Banc Ceannais na hÉireann
Central Bank of Ireland
Eurosystem

Research Technical Paper

Institutional Investment and Residential Rental Market Dynamics

Barra McCarthy

Vol. 2024, No. 1

Institutional Investment and Residential Rental Market Dynamics

Barra McCarthy[1]

February 23, 2024

Abstract

The impact of institutional investment on residential rental markets is a topic that draws much attention, but for which limited evidence exists. This paper aims to remedy this by analysing the direct and indirect impact of institutional investment in existing housing stock on residential rents at the property level, using Ireland as a case study. Ireland is particularly suited, as large portfolios of rental properties became available for purchase due to the Global Financial Crisis (GFC), enticing institutional investors into Ireland's rental market for the first time. By focusing on investment in existing stock, the analysis isolates the impact of institutional investment from the effect it can have on local rental markets when it is channelled to new rental housing – a supply effect which should reduce rents. The study finds that institutional investors increased the level of rents by 4.1 percentage points more than other landlords with comparable properties following purchase. Available evidence supports many different causal mechanisms, including bargaining power, institutional landlords providing a higher cost-higher quality service, and institutional landlords placing greater emphasis on profitability.

JEL-Classification: R30, G23
Keywords: Institutional Investment, Residential Rental Markets, Spatial Data

[1] Central Bank of Ireland, International Finance Division; barra.mccarthy@centralbankie. I thank Ronan Lyons for excellent comments on an early draft, Neill Killeen, Kitty Moloney, Peter Dunne, participants at the 2022 Irish Economic Association Conference and the 2022 American Real Estate and Urban Economics International Conference, and Central Bank of Ireland Economic Seminars. Thanks as well to Daniele Bertocchi. This views expressed in this paper are the authors and do not represent the views of the Central Bank of Ireland or European Central Bank.

Non-Technical Summary

Investment by institutional investors (i.e. financial institutions) in residential rental property is either presented as a key driver of rental market dysfunction, or as part of the solution to the problems present in many rental markets around the world. Despite the debate it generates, limited empirical evidence exists on its impact. This paper aims to help inform this debate by answering the following question: does institutional investment in existing (i.e. not new build) residential properties increase rents? By analysing investment in existing stock, the paper isolates the impact of institutional investment from the supply effect such investment can have when it is channelled into construction and purchase of new rental housing (i.e. new rental housing should lead to lower rents, relative to if it had not been constructed).

Institutional investment could increase rents directly if institutional investors charge more than other landlords for similar properties, and/or indirectly, if institutional investors behaviour influences other landlords.

There are many reasons why we might expect institutional investors to increase rents by more than other landlords. Due to their size, a month's lost rent is relatively less costly for them than it is for someone with just one property – allowing them to wait longer for a higher rent. They may also put less value on the things other landlords would, like having a good tenant, and more focus on profitability. Alternatively, they may have different costs to other landlords, which would be the case if they are offering higher quality, higher cost rental service (i.e. spending more money on repairs or services for the building). Finally, they may attract higher wage tenants, who may in turn bid up rents to a greater degree than for other landlords' properties.

To examine whether they do increase rents in excess of other landlords, I analyse Ireland's rental market following the 2007 property market crash. The crash led large portfolios of tenanted properties to transfer from developers and individuals to institutional investors. Individuals and developers ended up needing to sell or transfer these properties to creditors because they were purchased or constructed just before the crash, and they were no longer able to meet debt repayments for them. As these properties were relinquished due to timing, it's reasonable to assume that if they were not purchased by an institutional investor their rents would have followed the trend of similar properties which were not purchased. By comparing the changes in rents of properties purchased by institutional investors against similar properties that were not purchased, we can estimate the direct impact.

To answer the research question I construct a unique dataset from a variety of administrative and public data sources spanning 2007-2019. The core of this dataset is the Residential Tenancy Board's registered tenancy dataset, which provides data for individual properties' rents and characteristics. I construct a property level dataset for institutional investors' holdings using administrative and public data. These two datasets are combined with small area-level (census units of on average less than 100 properties) measures of residential building quality.

The data provide detailed information on the holdings of institutional investors. They held 12,643 units at the end of 2019, with 9,862 of these purchased as existing properties. These existing properties amount to around 4% of all properties with an Eircode in the dataset. The

data show that institutional investors' property differs to that of other landlords - it is more likely to be in Dublin, to have a lower number of bedrooms, and to be in area where a majority of properties are built after 2000 and have a building energy rating higher than B3. These differences between the properties institutional investors and other landlords own raise a valid concern that the direct impact I estimate will be due to where and what type of properties institutional investors chose to purchase, rather than the rents they chose to set.

To minimise this concern I limit my sample of properties not purchased by institutional investors to only those that are very similar to properties that are purchased. I only retain those that are within 500 metres of a property purchased by an institutional investor, which have the same number of bedrooms and is in a small area with the same level of building quality and age. I also limit the comparisons made to only between units with the same number of bedrooms, within the same county, and the same small area median decade of construction and building energy rating.

With this reduced sample, I estimate that institutional investors increased the level of monthly rents by 4.1 percentage points more than other landlords following purchase. This is an average difference between the percentage increase in the level of rents following purchase from their pre-purchase average for institutional and non-institutional landlord. Expanding with a simplified example, if one year after purchase institutional landlords properties' rents were 4 per cent higher and other landlords were 1 per cent higher, and after two years the corresponding figures were 10.2 per cent and 5 per cent, institutional investors would have increased the level of monthly rents by 4.1 percentage points more than other landlords (i.e. ((4%-1%)+(10.2%-5%))/2 = 4.1 percentage points). It should not be interpreted as institutional investors increasing rents by 4.1 percentage points more than other landlords every year (i.e. a difference in annual growth rates).

Analysing the data in euro terms this translates to an average €78 increase in monthly rents, which indicates that rents in properties purchased by institutional investors are 5.5% higher on average as a consequence of purchase. Limiting the comparison made to properties that exist within developments that contain both institutional landlords and other landlords, I find that institutional investors increase the level of monthly rents by 6.9 percentage points more than other landlords following purchase.

Analysis supports many explanations for why institutional investors charge a premium. Comparing online advertisements of institutional and non-institutional property, data show that institutional landlords take longer to change advertised rents in response to vacancy, and change rents by less when they do. This is suggestive, although far from conclusive, that institutional landlords have greater bargaining power. Analysis shows that institutional landlords are much less likely to leave rents unchanged between contracts than other landlords, which may suggest institutional landlords place more value on profitability relative to other landlords. Finally, the data are partially supportive of institutional landlords offering a higher quality, higher cost service, as the additional income institutional landlords gain from tenants is roughly equivalent to the expenditure they incur on repairs and improvements at an aggregate level. However, analysing a sub-sample of properties where properties' building energy ratings (BER), a proxy for unit quality, are reported, improvements in BERs do not explain the premium charged by institutional landlords. Other service elements, such as responsiveness to queries

and requests, or provision of onsite security, cannot be assessed against other landlords but are evident in some institutional investors marketing materials.

If institutional investors are increasing rents by more than other landlords, other landlords may reconsider how much they are able to increase their rents. If this is the case, landlords in the immediate vicinity of institutional landlords should be more impacted than those who are further away. To test for this indirect impact, I compare how rents change for properties within a given radius of a property purchased by an institutional investor against the change in rents for properties just outside that radius. The results suggest that landlords within 100 metres of an institutional investor raise monthly rents by 2.3 percentage points more than landlords within 100-500 metres following the purchase of a property by an institutional investor. This result decreases as the radius is expanded, and becomes inconsistent when comparisons are only made between landlords within the same local electoral area or electoral district. Whether an indirect impact is present is thus not clear from current methods, and further investigation is required.

In conclusion, the paper finds that institutional investment in existing residential property does increase rents by a modest amount. Further exploration is required to understand the indirect impact of their investment on other landlords.

1. Introduction

Investment in residential property by financial institutions, commonly referred to as institutional investors, is a topic that generates a significant amount of debate amongst policymakers and public commentators. It is either presented as a key driver of increasing rents and rental market dysfunction, or as part of the solution to the problems present in many rental markets around the world (Kollewe, 2017; Lyons, 2021; Gabor & Kohl, 2022). While anecdotes abound, causal analysis of its impact on rental markets is only beginning to be conducted.

This paper proposes to examine what impact institutional investment in existing stock has on residential rents, using Ireland's experience over 2007-2019 as a case study. Institutional landlords could impact rents directly, if they raise their rents by more than other landlords, and/or indirectly, if institutional landlords influence the rent setting behaviour of other landlords.

The nature by which most institutional investors entered the Irish property market gives rise to a quasi-natural experiment. Following Ireland's property market crash portfolios of individual homes and multifamily residential units, or loans on which they were secured, were sold to institutional investors by banks, receivers and state agencies. The data show these properties were overwhelmingly constructed or purchased between 2005-2008. If they were sold or surrendered by their former owners because they were purchased or built just before the market collapsed, then it is reasonable to assume that in the absence of ownership changing, their rents would have followed the trend of comparable units within the same area. Thus, their loss can be considered exogenous to future rental growth, conditional on property characteristics.

While this is an assumption, it is one for which a strong argument can be made. Units were lost by their owners because Ireland's property market crashed. The income the units provided stopped, or fell to the point it could no longer meet the cost of financing purchase, and/or the units could not be sold at a price that would pay back the debt used to purchase/construct them. This reversal, and the lax lending standards that facilitated the preceding boom, were macro factors, unlinked to the potential future rental growth of the properties in question. They potentially were related to property characteristics that impact rents, like location and quality, but we can limit the variation we are exploiting to only that between units with the same characteristics in the same area. In essence, a two bed apartment purchased in 2007 and lost by its owner as a result of the crash should not have a different expected rental growth to a two bed apartment in the same block purchased in 2003 and retained.

Conditioning on property characteristics is important because these properties are not a representative sample of Ireland's rental stock. They are smaller, located in areas with more recent construction and better building energy standards (a proxy for building quality), and are concentrated in Ireland's capital, Dublin. Even within Dublin, there is substantial variation in ownership across the city, with some areas seeing institutional investors account for a substantial portion of agreed tenancies in a given year, while others see much less.

