%, compared to 4.68% in the baseline model (Table 3). Within-property changes account for 59% of the average hedonic-based growth (up from 53% in our baseline estimates), while the net entry accounts for 27% (down from 33%). The difference in standardised rent levels between entrants and exits is now 8% (compared to 11% in the baseline results). Overall, these results indicate there was some increase in quality due to new-build high-efficiency properties. Nevertheless, even after more precise measurement of quality differences, the substantial gap between the RRI and hedonic estimates of rental inflation remains, primarily driven by entry and exit dynamics outlined in section 5.3.3.

6 Conclusion

In this paper we have constructed a repeat rent index (RRI) using newly available administrative data from Ireland's rental regulator, providing near population-level coverage of both new and ongoing tenancies, to complement the existing hedonic approach. Applying both methods to the same dataset enables us to isolate methodological effects from data source differences and to assess how index choice influences rental inflation estimates. A formal decomposition of hedonic estimates into within-property growth and entry and exit effects reveals additional insights into the differing underlying drivers of rental inflation.

Several key findings emerge. Nationally, annual rental inflation measured by the repeat rent index (RRI) is consistently 2.1–2.8 percentage points lower than hedonic estimates over the period Q2 2022–Q1 2025. Our decomposition shows that within-property rent growth accounts for 53% of hedonic inflation, while entry and exit effects, contribute roughly one-third. Properties entering the market exhibit standardised average rents around 11% higher than those exiting, highlighting

[17] (Kren et al., 2025) showed that there is some self-selection in reporting of BER, therefore this sample is not representative of the entire Irish rental market.

the role of property turnover in shaping hedonic estimates. Importantly, these differences are not explained by sample selection bias, as the hedonic and RRI samples are very similar in observable characteristics. At the local level, RRI inflation was low in rent-controlled urban markets compared to hedonic figures, which often exceeded rent cap thresholds. In contrast, non-urban areas without controls showed higher within-property growth and limited churn. These findings illustrate the need to distinguish market turnover effects from affordability pressures on sitting tenants in local markets.

Our findings have important implications for both measurement and policy. While each index captures distinct aspects of rental sector dynamics, their combined use offers a more comprehensive understanding of market conditions and affordability developments. The hedonic index is best suited to capturing market-level dynamics and the influence of churn on average rent levels. This reflects broader market pressures such as strong population growth, supply shocks, elevated construction costs, and high interest rates. In contrast, the RRI better represents household affordability trends and compliance with rent caps by isolating within-property changes. This distinction is particularly critical in regulated markets where policymakers are trying to balance short-term evaluations of rent control regulations with longer-term goals of increasing housing supply and ensuring the sustainability of the private rental sector. Although shelter costs account for a smaller share of Ireland's CPI than in countries like the US, the divergence between hedonic and RRI estimates illustrates the importance of index selection for accurate inflation measurement and cost-of-living assessments. Further research on these implications in the Irish context would be valuable.

Finally, this research highlights the value of comprehensive administrative data for improving rental market measurement, enabling more nuanced insights which are critical for researchers and policymakers seeking to understand complex rental market dynamics and design effective policy interventions. While our findings provide important contributions, several limitations should be noted. First, the analysis is based on a relatively short time frame, which constrains our ability to examine longer-term trends and the persistence of entry and exit effects on inflation estimates. Second, the period covered coincided with significant structural developments, including the delivery of large-scale build-to-rent schemes in parts of Dublin, which may have amplified compositional effects. Future research should explore how these dynamics evolve over a longer time horizon.

References

Abramitzky, R., Boustan, L. P., & Storeygard, A. (2025, March). *New data and insights in regional and urban economics* (NBER Working Paper No. 33561). National Bureau of Economic Research. https://www.nber.org/system/files/working_papers/w33561/w33561.pdf

Abramson, B., Han, L., & De Llanos, P. (2025). Monetary policy and rents. *SSRN 5353365*. https://doi.org/10.2139/ssrn.5353365

Adams, B., Loewenstein, L., Montag, H., & Verbrugge, R. (2024). Disentangling rent index differences: Data, methods, and scope. *American Economic Review: Insights*, *6*(2), 230–245. https://doi.org/10.1257/aeri.20220685

Adams, B., & Verbrugge, R. (2025). Location, location, structure type: Rent divergence within neighborhoods. *Journal of Housing Economics*, 102081. https://doi.org/10.1016/j.jhe.2025.102081