The paper combines newly developed difference-in-difference methodology (Wooldridge, 2021) with exact matching and interaction fixed effects to answer the research question. Institutionally and non-institutionally owned properties are exactly matched across property

characteristics and hyperlocal measures of building stock quality and age to units within a close proximity of them. The preferred specification finds that institutional investors raised the level of monthly rents by 4.1 percentage points more than other landlords following purchase, relative to similar properties in close proximity. Provisional evidence also suggests that institutional landlords may indirectly impact rents by leading other landlords to increase rents as well, although this effect is not present when comparisons are only made between units within geographic areas more granular than an Irish county.[2]

This is an important question for housing policy, both in Ireland and internationally. Institutional investment is generally acknowledged as something that can support residential construction, but whether there is a cost to investment in existing stock is yet to be fully assessed with property level data. In the absence of evidence, the perceived malign influence of institutional investment is something that has underpinned recent rent control legislation in Germany, the Netherlands and Spain (Dombay, 2021). In Berlin, citizens voted to expropriate the properties owned by institutional investors in a non-binding referendum (DW, 2021). It has also influenced the European Green party to fund research developing proposals for EU regulation of institutional landlords (Gabor & Kohl, 2022).

In Ireland, the impact of institutional investors has also been a topic of housing policy discussion. The Department of Finance (2019) has raised the question of whether institutional investors have pricing power in the rental market, and questions on institutional investors market impact have been raised in the Dáil (i.e. Ireland's parliament) (Dáil Debates, 2019). Housing policy developed in response to institutional investment has consisted of additional taxation introduced in 2021 for institutional investors purchasing new homes in bulk (although not apartments) (Horgan-Jones, 2021).

Focusing on investment in existing stock allows a more credible analysis of institutional investments impact and its mechanisms. There is no new supply effect from investment in existing units.[3] This ensures all potential mechanisms through which institutional investment may affect rents should have the same sign, with four possible channels. First, institutional investors size ensures that they have repeated interaction in the rental market, so may be better informed as to what the upper limit of renters' willingness to pay is. Further, their size confers them with bargaining power to charge such a rent and the ability to wait longer to receive that rent. This would be in the same vein as Gallin and Verbrugge's (2019) model of differential rent stickiness between institutional (who they refer to as multi-unit) and other landlords. Second, institutional landlords and other landlords may have different objective functions (i.e. value different things). For example, having a 'good tenant' may be more important to non-

[2] The Republic of Ireland is divided into 26 counties. There is no more than one city within each county. The majority of counties have no settlement that reaches the status of city, instead having towns. For context, the Ireland's largest town (Drogheda) had 14,558 dwellings occupied by households in 2016.
[3] A number of recent papers (Li, 2021; Pennington, 2021; Asquith et al, 2023) have provided empirical support to the idea that rental markets follow the laws of supply and demand –(i.e. an increase in the supply of new rental housing leads rents in local rental markets to be lower than they otherwise would be

institutional landlords, and they may be willing to accept less in rent to retain such tenants.[4] Third, institutional investors may also operate a higher cost, higher quality business model, with this reflected in higher rents.[5] Fourth, if they do offer a higher quality rental service, institutional landlords may see higher wage tenants sort towards them (i.e. they may self-select to apply for institutional units only), which could also bid up rents. Data are not sufficient to determine which explanation is most valid, but what is available can be used to assess their reasonableness.

This paper adds to the small literature that investigates how institutional landlords impact rental market outcomes. Gurun et al. (2022), analyse the impact of mergers between US institutional investors on residential rents, by comparing neighbourhoods where both companies in a merger held properties to those in which only one company did using a difference-in-differences design. Where institutional investors gain more than 5 properties, they find that rents at the neighbourhood level increase by 0.52% after a merger, while rents at the property level increase by 1.2%.[6] Lambie-Hanson, Li and Slonkosky (2019) and Garriga, Gette & Tsouderou (2021) both analyse the impact of US institutional investment on housing market outcomes using a shift share IV strategy. Both papers find a positive effect on rents at the county and metropolitan statistical area, respectively, although they are only significant in Garriga, Gette and Tsouderou. Raymond et al. (2016) study whether landlord heterogeneity plays a role in evictions, finding that corporate landlords are 8% more likely to evict tenants than other landlords. This work builds on these papers by focusing on how rents at a property level change following purchase by institutional investors from non-institutional landlords.

This work also takes inspiration from papers seeking to understand how landlord heterogeneity mediates the impact of rent control. O'Toole, Martinez-Cillero and Ahrens (2021) analyse the impact of Irish rent control legislation, and assess whether corporate landlords (of which institutional investors are a subset) reacted differently to other landlords.[7] They find that nominal rigidities in rent were less likely to occur for company landlords following the policy change and that they were more likely to increase rents by the regulatory maximum relative individual landlords. In a similar vein, Diamond, McQuade and Qian (2019) investigate whether landlord type (individual or corporate) affected responses to introduction of rent control legislation in San Francisco. They find that corporate landlords are 4 times more likely to reduce their supply of rental properties, either via sales or renovations so that they can be excluded from rent control measures. This paper will add to this by focusing directly on differences in behaviour between institutional and non-institutional landlords.

[4] This could also include other landlords placing greater value on capital gains associated with the sale of the property, amongst other factors.
[5] This could include larger amounts of money spent on repair and maintenance or greater responsiveness to tenant requests and queries. Better amenities are possible, but less likely given that in building gyms, entertainment rooms and other amenities were not common features of developments built in the mid-late 2000s.
[6] Austin (2023) employs a similar identification strategy in the US city of Atlanta with neighbourhood level rent indices and finds that rents in areas where both landlords owned units increased by 3.5% more on average following mergers.
[7] There is no common definition of corporate landlord, but generally it is applied to any landlord which is a company. This company could have a single shareholder and own a single property, or be REIT with thousands of shareholders and thousands of properties.

The paper proceeds as follows. Section 2 provides background and context on Ireland's housing market and institutional investors' role within it. Section 3 describes the sources of data used for analysis, methods used to enrich them and the resulting dataset. Section 4 describes the identification strategy for the paper. Section 5 outlines the results for the direct impact of institutional investment on residential rents, considers treatment effect dynamics and evaluates potential explanations for why the treatment effect may be present. Section 6 investigates the possibilities of institutional landlords influencing other landlords. Section 7 concludes.

2. Context: Institutional Investors and Ireland's Rental Market

To present the analysis that follows in its proper context, it is first required to define what is meant by the term 'institutional investor', and provide a brief account of Ireland's housing market since the 2000s and their place within it. This context is relevant to the research design's credibility.

2.1 Who Are Institutional Investors and What Have They Purchased?

In the context of this paper, I define institutional investors as any collective investment vehicle (i.e. investment fund), or their subsidiaries, investing in residential property. In Ireland, this population consists of Irish resident authorised investment funds and real estate investment trusts (REITs) (see Chart 1). There are also a smaller, but still significant number of properties which are held by Irish and foreign companies acting as subsidiaries of a fund, REIT or other collective investment vehicle. Funds are also used in this way e.g. a foreign fund sets up an Irish fund which it uses invests in Irish property.

This definition of institutional investor does not include pension funds, insurance companies or developers who own large portfolios of rental property. Insurance companies and developers are not collective investment vehicles, while pension funds are, but complete, granular data on their holdings of property is unavailable.

The most significant absence is likely developers. One single developer owns at least 1,300 units (Mulligan, 2022), and company identification codes in the data the paper uses suggests that there are a small number of others who own more than 100 rental units. Insurance company data suggests they hold approximately 120 residential units directly on their balance sheet at end 2019.

While pension funds hold significant amounts of non-financial assets, €1bn at end 2019, this is not necessarily residential property. Of the €0.2bn of this €1bn subject to granular reporting, no property is categorised as residential. A further €0.2bn of the €1bn is held by small schemes, typically consisting of a single member, which would not be collective investment vehicles. If the remaining €0.6bn is solely Irish residential property, then it could translate to a maximum of

2,307 units absent from my sample based on the average sale price of a home at end 2019.[8][9] However, this €0.6bn covers both Irish and non-Irish non-financial assets, so the number of units held is likely much lower than 2,307. Moreover, unlike developers, pension funds are not a clear owner of multiple properties in the paper's data.[10]

At end 2019, institutional investors' total holdings of residential property amounted to 12,643 units. Of this total, only 9,862 are ultimately in the treatment group. New builds purchased or constructed by institutional investors are excluded, as there is no change in ownership that we can use to estimate the impact of purchase by an institutional investor (i.e. they are always treated). In addition, given that such units contribute to new supply of rental property, their inclusion would distort the estimation of any treatment effect.

Chart 1: Approx. number of units held by institutional investors by investor type, end 2019.

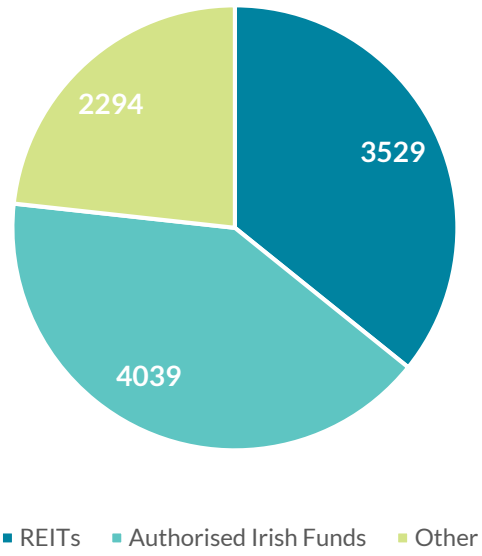

Source: Author's calculations
Notes: Sample only includes units included in treatment group (existing stock).

While institutional investors have continuously invested in rental property for the last decade, Ireland's rental market is still dominated by small landlords. Landlords holding one or two units still make up the majority of properties held, approximately 53.5% of private tenancies (Amárach Research, 2021). In contrast, landlords holding 100 properties or more, a rough proxy capturing most institutional investors, account for around 6% of private tenancies (ibid). All of

[8] The scope of pension fund data extends to all pension funds resident in Ireland, ranging from self-administered schemes to large corporate pension funds.
[9] Reported to be roughly €260,000 at end 2019.
[10] In the event that there are significant numbers of pension fund properties in the control group, this may bias the results.

this suggests that institutional investors do not have substantial market power at a national level.

While these units are spread across Ireland, they are most located in cities, especially Dublin. Of the 9,862 properties held at end 2019, 8,366 are located in Dublin, with another 101 properties being located in the commuter counties surrounding Dublin. Cork, Ireland's second largest city, had 770 properties owned by institutional investors at the end of 2019. To put this in perspective, this represents approximately 7% of Dublin's private rental sector stock, and 2% of Cork county's private rental stock.[11]

However, at a local level their market power may be more significant. This is apparent when looking at the percentage of new tenancies registered by institutional landlords at the level of local electoral areas (LEAs), a sub-county level administrative unit in Ireland. In 2018, institutional landlords collectively accounted for more than 10% of tenancies in 11 local electoral areas, and in two of these electoral areas they held more than 20%.[12] The LEA with the highest share of new tenancies agreed with institutional investors is Glencullen-Sandyford LEA, with 35% agreed with a number of different institutional investors (chart 3). This concentration is primarily a feature of Dublin, with only Cavan-Belturbet, Westport and Kinnegad LEAs seeing shares of new tenancies agreed by institutional investors above 10%. Kinnegad, and to a lesser extent Westport's, inclusion is primarily a function of a small denominator.

Chart 2: % of new tenancies agreed with institutional investors in 2018, by local electoral area

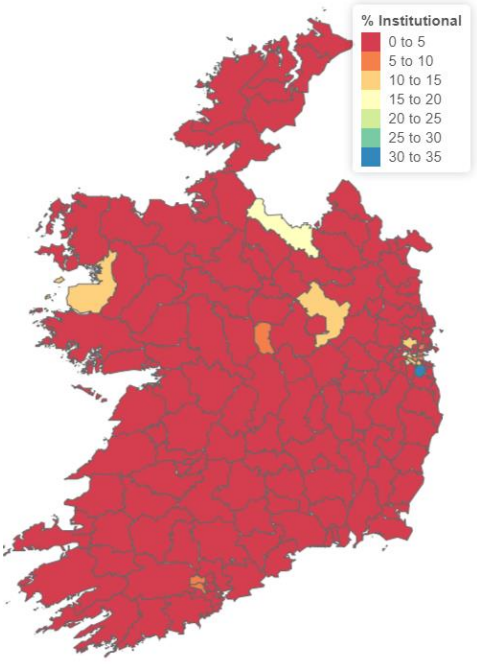

Source: Residential Tenancy Board, Author's calculations

Chart 3: % of new tenancies agreed with institutional investors in 2018 within Dublin, by local electoral area

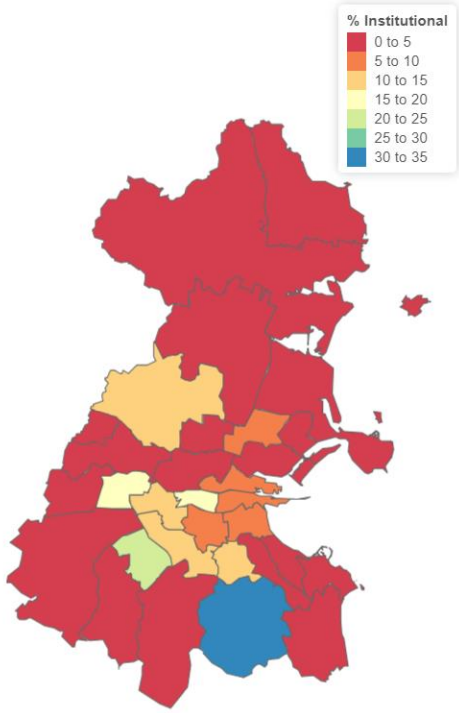

Source: Residential Tenancy Board, Author's calculations

[11] The total private rental stock figure is taken from the 2016 Census collected by the Central Statistics Office

[12] 2018's data is used as it is the last year for which a full year of data is available.

2.2. Institutional Investor's entrance into Ireland's rental market

Prior to the global financial crisis, institutional investors were absent from Ireland's residential property market. Ireland's property market was developer-led, and their business model relied on selling what they constructed, primarily to homeowners, but also to buy-to-let investors (typically small scale). The period in the run to the crisis was characterised by a housing boom, as increasing housing supply was met with increased demand and accommodative credit conditions, seeing residential property prices rise alongside rents.

Ireland's property market turned in 2007, with serious consequences for households, banks and developers. Property prices fell sharply, with the average national property price declining by 55% between 2007 and 2013 (see Chart 4). The national average rent fell by 25%, reaching its trough 2012 (see Chart 4). Ireland's domestic banks had significant exposures to loans secured on Irish property. As demand for new developments vanished, developers became unable to pay their loans. Likewise, individuals, including buy-to-let investors, struggled to meet their loan payments due to a fall in their income (rental and non-rental).

Chart 4: New supply, sale price and rent of residential property in Ireland.

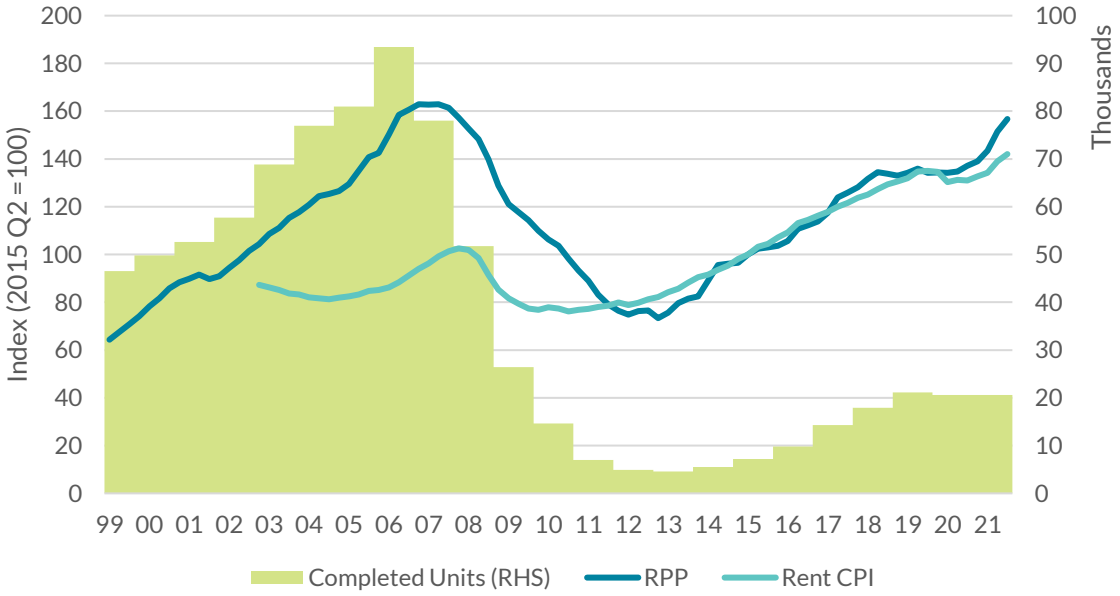

Source: CSO, ESRI
Notes: Both the RPP and supply are blended series. The RPP extends back to 2005, but a further historical series is available from the ESRI stretching back to 1990s. The first index value for the CSO series is adjusted back using the growth rate of the ESRI measure. For new supply, the data switches from ESB connections data to new dwelling completions data in 2011.

The Irish government intervened in the banking sector in an attempt to protect the stability of the Irish financial system and to support the provision of credit. One key element of this intervention was the creation of the National Asset Management Agency (NAMA) in 2009. NAMA purchased property-related loans (often made to property developers) from participating banks. The acquisition price paid by NAMA for the loan was at a discount to the par

value of the loan. Post-acquisition of the loans by NAMA, borrowers would have to work out a repayment plan with NAMA, which normally involved the disposal of the collateral securing the loans. Typically, borrowers remained in control of the underlying collateral before it was disposed of.[13]

Where tenanted residential property formed part of the collateral for the loans, its sale to institutional investors was viewed by NAMA as the commercially optimal sales strategy. In advance of the disposal of these residential properties, NAMA typically ensured that the individual properties were rented out and tenanted to generate secure rental income and enhance their disposal value. In addition, NAMA considered that selling to institutional landlords was less likely to result in existing tenants being evicted (Daly, 2016). Institutional investors preferred to purchase portfolios of tenanted, completed units in multifamily developments (henceforth, multifamily units) as they were already producing an income stream (Dáil Éireann, 2016). Thus, portfolios of properties auctioned on the open market often sold to institutional investors (especially larger portfolios), and these properties form the core of the dataset.

However, NAMA debtors were not the only source of property sold to institutional investors. Banks that decided against selling non-performing loans secured on property to NAMA seized the collateral and sold it to institutional investors. Typically, it seems these properties also were tenanted when sold, and consisted of multifamily units (O'Donovan, 2012; Fagan, 2013). It also appears that some of what institutional investors purchased was portfolios of single-family homes (i.e. one off properties, not clustered within a development), potentially owned as buy-to-let properties, where the landlord had defaulted on the mortgage. This was in contrast to NAMA, who did not acquire retail mortgages and thus did not sell such portfolios.

Residential properties were sold by NAMA debtors and banks because they were developed or purchased at the peak of Ireland's property market bubble. Developers and buy-to-let landlords no longer able to repay their debt surrendered the properties used as collateral to their lenders, or sold them in an attempt to pay down their debt. Data on the year of mortgage origination for properties in my dataset is unavailable, but trends in the Irish mortgage market show there is a sharp discontinuity in voluntary surrender and defaults for mortgages originated over 2005-2008 (Chart 5). Data is available for the year of construction for most multifamily units in my dataset, which show an even sharper discontinuity for the 2005-2008 period (Chart 6). If these properties transferred to banks/NAMA because they were purchased or constructed just before the largest economic shock in Ireland's history, then this should be an exogenous source of variation, once property characteristics and location are controlled for.

As the stock of residential property being sold by distressed debtors and NAMA declined, institutional investors started changing their focus to purchasing new stock (Daly, 2022). This was aided by changes to planning policy to promote the construction of build-to-rent units, targeted at institutional investors, who were identified as an additional source of capital to help finance residential development (Rebuilding Ireland, 2016; Walsh & NicAongusa, 2016). Institutional investors also responded by constructing new residential units themselves, either

[13] Where debtors did not engage with NAMA receivers would have been appointed and would have taken control of the collateral.

in developments that were partially finished, or in new sites. This practice pre-dated the turn to investment in new stock, and involved both build-to-rent and build-to-sell activities by institutional investors. As mentioned in the Section 2.1, these units do not form part of the sample.