Ambrose, B. W., Coulson, N. E., & Yoshida, J. (2015). The repeat rent index. *Review of Economics and Statistics*, *97*(5), 939–950. https://doi.org/10.1162/REST_a_00500

Ambrose, B. W., Coulson, N. E., & Yoshida, J. (2018). Reassessing Taylor rules using improved housing rent data. *Journal of Macroeconomics*, *56*, 243–257. https://doi.org/10.1016/j.jmacro.2018.03.001

Ambrose, B. W., Coulson, N. E., & Yoshida, J. (2023). Housing rents and inflation rates. *Journal of Money, Credit and Banking*, *55*(4), 975–992. https://doi.org/10.1111/jmcb.12971

Bailey, M. J., Muth, R. F., & Nourse, H. O. (1963). A regression method for real estate price index construction. *Journal of the American Statistical Association*, *58*(304), 933–942.

Boesel, M., Chen, S., & Nothaft, F. E. (2021). Housing preferences during the pandemic: Effect on home price, rent, and inflation measurement. *Business Economics (Cleveland, Ohio)*, *56*(4), 200. https://doi.org/10.1057/s11369-021-00241-4

Bricongne, J.-C., Meunier, B., & Pouget, S. (2023). Web-scraping housing prices in real-time: The covid-19 crisis in the uk. *Journal of Housing Economics*, *59*, 101906. https://doi.org/10.1016/j.jhe.2022.101906

Case, K. E., & Shiller, R. J. (1989). The efficiency of the market for single-family homes. *The American Economic Review*, 125–137.

Clark, J. (2022). Methodology: Zillow observed rent index (ZORI). https://www.zillow.com/research/methodology-zori-repeat-rent-27092/

Coffey, C., McQuinn, K., & O'Toole, C. (2022). Rental equivalence, owner-occupied housing, and inflation measurement: Microlevel evidence from Ireland. *Real Estate Economics*, *50*(4), 990–1021. https://doi.org/10.1111/1540-6229.12360

Contat, J., & Larson, W. D. (2024). A flexible method of housing price index construction using repeat-sales aggregates. *Real Estate Economics*, *52*(6), 1551–1583. https://doi.org/10.1111/1540-6229.12474

Corrigan, E., Foley, D., McQuinn, K., O'Toole, C., & Slaymaker, R. (2019). Exploring affordability in the irish housing market. *The Economic and Social Review*, *50*(1, Spring), 119–157. https://www.esr.ie/index.php/esr/article/view/1121

Crone, T. M., Nakamura, L. I., & Voith, R. (2010). Rents have been rising, not falling, in the postwar period. *The Review of Economics and Statistics*, *92*(3), 628–642. https://doi.org/10.1162/REST_a_00015

De Haan, J. (2013). Repeat sales methods. In E. Name (Ed.), *Handbook on residential property prices indices (rppis)* (pp. 65–71). Eurostat.

De Haan, J., & Diewert, E. (2013). Hedonic regression methods. In E. Name (Ed.), *Handbook on residential property prices indices (rppis)* (pp. 49–64). Eurostat.

Disch, W., & Slaymaker, R. (2023). Housing affordability: Ireland in a cross-country context. *ESRI Research Series 164*. https://doi.org/10.26504/rs164

Gillespie, T., Kren, J., Lyons, R. C., & O'Toole, C. (2025). The supply side effects of rent controls: Evidence from Ireland. *Journal of Housing Economics*, 102099. https://doi.org/10.1016/j.jhe.2025.102099

Gindelsky, M., Moulton, J., & Wentland, S. A. (2021). Valuing housing services in the era of big data: A user cost approach leveraging zillow microdata" in" big data for 21st century economic statistics. *University of Chicago Press, Chicago, IL*.

Hill, R. J. (2013). Hedonic price indexes for residential housing: A survey, evaluation and taxonomy. *Journal of Economic Surveys*, *27*(5), 879–914. https://doi.org/10.1111/j.1467-6419.2012.00731.x

Hill, R. J., & Trojanek, R. (2022). An evaluation of competing methods for constructing house price indexes: The case of Warsaw. *Land Use Policy*, *120*, 106226. https://doi.org/10.1016/j.landusepol.2022.106226

Kholodilin, K. A. (2024). Rent control effects through the lens of empirical research: An almost complete review of the literature. *Journal of Housing Economics*, 101983. https://doi.org/10.1016/j.jhe.2024.101983

Kim, K. (2024). The repeat rent model in the presence of variable deposits. *Applied Economics Letters*, 1–8. https://doi.org/10.1080/13504851.2024.2334927