Chart 5: No. mortgages ending in voluntary surrender or repossessions of property by year of origination

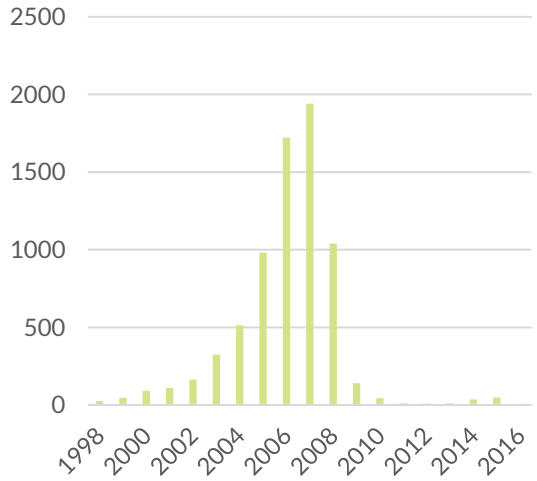

Source: Central Credit Register, Central Bank of Ireland
Note: Data captures any BTL or PDH mortgage still on books from 2017 onwards where collateral was surrendered or repossessed

Chart 6: No. units held by institutional investors by year of construction

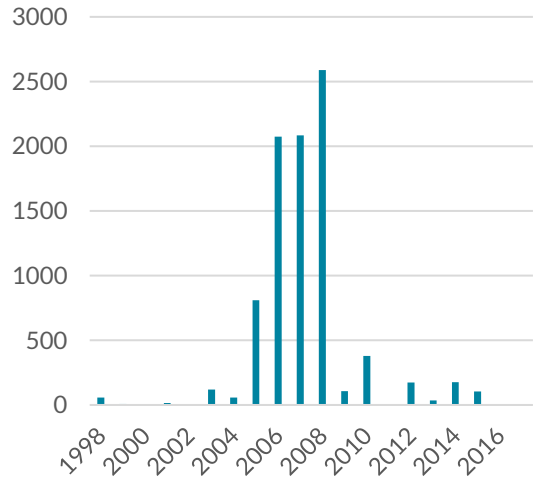

Source: Author's calculations

3. Data

This paper combines a number of novel datasets with publicly available data, and enriches them with the assistance of machine learning algorithms so that they can be connected. This section outlines the data sources used, data quality issues faced, and solutions developed to overcome those issues, before finishing with a comparison of the treatment and control group.

3.1 Sources

3.1.1 Residential Tenancy Board Dataset

The core dataset for this paper is the Residential Tenancy Board's (RTB) Registered Tenancies dataset (RTD). This dataset includes all records of tenancies registered by landlords over the period 2007-2019Q2, and includes around 1.24 million observations for private landlords. Every time a new tenancy is agreed between landlord and tenant, the landlord is required to send a registration form with property, rent amount, and tenant details to the RTB. Over the sample period, observations were only recorded when tenancies are agreed, this means the data form an unbalanced panel, with observations primarily consisting of changes between tenants, rather than for the same tenants over time. On receipt, 77% of observations had Eircodes -

unique, property-level identification codes. Eircodes are required to identify the same unit in a panel over time and to match the RTD to other sources of data.

I used a variety of strategies to enhance the Eircode coverage and accuracy of the RTD. To help fill the 21% of Eircode entries that are null in the RTD, I utilised a genetic machine learning fuzzy matching algorithm to predict Eircodes on the basis of the address provided in the RTD. The algorithm was trained on the existing RTD data where Eircodes were filled in.

The accuracy of the matching algorithm was tested against a sample of 14,320 observations that were predicted to be owned by institutional landlords. The algorithm correctly predicted the Eircode for 11,922 observations, translating to an error rate of 17%. However, the error rate improves if we apply a threshold value which the score for each match must be above. For example, a threshold of 1,000 yields an error rate of 3%, while a threshold of 2000 yields an error rate of 0.3%. However, there are more than 5 times as many matches with a threshold of 1000 relative to 2000 (110,000 vs 21,000). This improves the Eircode coverage to 87% of observations.

Due to the tendency for the data to be unbalanced, a choice was made to prioritise sample size over possible measurement error, and adopt a threshold matching score of 1,000. This can lead to a downward bias in the treatment effect if it results in institutional and non-institutional properties being categorised as the same property. It can also lead to greater imprecision in estimates if different properties owned by a given landlord type are identified as the same property (i.e. two different properties owned by non-institutional landlords identified as the same property). The sample reduces to 1.12 million observations after enhancing the data with the algorithm and limiting it to observations with Eircodes.

I also manually search the data to minimise the possibility of measurement error in landlord type. Non-existent Eircodes are present in the data, and their inclusion can result in incorrect ownership status being assigned to them. To counter this I search the RTD by address for the treated properties in the sample that already received Eircodes. I then checked the correct Eircode for each address against the Eircode in the RTD. Finally, I checked cases where the second to fifth best algorithm score was a property purchased by an institutional investor. This led to approximately 5,000 additional Eircodes being imputed into the data.

While Eircodes may be unique to a property, a property may be divided into many sub-units that do not have their own Eircodes. To get around this issue, I identified cases where the same Eircode was observed multiple times for properties with different numbers in their address – e.g. flat 1 123 Fake Street, D01FAKE, flat 2 123 Fake Street, D01FAKE. To do so, I stripped all non-numeric characters from the address string and identified the mode of the remaining numbers in the strings (i.e. 123). This was also performed for sub-units which were distinguished with non-numeric characters – e.g. Top Flat 123 Fake Street, D01FAKE, Bottom Flat 123 Fake Street, D01FAKE. I then appended the non-modal number in each string or an alphabetic code to the Eircode, rendering them as unique (i.e. 1_D01FAKE, 2_D01FAKE, TF_D01FAKE, BF_D01FAKE).

Recent results from the 2022 Census in Ireland have raised a question as to whether data may be missing from the registered tenancies dataset. The number of households that reported renting from the private rental sector in the 2022 Census is approximately 20% higher (approx.

50,000 units) than the number of tenancies registered with the RTB that year for private rental sector accommodation (Joint Committee on Housing, Local Government and Heritage, 2023). Some of the difference may stem from compositional issues, such as households in the Census incorrectly classifying themselves as being in the private rental sector, or certain type of student accommodation tenancies not considered as private tenancies in the RTB data (ibid). However, it's also possible that some of it is emerging from non-reporting by landlords, which raises questions for identification. The RTB and the Central Statistics Office are working together to assess the cause of this discrepancy.

3.1.2 Novel Dataset of Institutional Investor Holdings

The analysis requires the identification of which specific properties were purchased and sold (if they are ever sold) by institutional investors, and when the transaction took place. A variety of sources was used to create a novel dataset with this information. The majority of institutional investor purchases are by investment funds. These funds must provide annual financial statements to the Central Bank of Ireland, which list the addresses or development name of properties that were purchased. The financial statements of real estate investment trusts (REITs) provide similar information, and are publicly available.

We also account for residential property purchased by entities other than Irish domiciled investment funds and REITs. For example, investors purchase through non-fund, Irish domiciled entities or directly through foreign domiciled entities. To account for this type of purchasing, we consult industry quarterly and annual reports on the largest purchases of residential property, and use Origin Capital's weekly round-up of property news to capture anything we may have missed in other sources.

The date when properties were purchased and sold, and where development names are provided instead of unit addresses, which specific units in a development are purchased, are required to complete the dataset. Dates of purchase and sale are available in Ireland's property price register. In most cases, the properties listed in financial statements and other sources are multi-family developments, which must be broken out into individual units. This is possible with the data in the property price register, providing units are listed individually.

In cases where multiple units are listed as one item in the property price register, I utilise a variety of alternative data sources that provide detail on which exact units were purchased. This includes estate agents promotional material for developments, landlord and property data from the RTD, company registration office charge documents and newspaper articles. Where all units in a development are purchased, I can use the Eircode database to find addresses and Eircodes for all units in the development.

3.1.3 Other Datasets

The quality of a property will be a significant determinant of its rent. While direct measures of quality at the property level are not sufficiently populated in the Registered Tenancies Dataset, variables proxying quality are available from the Sustainable Energy Authority of Ireland (SEAI) at the small area level. Small areas are the most disaggregated spatial unit of data collection in Irish national statistics, corresponding to less than 100 individual properties on average.

The SEAI's Building Energy Rating of Residential Homes Register contains data to produce two measures of quality at the small area level. The first is the building energy rating (BER) for properties in a small area. This is a categorical variable that grades properties according to how energy efficient they are. It encompasses features like windows, doors, insulation and heating, all of which contribute to the quality of a home. Research has shown that higher BERs are associated with higher rents and property values (Hyland, Lyons, & Lyons, 2013). Furthermore, new homes being built have a minimum BER of A3 since April 2019, and regulations improving building standards over the period have corresponded with improved BERs (Irish Concrete, 2009). The second is the date of construction of units in small areas.

These data are anonymised and cannot be linked to units. To overcome this, I take the median for each over the sample period within each small area. Given the granularity of small areas, the median should generally be a good proxy for the quality of individual units.

In addition, I collect data from a number of sources to assess the plausibility of explanations for why a premium exists (section 5.3). I record the expenses from the profit and loss accounts of authorised investment funds and real estate investment trusts to see how they compare to additional income. The financial statements of REITs are public, while investment funds have to submit financial statements privately to the Central Bank of Ireland. From these documents, I collect data on any type of expense that classifies as refurbishing, improvements or renovation, and other undefined capital expenditures.

I use advertisement data from Daft.ie, Ireland's primary property advertisement website, to assess whether institutional investors behave differently to other landlords online when advertising property. Specifically, I look at the frequency, value and average time taken for decreases in asking rents. In doing so, I limit the observations to only those with Eircodes, and see rent changes after 1 week to 6 months. This should remove mistakes in advertised rents, and observations that should not be included.

3.2 Dataset

Once all sources are joined the resulting data structure is an unbalanced panel, with 13 years of data (2007-2019) and 253,304 unique properties totalling to 834,877 observations. This is smaller than 1.23 million figure previously mentioned. This occurs for six reasons: first, I remove all observations missing Eircodes as they cannot form part of any panel analysis. Second, I remove student accommodation. It is arguably a market segmented from the rest of the property market, and it has data quality issues.[14] Third, I remove all observations with monthly rents greater than 10,000 and less than 200 - this is the outlier threshold used by the RTB. Fourth, Wooldridge's (2021) difference-in-difference method requires that any treated units remain treated for the whole period, and does not accommodate units that are in a treatment group year, but are not yet treated. Fifth, I only select the latest observation in each year as the analysis uses annual frequency. Sixth, I remove all Eircodes that only appear once in the data –

[14] The frequency of rent payments vary from year to year, making comparison difficult. In addition, multiple units may be let one year under the same Eircode, while the following year they are subdivided. Both factors lead to substantial variation in rents that do not represent real changes in rent per unit.

they cannot form part of any panel analysis. Table 4 in Appendix A outlines the contribution restriction makes to reducing the sample.

The resulting data show that properties held by institutional investors do differ substantially from those held by other landlords (see Table 1). They are much more likely to be located in Dublin, built in small areas with a median decade of construction of 2000 or greater and to have median BERs of B3 or above. They are also around half a bedroom smaller on average, and are nearly twice as likely to be in rent pressure zones (areas subject to rent control – expanded on later).

Table 1: Summary Statistics for Institutional and Non-Institutional Landlords

	Institutional	Non-Institutional	Difference
Share in Dublin	0.83	0.42	0.42 ***
No. Bed	1.93	2.51	-0.58 ***
Share Rent Pressure Zone	0.19	0.09	0.10 ***
BER B3 or Above	0.44	0.06	0.38 ***
Decade of Construction >= 2000	0.82	0.42	0.40 ***
N	27,775	807,085	

Source: Author's calculations
Notes: Share in Dublin refers to share of observations relating to properties located in Dublin. No. Bed refers to the mean number of beds of observations in each group. Share Rent Pressure Zone refers to the share of observations located in a rent pressure zone. BER B3 or above refers to the share of observations which are from residential units located in a small area with a median BER of B3 or higher. Decade of Construction >= 2000 means the share of observations located in small areas with a median decade of construction of 2000's or more recent. Significance level displayed as * p<0.1; **p<0.05; ***p<0.01.

There are also features of the data that are beneficial to the research. The treatment group is a small part of the total sample, amounting to only 27,775 of the total 834,860. This facilitates exact matching, as large numbers of control units dissimilar to treatment units can be dropped while still retaining a large sample. In addition, the average frequency of observations per unit is very similar for treatment and control units. Treatment units see on average 3.33 observations while control units see 3.31. Thus, they are roughly seeing similar levels of turnover between tenants.

4. Identification Strategy

The paper aims to answer whether institutional investors increase rents more than other types of landlords following purchase. This type of question is amenable to a staggered difference-in-difference (DiD) design. However, by itself DiD is likely insufficient to produce unbiased estimates of the treatment effect. The properties that institutional investors own are clearly different to those owned by other landlords.

Therefore, to make identification more credible we make three further additions. First, property characteristic-location-year fixed effects are included so that only variation amongst similar properties in the same area at the same time is being used to estimate the impact of institutional investment on rents. Second, the treatment and control group units are exactly matched on number of beds, median small area BER and median small area decade of construction, as well as requiring that all treatment/control units be within a specific distance of their nearest matched control/treatment unit. Third a specification is run only on properties that exits within developments containing both treatment and control units.

4.1 Difference-in-Differences

Recent literature has demonstrated two-way fixed effects difference-in-difference estimators have a number of potential problems with their use (Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Boryrusak & Jaravel, 2017). Estimates are biased in the presence of treatment effect heterogeneity, and negative weights can lead to the sign of a treatment effect flipping. To avoid these issues I use (Wooldridge, 2021) two-way Mundlak difference-in-difference estimator (also referred to as extended two way fixed effects).

Wooldridge's methods are preferred as they are simpler and more flexible. They build upon the existing two-way fixed effects framework by replacing a single difference-in-difference treatment effect with an extended set of period-treatment cohort (i.e. treatment group year) treatment effects, which are aggregated up into desired estimators. They can handle interaction fixed effects and multiple treatments, whereas other models (and the packages supporting them) tend to be more rigid.

Its ability to account for multiple treatments is important, because from late 2016 rent control was introduced in a staggered way across Ireland. If a local electoral area saw rent increasing by more than 7% annually in 4 of the last 6 quarters, and, that the average rent for the area was above the national average, an area would be designated as a rent pressure zone (RPZ). Properties in RPZs were limited to a 4% annual rent increase per annum, with exceptions where significant improvements have been made to the property. As RPZ limit landlords' abilities to increase rents, it's clear that they need to be controlled for, as otherwise the results will include differential exposure to rent pressure zones between institutional investors and other landlords.[15]

Using Wooldridge's estimator in this setting raises a particular issue for how to manage observations that appear for a treated unit in its year of treatment before it is actually treated. For example, a unit may be purchased in December 2014, but the only observation for that unit in 2014 is in June. Because the data is aggregated to an annual format the unit is considered as forming part of the 2014 treatment group.

[15] It's also possible that rent pressure zones have influenced institutional investor behaviour. Not increasing rents by the regulatory limit means that the rental income received on a property is permanently lower, which will ultimately affect the sales value of their investment. If institutional investor place greater importance on this than other landlords, then it would form part of the 'different objective functions' explanation for why institutional investors may increase rents by more than other landlords.

Regardless of how such units are managed, they will introduce some form of bias. If left in the sample they will be classified as treated and exert a negative bias on the estimate. If they are moved into the next treatment group year (e.g. our example is treated in 2015) this will also bias estimates if treatment effects vary across treatment groups, and doing so makes regression output more difficult to interpret.

Finally, there is the option to remove such observations, which may also result in biased estimates. Units that ultimately should form part of a pre-treatment average will be removed, and if they differ on average from other pre-treatment observations then estimates will be biased. However, removing units only introduces one source of bias, and does so in a less messy way than adjusting treatment year, so this approach is chosen.[16]

To use Wooldridge's estimator we need a slightly modified parallel trends assumption, such that in the absence of units in a treatment cohort (i.e. properties purchased in 2014) being purchased by institutional landlords their rents would have followed the same trend as the control group. In addition, the method requires a no anticipation assumption, such that we see no divergence between a given treatment cohort group and the control group in the pre-treatment period. When covariates are included, these assumptions become conditional upon them.

4.2 Matching and Other Considerations

The parallel trend assumption is not entirely credible in light of the differences observed between treatment and control units in section 4. There are many factors affecting local rental market demand and supply conditions that vary across the country and also determine rents.

One method to deal with this concern is using a combination of exact and distance matching of treatment and control units to make the groups more comparable. Specifically, I limit the sample to only those control units which are within 500/250/150 metre of a treatment unit which has the same number of beds, same median decade of construction and same median BER. All units without a match are discarded. This is also repeated for the treatment group. By limiting my sample to similar quality units that are in close proximity to each other, the presence of selection bias should be reduced.

To elaborate on what this matching method, consider the example illustrated in Chart 7 below. In it there is one treatment unit, T (the red dot), and a number of control units, C1-C5. The text lists their number of beds, median decade of construction and BER for their small area. The circle is a 500m radius around the treated unit. Imagine this is the entire sample. The matching procedure works by first identifying whether each control unit has a match to another property in the sample with the same number of beds, median BER and decade of construction. Only control units with matches are retained (in this example C1 and C5). It then calculates the distance between the retained control unit and their matched treatment unit(s), and selects the minimum distance amongst all matches for each control unit. If this is less than 500m it is

[16] Doing so leads the number of treatment group observations to shrink by approximately 6% (approx. 1,900 observations).

retained, otherwise it is not. After this only C1 would be retained, as it exists within 500m of the treatment unit.

Chart 7: Graphic illustration of exact matching procedure.

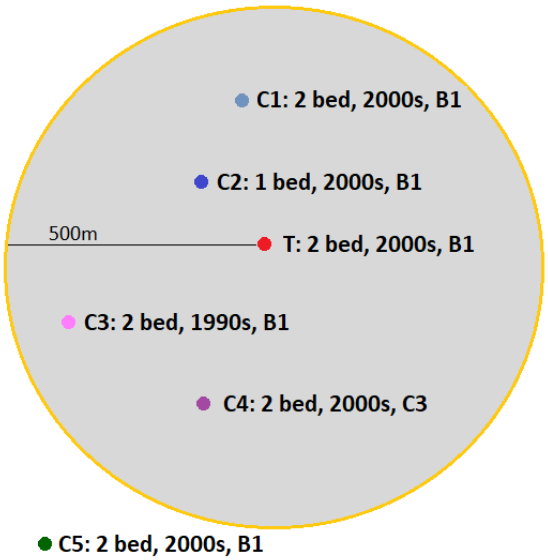

This many to one matching method makes the treatment and control groups more similar, but it does not eliminate differences between them entirely. The differences highlighted in section 3.2 are reduced but still sizable, with the exception of no. bedrooms which is near identical (see Table 5, Appendix A).

This leaves room for reasonable doubt about the unbiasedness of any results. To minimise them, I use year-county-number of bedrooms-median small area BER – median decade of construction BER interaction fixed effects (henceforth, "year-county-property characteristics fixed effects"). This should ensure that only variation within groups of similar properties is being used to estimate the treatment effect. I also run a specification only looking at properties that exist within developments that are shared by institutional and other landlords, and focus only on within development variation for properties with the same number of beds.

4.3 Threats to Identification.

Asides from local supply and demand conditions, tenant, property and landlord characteristics will also matter in determining rents. Landlord characteristics cannot be separated from landlord identity, which is what we are analysing the impact of. Year-county-property characteristic fixed effects will absorb time varying impacts of time invariant property characteristics and location on rent. If property characteristics do vary over time, this is due to landlord choices (i.e. improving the quality of a rented property) or wear and tear/accidents (fire, property damage, etc.). If institutional landlords are more likely to initiate improvements, they do so with a view to raising the rental value of the property, and therefore this forms part of the treatment effect. The degree to which institutional landlords increase rents due to improved property quality is an interesting nuance to the results, but cannot be determined from the available data.

If institutional investment influences other landlords, estimates of the treatment effect will be underestimated. When setting rents, it would be natural for other landlords to research what asking rents landlords of similar properties nearby are advertising. If institutional investors have the bargaining power to charge higher rents than other landlords, this may affect other landlords' expectation of what an achievable asking rent is. In this way, institutional investors may act as price leaders in local rental markets. If this is the case, the difference-in-difference estimates will be biased downwards. The possibility of such spillover effects is tested for in section 6.

In certain cases, institutional landlords may purchase existing units and develop, or complete construction of, new units within close proximity of them. Assuming that the supply effect of the construction outweighs any amenity effect on rents, this new supply should decrease rents in the immediately surrounding area relative to if it had not been constructed. This would exert a downward bias on the treatment effect. More generally, if the locations of new developments and investment in existing units are correlated, then this too could exert a downward bias on the treatment effect.