Kren, J., Kenny, E., O'Toole, C., Shiel, E., & Slaymaker, R. (2025). Exploring investment requirements for energy efficiency upgrades in the private rental sector. *ESRI Research Series 205*. https://doi.org/10.26504/RS205

Lawless, M., McQuinn, K., & Walsh, J. (2018). Identifying rent pressures in your neighbourhood: A new model of Irish regional rent indicators. *The Economic and Social Review*, *49*(1, Spring), 73–92. https://www.esr.ie/index.php/esr/article/view/860

Malpezzi, S., et al. (2003). Hedonic pricing models: A selective and applied review. *Housing economics and public policy*, *1*, 67–89.

Melitz, M. J., & Polanec, S. (2015). Dynamic Olley-Pakes productivity decomposition with entry and exit. *The Rand Journal of Economics*, *46*(2), 362–375. https://doi.org/10.1111/1756-2171.12088

O'Toole, C., Martinez-Cillero, M., & Ahrens, A. (2021). Price regulation, inflation, and nominal rigidity in housing rents. *Journal of Housing Economics*, *52*, 101769. https://doi.org/10.1016/j.jhe.2021.101769

Slaymaker, R., Kren, J., & Devane, K. (2024). An assessment of property level rental price growth in Ireland. *Jointly-published Reports 10*. https://doi.org/10.26504/jr10

Slaymaker, R., Roantree, B., Nolan, A., & O'Toole, C. (2022). Future trends in housing tenure and the adequacy of retirement income. *ESRI Research Series 143*. https://doi.org/10.26504/rs143

Spiegel, S.-S. (2025). Price indices for Austrian municipalities-hedonic models based on microlevel data. *The Journal of Real Estate Finance and Economics*, 1–44. https://doi.org/10.1007/s11146-025-10032-9

A Appendix: Robustness checks

Figure 7: Robustness check: re-weighting

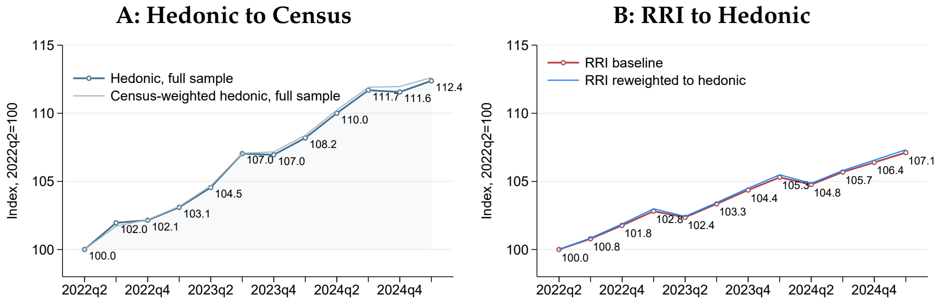

Figure 8: Comparison of RRI baseline and RRI with heteroskedasticity adjustment estimates

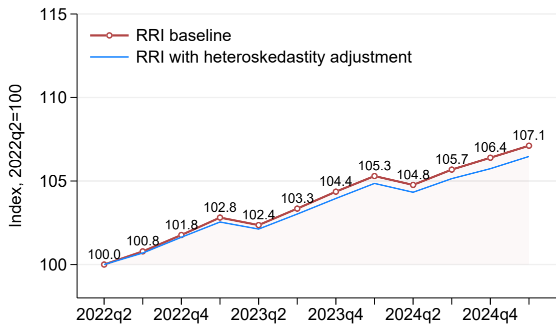

Table 5: Decomposition of differences with energy efficiency controls in hedonic regression using only the subsample of registrations

		Decomposition of difference					Entry		Exit	
Period	Hedonic	Overlap	Net entry	Outliers	Resid.		Share	Rel. rent	Share	Rel. rent
2023q2	4.07	2.27	1.18	0.24	0.38		0.21	4.59	0.08	-2.9
2023q3	4.42	2.55	1.25	0.23	0.39		0.14	6.93	0.09	-3.04
2023q4	3.81	2.49	1.16	0.24	-0.08		0.13	6.87	0.09	-2.81
2024q1	4.25	2.48	1.03	0.28	0.46		0.12	6.83	0.11	-2.03
Total	4.14	2.46	1.11	0.25	0.32		0.15	5.87	0.09	-2.58
%Total	100%	59%	27%	6%	8%					

Notes: Subsample of registrations with non-missing BER energy rating.