Another challenge for identification is that institutional investors may have purchased properties that had rents that were below the equilibrium level for comparable units. If this is the case, then the rents for such properties would return to market levels once they were sold to any new landlord. Similar to O'Toole et al. (2021), I control for this by including a measure of the first observed value of the rent divided by the average rent for that period for properties with the same number of bedrooms in the same area – referred to henceforth as the relative rent.

Finally, if landlords are deliberately not registering tenancies with the RTB then this also presents a threat to identification – although the direction of the bias is unclear. For example, landlords who intended to break rent pressure limits may be more likely not to register tenancies, and to increase rents by more than other landlords. On the assumption these landlords are more likely to be in the control group, their absence from the data would bias the estimate of the impact of institutional investors upwards. In contrast, if this is occurring because properties are being rented informally to friends or family, and this is occurring more among the control group, then the absence of these tenancies would bias estimates downwards.

5. Results

5.1 Results

Table 2 presents a range of difference-in-differences specifications estimating the impact of institutional investment on the level of a properties rent. Specifications (1)-(3) use the full sample, while specifications (4)-(6) used the matched sample. Specifications (1) and (4) are baseline models with unit and year fixed effects; specifications (2) and (5) add county-year-property characteristic fixed effects; specifications (4) and (6) add rent pressure zone and relative rent variables. Specification (7) reduces the sample again to only consider variation between treatment and control units that exist within the same development using development-no. bedrooms-year fixed effects.

All specifications show a clear and positive effect of institutional investment on the rents of properties they purchase. Specification (1) suggests that after purchase, the monthly rent of properties purchased by institutional investors increases by 9.0 percentage points more than properties not purchased by institutional landlords. Limiting comparison to properties within the same county and year that share the same property characteristics (specification 2), we see that this drops to

Table 2: Staggered DiD Regression Specification

Dependent Variable: Log of Monthly rent

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treatment	0.09 (0.009)***	0.042 (0.009)***	0.034 (0.007)***	0.047 (0.01)***	0.044 (0.01)***	0.041 (0.006)***	0.069 (0.008)***
Rent Pressure Zone	No	No	Yes	No	No	Yes	No
Relative rent	No	No	Yes	No	No	Yes	No
Year – County – Property Characteristic FE	No	Yes	Yes	No	Yes	Yes	No
Year – Development – No. Bed FE	No	No	No	No	No	No	Yes
N	834,860	834,359	834,359	141,747	141,747	141,747	37,620

Notes: Table 2 reports estimates of the impact of purchase by an institutional landlord on a residential property's rent. The dependent variable is the log of monthly rent (€ per month). All regressions are performed with Wooldridge's two-way Mundlak (2021) difference-in-difference estimator, with individual group-year effects aggregated into a single treatment effect. Specifications (1)-(3) use the full sample, while specifications (4)-(6) use the sample matched on property characteristics and distance (described in section 4.2). Specification (7) limits the sample to units that are located in a development containing both treatment and control units, and focuses on within development variation. Unit and year fixed effects are included in all specifications. Rent Pressure Zone is a control for the introduction of rent pressure zone legislation at the local electoral area level, and is also a difference-in-difference estimator. Lagged relative rent measures the log of the previous rent of a property relative to the average of similar properties in its local electoral area. Standard errors are in parentheses and significance level displayed as * p<0.1; **p<0.05; ***p<0.01.

an increase of 4.2 percentage points, with this dropping further to 3.4 percentage points once rent pressure zones and properties' starting relative rent are controlled for in specification (3). The drop in estimates between specifications suggests that specification (1) is primarily capturing differences between the type of properties purchased by institutional investors and other landlords, rather than differences in their rent setting behaviour.

The results in the matched sample show that the findings hold once we limit our analysis to only similar units in close proximity. Specification (6) suggests that when only using a sample of treatment/control properties with matching control/treatment units within 500 metres (and using the controls used in specification (3)), the purchase of a property by institutional investors leads rents to increase by 4.1 percentage points relative to properties with owned by other

landlords. Specifications (2) and (5) and (3) and (6) are near identical, while the difference in estimates between (1) and (3) are more considerable.

Specification (6) is the preferred specification as it strikes the best balance between external and internal validity. It includes all controls and uses matching, but only reduces the number of treatment group observations by 4%. While specification (7) likely has stronger internal validity, it only includes treated units in multi-unit developments where institutional investors own some but not all units, and therefore it is uncertain as to how generalizable its estimate would be - 33% of treatment group observations are missing in specification (7).

The results are robust to changes not presented in the above table. The results are positive and significant with a level dependent variable, when using LEAs and electoral districts instead of county in the interaction terms, and when the distance metric for matching is changed to 250 or 100 metres.

The results are similar when the sample is limited to only units within the same development (specification (7)). Institutional landlords raise rents by 6.9 percentage points more than other landlords within the same development following purchase. It might be expected that the figure for specification (7) would be lower than specification (2) or (5), both due to the fact that institutional investors are more likely to have an indirect impact on other landlords within their building than those outside, and because specification ensures a more like-for-like comparison between units. Differences in the treatment group between specifications account for some but not all of this difference in estimates. Specification (2) returns an estimate of 5.3 percentage points instead of 4.2 percentage points when using specification (7)'s treatment group.

Considering the estimate of specification (6) in levels rather than logs produces an estimate of €78. To contextualise this figure, it can compared to the average rent for units purchased by institutional investors – approximately €1,500 in monthly rent. This indicates that rents in properties purchased by institutional investors are 5.5% higher on average as a consequence of purchase.[17]

Whether taken as 4.1 percentage points or 5.5%, these estimates are larger than the estimates of Gurun et al. (2022), who find a 1.2% increase in rents following purchase by institutional investors. However, their study focuses on institutional investors purchasing from each in other in mergers, so one would expect the impact to be smaller as it relates to changes in local market power, rather than landlord type. In regards to its overall importance in explaining variation in rent it accounts for 0.08 standard deviations of monthly rent in the sample.

5.2 Relative Time Event Study Estimates

Event study estimates facilitate the analysis of pre-trends and demonstrate the dynamics of the treatment effect. While the absence of pre-trends does not guarantee parallel post trends, the presence of differing pre-trends is a cause for concern.

[17] This 5.5% is derived from: treatment effect/(post-treatment average rent for treated units – treatment effect).

I follow the method suggested by Wooldridge (2021) for calculating pre-trends. It calculates a difference-in-difference estimate for each treatment group-year combination relative to the initial period in the sample. In my case, I do this for the first two periods (2008, 2007) due to the limited number of observations for the treatment group in each year. Limiting the sample to starting from 2008 and using that as the first year does not substantially change the results. Effect sizes are larger, but significance does not change. The parameters calculated for each group-year combination are then aggregated for each relative period to the treatment date. A positive figure indicates that the difference between treatment and control units at that date is higher than it was in 2008-2007, and a negative figure means its lower. Tracking these estimates over relative periods gives us an idea as to how they are changing year to year.

The post-treatment effects are calculated differently. They use the full pre-treatment sample rather than one year, but are also estimated on a treatment group-year basis, and then aggregated up to relative time periods. This means that the figures being compared between -1 and 0 are not calculated on the same basis, and this should be borne in mind when looking at the charts.

Chart 8: Specification (6) in event study form

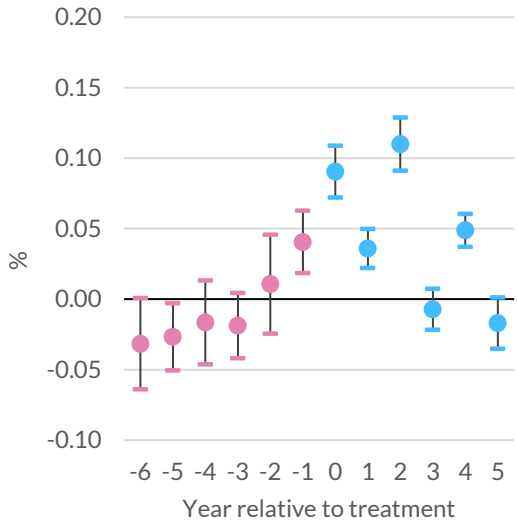

Source: Author's Calculations
Notes: 0 is year of treatment. Circles are treatment effects, bars are 95% confidence intervals. Objects shaded pink are in the pre-treatment estimates, objects shaded blue are estimates of the treatment effect

Chart 9: Specification (7) in event study form

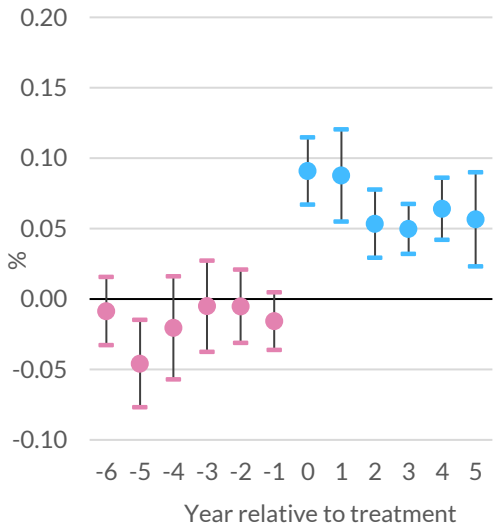

Source: Author's calculations
Notes: 0 is year of treatment. Circles are treatment effects, bars are 95% confidence intervals. Objects shaded pink are in the pre-treatment estimates, objects shaded blue are estimates of the treatment effect

Substantial evidence of pre-trends is limited, although many specifications show an increase in rents immediately prior to purchase, which could indicate anticipation effects. Only specification (1) presents pre-trends that indicate a persistent increase in rents over the whole pre-treatment period, with 6 of 10 pre periods statistically significant. This is not surprising, given it is the baseline model. With the exception of specification (7), all specifications show a significant positive increase in period -1.

If landlords are raising rents on properties in advance of selling them to institutional investors this would violate the no anticipation assumption, leading to a downward bias in the treatment effect. This seems probable, as the positive and significant increase at period -1 is driven by properties treated at the end of the sample that were purchased by construction companies, redeveloped, rented out and then sold on to institutional investors. Removing these properties results in a much smaller, statistically insignificant treatment effect at period -1.[18]

The event study version of specification (7) shows no evidence of pre-trends, and a declining treatment effect. With the exception of one period, pre-trends are consistently insignificant, closer to zero than for the other specifications, and fluctuate between being positive and negative. The post-treatment effects suggest that following the first two years an institutional landlord takes ownership of a property, the differential charged relative to other properties owned by landlords in the development starts to decline. This may suggest that other landlords respond to institutional investors by raising the rents they charge as well, although not by enough to erase the difference.

Adjusting the matching process does not lead pre-trends to meaningfully differ. To demonstrate the pre-trends are not unique to the specification chosen a number of adjustments to the matching process are made. This includes: i) using a 250m radius instead of a 500m radius; ii) using a 750m radius instead of 500m radius; iii) median small area decade of construction removed from matching; iv) median small area BER removed from matching (see Charts 20-23). In 2 of 4 alternative specifications, this leads an additional pre-treatment effect to become insignificant.

The sharp jump in treatment effects observed immediately at date of treatment are supportive of this paper's claim that the effect observed is causal. If it was the case that what is being captured is simply differing trends across areas in which units are located, we would expect a persistently increasing impact of institutional investment across the post-treatment horizon. Instead, we see an initial jump at year of purchase, and a volatile estimate afterwards, likely due to the changing composition of the treatment effect between years and a possible fading of the treatment effect over time.

The results may suggest that the treatment effect only holds over the medium term, although this also in part be due to the shrinking of the treatment group over time. For specification (6) (and (3)-(5)), we see that 5 years after treatment the effect generally becomes statistically insignificant.[19] This would suggest that institutional investors are able to increase rents in excess of other landlords for a period of time, but not permanently. This could be due to affordability constraints being reached by institutional investors, due to other landlords adapting to the behaviour of institutional investors, or simply because institutional investors are

[18] A concern for identification may be that this is endogeneity, with institutional investor purchasing properties because their rents were already rising. While this may seem reasonable, one has to consider that in the absence of the entrance of institutional investors into the rental market it's questionable whether these portfolios of property would have been put together and refurbished, as there would have been no buyer for them. If the creation of these portfolios was created by their entrance, then this situation more closely aligns with anticipation of treatment.

[19] In contrast, the treatment effect seems to be much more stable in the within development specification (i.e. (7)), and stays positive and significant until 7 years after treatment.

able to move faster to the equilibrium rent for a property than other landlords. However this dissipation could also be compositional – more than half of the units in the treatment group do not have observations by period 5 or later (i.e. units purchased from 2016 on do not have observations in period 5 or after).

5.3 Exploring Possible Mechanisms

Understanding why institutional investors increase rents more than other landlords is important if we are to understand consumer welfare implications. Data is available to roughly tease out the plausibility of three explanations for why this premium might arise. They include bargaining power, different objective functions and providing a higher quality service and attracting higher income tenants. Data provides support to them all, or at least does not contradict them.

One way to assess if institutional investors have greater bargaining power is to look at their pricing behaviour relative to other landlords. To do so, we can use daft.ie's data on asking rents. If institutional landlords do have pricing power we would expect that 1) they are less likely to decrease rents; 2) when they do decrease rents, they do by less; and, 3) it takes them much longer to decrease rents.

Analysis shows that there are differences in behaviour between landlord types, providing some support to bargaining power being an explanation.[20] Institutional landlords on average take an additional 30 days to decrease rents on an advertised property which has not yet been let relative to other landlords, who take on average 29 days to do so. Comparing sample means, institutional investors are 7 percentage points less likely to decrease rents, but this disappears once year dummies are included.[21] Considering a sample of rent decreases, institutional landlords decrease their rents by approximately 1 percentage point less than other landlords, although this is only significant at the 10% level.

While we do not observe landlords' objective functions, there are differences in behaviour that are difficult to explain if we think that all landlords' value the same things equally when setting rents. One such behaviour is nominal rigidity in rents (i.e. rents not changing from contract to contract), despite market rents increasing over time alongside inflation. Looking at the data, we see nominal rigidity is much more common amongst non-institutional landlords, with 18.4% of contracts seeing no change in rent from their previous level, while the corresponding figure for institutional landlords is 6.4%.

This behaviour is evident beyond comparing samples means, remaining apparent when considered in a difference-in-difference design. Running specification (2) as a linear probability model with a dependent variable equal to 1 when rents do not change between contracts, we find that following purchase by an institutional landlord, the probability that rent remains unchanged between contracts decreases by 12 percentage points (see Appendix A – Table 6 – specification (16)). This is suggestive that non-institutional landlords may be placing more value

[20] See Appendix A – Table 6 – specifications (13)-(15) for regression results.
[21] While this could also reflect institutional landlords being better informed about what the maximum acceptable asking rent is, one would expect other landlords to be more likely to underestimate rather than overestimate this rent due to the greater cost that vacancy imposes on them

on things other than rental income, such as whether they have a 'good' tenant, a reason suggested for nominal rent rigidity in Gallin and Verbrugge (2019) and O'Toole et al. (2021).

We can explore whether the premium institutional investors charge reflects a higher quality accommodation service by comparing it to the money institutional landlords spend on improving their properties. We cannot compare the expenditure of institutional landlords on property to non-institutional landlords due to an absence of data for the latter. If institutional landlords were to spend nothing on improving the properties they purchase, but still increase rents by more than other landlords, then the argument for that premium reflecting a higher quality service would not be supported by the data.

A number of assumptions and caveats are required when making this comparison. Data is only available for REITs and Irish resident funds. Sometimes 'improvements' are specifically identified in profit and loss statements, while at other times they are tied up in a capital expenditure item which could contain irrelevant expenditures.[22] Another challenge is that individual entities may own commercial and residential property – where this occurs, expenditure is assigned according to the share of residential property in total property holdings for an entity, unless it is clear in the financial statement that expenditure relates to a specific property.

The additional income (taking specification 6's estimate) gained by institutional investors matches the expenditure clearly identified as relating to improvements. Additional income of €31mn is matched with expenditure on improvements of €31mn.[23] When undefined capital expenditures are added to improvements, total expenditure rises to €82mn, more than twice the additional income gained by institutional investors. The data does not contradict that institutional landlords may be providing a higher quality, higher cost service.

Expenditure exceeding the income gained from the premium does not mean an investment will make a loss. This expenditure should improve the capital value of the property as well as its rental income over the lifetime of their investment. On a per annum basis, institutional investors are earning €8mn in additional income relative to if their properties had been retained by other landlords. Assuming that half of their improvement and capital expenditure is realised in higher property values when they decide to sell, they would breakeven on expenditure on improvements in 2 years and capital expenditure in 5, both within the average lifespan of a real estate fund (5-7 years).[24]

The 'higher quality, higher cost' mechanism can also be assessed by controlling for improvements in the building energy rating (BER) of individual units. There is a sub-sample of approximately 128,000 units for which BERs are available. As an improvement in the BER can be considered a binary event it too can be controlled for as a treatment, with the unit being

[22] Capitalised acquisition costs, development expenditures on non-residential properties, etc.
[23] Improvements include any expense items labelled 'improvements', 'addition of fixtures and fittings' and 'refurbishments'. Additional income gained = average no. months property owned * no. properties * specification 6 estimate.
[24] The choice of 50% is arbitrary, and the period required to hold the property until a landlord breaks even on expenditure will increase as if returns are lower.

considered treated if it observed an improvement in its BER.[25] If the premium were to significantly reduce after controlling for improvements in BER then this would give support to the 'higher quality, higher cost' mechanism.

Controlling for improvements in BERs sees a slight increase in the premium institutional investors charge over landlords. Specifications (3) and (6) are replicated on the smaller sample of units which report BERs. Results from these regressions are then compared to results from regressions which add a control for improvements in BERs. Specifications which control for improvements to BER show slightly higher estimates (see Table 7 – Appendix), which would not be supportive of the premium arising from improvements in unit quality made by institutional landlords.

The non-physical element of the higher quality, higher cost service explanation is not something that can be evaluated. For example, there is no data to assess average speed of response to queries, or how likely different landlord are to act on requests by tenants. Marketing materials for developments owned by some institutional landlords highlight security and emergency maintenance services, and these may be something that tenants are willing to pay more for.

If they do provide a higher quality service, they may also attract higher wage tenants, which could also lead to higher rents. This would occur if high wage tenants compete with each other for property owned by institutional investors and bid up its rent in the process. Unlike landlord who only own a single property, institutional landlords have to internalise the externality a bad tenant poses on a development. Given uncertainty for landlords over tenant characteristics, institutional landlords may use wages as part of their screening process, or set rents in an equivalent way (i.e. set rents at levels only high wage tenants could pay), which would further ensure that tenants offered a rental property owned by institutional investors would have higher wages.

Sufficient data is not available for this to be explored directly. However, if higher wager tenants sort towards each other and institutional investors attract higher wage tenants, then institutional investment in existing stock may lead to increased rents in properties in the immediate vicinity of institutional investors.

6. Further Analysis

6.1 Indirect Impact of Institutional Investment

This section tests whether institutional landlords have an impact on landlords in their immediate vicinity. The implicit assumption in section 5 is that they do not, as otherwise the results are likely to be biased downwards. This is a strong assumption, given that institutional landlords do

[25] This is somewhat of a simplification, as BERs can increase by varying degrees of intensity. However, continuous treatment DiDs require very strong assumptions to identify the average treatment effect (Callaway, Bacon-Goodman, & Sant'Anna, 2021), so this is treated as a binary treatment. In addition, if a unit has multiple improvements in its BER, only the first is considered. Around 4% of improvements occur in units that have already seen at least one improvement, so their omission should not be too impactful.

increase rents by more than other landlords and the asking rents they charge would be visible to other landlords on rental market advertisement websites.

It seems unlikely that non-institutional landlords would pay no attention to the rents set by institutional landlords, although how much of an influence institutional landlords would have is unclear. One would expect that this influence would be increasing in the share of rental property in the local area owned by institutional investors, and with the substitutability of non-institutional landlords' property with that owned by institutional landlords against which they are comparing. The analysis that follows does not incorporate these two factors, but as a starting point investigates if there is evidence of an indirect impact of an institutional landlord purchasing a property on any non-institutional properties in close proximity. If an effect is observed irrespective of these two factors, then it is more likely to be present once they are accounted for.

If institutional investment results in an improvement in local amenities, this too may lead to an increase in rents of properties owned by other landlords in their immediate vicinity. This improvement in amenities is unlikely to directly result from investment in amenities by institutional landlords, as the paper focuses on investment in existing stock. However, there are a very limited number of cases where investment in existing units was associated with completion/construction of other residential/commercial units nearby. It is also possible that institutional investment may attract new commercial development in the nearby vicinity, if businesses believe that institutional landlords will attract higher wage tenants. Tenant sorting amongst high wage tenants, as outlined at end of section 5.3, could also have a similar impact.

To assess whether there is an indirect impact, we can follow the "ring" approach used by Asquith et al. (Li, 2021)(2023) and Pennington (2021) who have studied the impact of new rental supply on local rental markets.[26] Their general approach is to assume that the impact of new supply only exists within a certain radius of its location, and compare units within that radius to other units within a limited range outside this radius (i.e. within 0-100m vs 100-250m).

We use these methods with a variety of radii (100 to 300 metres, in cumulative 50 metre increments) for properties within 500 metres of a property that is purchased by an institutional investor. To elaborate, any unit which is within X metres of an institutionally owned property is considered treated, and units from X-500 metre are the control group. These units' radii are arbitrary, and it is possible that the treatment impact would extend beyond them. The treatment date for these units is the year the institutionally owned property was purchased. If subsequently, more units are purchased within a X metre radius this does not change the treatment status of the treated units. As in specification (2), I use the two-way Mundlak estimator with county-year-property characteristic interaction fixed effects.

The results of these regressions are shown in the table below (Table 3). The treated within row informs the reader of the distance from a unit purchased by an institutional landlord within which a unit owned by a non-institutional landlord is considered treated. For example, in specification (8) properties within 100 metres of units purchased by institutional investors are considered treated, while units within 100-500 metres are considered the control group.

[26] Originally developed in Campbell et al. (2011) to study externalities of foreclosures/forced sales.

All specifications suggest that following investment by institutional landlords, non-institutional landlords in close proximity to these purchased units increase their rents by more than non-institutional landlords who are further away. Considering specification (8), it appears that non-institutional landlords within 100 metres of a property purchased by an institutional investor increase monthly rents by 2.3 percentage points more than non-institutional landlords with properties within 100-500 metres of a property purchased by an institutional investor. This estimate varies with the radius of the treatment group, but always remains positive and statistically significant.

Table 3: Staggered DiD Regression Specification – Indirect Impact

Dependent Variable: *Log of Monthly rent*

Variable	(8)	(9)	(10)	(11)	(12)
Treatment	0.023 (0.007)***	0.021 (0.007)***	0.010 (0.005)**	0.012 (0.06)**	0.015 (0.006) ***
Year – County – Property Characteristic FE	Yes	Yes	Yes	Yes	Yes
Treated within	100m	150m	200m	250m	300m
N	394,823	394,823	394,823	394,823	394,823

Notes: Table 3 reports estimates of the impact of an institutional landlord purchasing of a property on the rents of non-institutional properties within x metres of the purchased property, where x ∈ {100, 150, 200, 250, 300}. The control group is all units within x-500 metres of properties purchased by institutional investors. The dependent variable is the log of monthly rent. All regressions are performed with Wooldridge's two-way Mundlak (2021) difference-in-difference estimator, with individual group-year effects aggregated into a single treatment effect. Unit and year fixed effects are included in all specifications. Standard errors are in parentheses and significance level displayed as * p<0.1; **p<0.05; ***p<0.01

Tenant preferences for property tend to be location specific. Therefore, the substitutability of other landlords' property with existing units purchased by an institutional investor would be expected to decay with distance, and thus we would expect the indirect impact to decline when using a larger radius as it would include units which see a weaker treatment effect or no treatment effect. This seems to be the case when moving from radii of 100 metres to 200 metres, although afterwards the effect starts increasing again, although not above estimate in the 100 or 150 metre treatment radii specifications.

An indirect effect on rents will likely have a broader impact on the rental market than the direct effect of institutional investment. Within 100 metres of properties owned by institutional landlords, there are roughly 22,000 properties owned by other landlords, with the figure growing to 44,000 within 200 metres and 62,000 within 300 metres. This would correspond to 7%, 14% and 20% of the total stock of Irish private rental property (based on 2016 Census).

Event study specifications of these regressions show that the indirect impact of institutional landlords on other landlords is not immediate (see charts 10, 11 and Appendix A – Charts 17-

19). In most specifications, no treatment effect is significant before relative period 2. Pre-treatment effects are still insignificant in the majority of cases. Importantly, there is no evidence of an increasing positive trend in the run up to treatment.

The dynamics of the treatment effect do not have an obvious interpretation. Unlike the direct impact, the hypothesised indirect impact may take time to appear. Non-institutional landlords are generally hypothesised to be risk averse, so it may take them repeated observations of institutional landlords asking rents to become comfortable with increasing rents themselves. Similarly though, amenity effects and tenant sorting would not be instantaneous.

Chart 10: Specification (8) in event study form

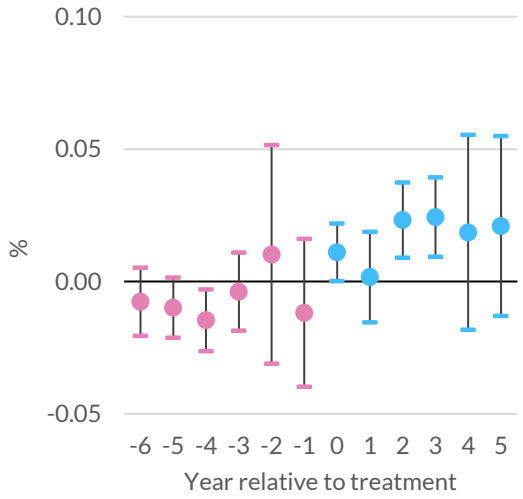

Source: Author's Calculations
Notes: 0 is year of treatment. Circles are treatment effects, bars are 95% confidence intervals. Objects shaded pink are in the pre-treatment estimates, objects shaded blue are estimates of the treatment effect

Chart 11: Specification (12) in event study form

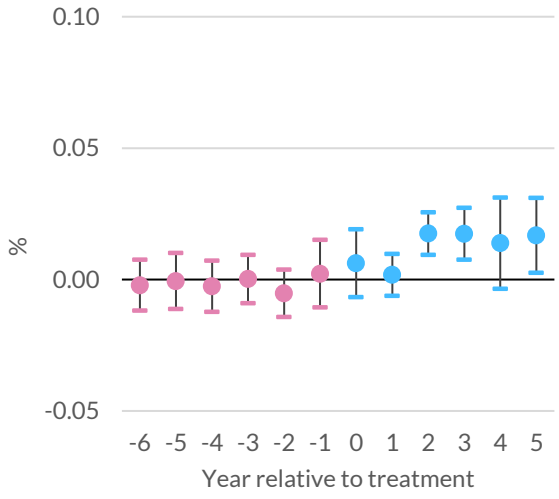

Source: Author's calculations
Notes: 0 is year of treatment. Circles are treatment effects, bars are 95% confidence intervals. Objects shaded pink are in the pre-treatment estimates, objects shaded blue are estimates of the treatment effect

Unlike the results in section 5, the estimates presented are not wholly robust to moving to a more granular measure of geography. Replacing county with electoral district or local electoral areas in the interaction fixed effects results in many specifications losing statistical significance, with effect sizes decreasing, and in a limited number of cases flipping sign.

On this basis, it is difficult to conclude whether the estimated impacts are capturing a causal effect, or differences in the trends of rents between treatment and control units within counties. Further investigation is needed on this to before definite conclusions can be drawn.

7. Conclusion

In conclusion, this paper provides evidence that institutional investors increase rents in excess of other landlords with comparable properties following purchase. The results suggest that institutional landlords increase monthly rents by around 4.1 percentage points more than other landlords. An initial comparison against available data supports many different mechanisms for

why this premium exists, including bargaining power, differing objective functions, and institutional investors offering a higher quality, higher cost service. There is also suggestive evidence that institutional investors may be having an impact on other landlords, but this issue requires further investigation before firm conclusions can be drawn.

Overall, the data and results of the paper suggest that institutional investment in existing stock in Ireland has had a modest direct impact on Ireland's rental market up to 2019. While the additional increase in rent over other landlords are non-negligible, the existing properties they have purchased constitute a small portion of the total rental stock. Furthermore, available data suggests that this observed increase is likely not solely due to bargaining power, but also differences in landlords preferences (i.e. nominal rigidity) and the provision of higher quality rental services (i.e. expenditure on improving properties). This makes it difficult to draw a firm conclusion about what impact this investment has on consumer welfare, and the premium may dissipate over the medium term. Their indirect impact may be more extensive, but its presence is uncertain, and its impact is smaller.

This paper also leaves open future opportunities for investigation. Further investigation of the mechanism underpinning institutional landlords charging a premium would allow us to better understand the welfare implications of this kind of investment.

Bibliography

Amárach Research. (2021, July 14). *Small Landlords Report.* Residential Tenancy Board. Retrieved from Residential Tenancy Board: https://www.rtb.ie/images/uploads/general/RTB_Small_Landlord_Report_July_2021.pdf

Asquith, B., Mast, E., & Reed, D. (2023, April 29). Local Effects of Large New Apartment Buildings in Low-Income Areas. *The Review of Economics and Statistics*, 359-375. Retrieved from SSRN: https://direct.mit.edu/rest/article/105/2/359/100977/Local-Effects-of-Large-New-Apartment-Buildings-in

Austin, N. (2023). Keeping Up with the Blackstones: Institutional Investors and Gentrification. Retrieved from https://afajof.org/management/viewp.php?n=19544

Borusyak, K., & Jaravel, X. (2017, May 8). Revisiting Event Study Designs. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2826228

Butts, K. (n.d.). Retrieved from Ar.

Butts, K. (2021, October 19). Retrieved from arXiv: https://arxiv.org/abs/2110.10192

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200-230. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S0304407620303948

Callaway, B., Bacon-Goodman, A., & Sant'Anna, P. (2021). Difference-in-Differences with a Continuous Treatment. Retrieved from https://psantanna.com/files/Callaway_Goodman-Bacon_SantAnna_2021.pdf

Committee of Public Accounts. (2013, September 26). Retrieved from Oireachtas.ie: https://data.oireachtas.ie/ie/oireachtas/debateRecord/committee_of_public_accounts/2013-09-26/debate/mul@/main.pdf

Committee On Housing And Homelessness. (2016, May 12). Retrieved from data.oireachtas.ie: https://data.oireachtas.ie/ie/oireachtas/debateRecord/committee_on_housing_and_homelessness/2016-05-12/debate/mul@/main.pdf

Conefrey, T., & McIndoe-Calder, T. (2018, April). *Where are Ireland's Construction Workers?* Retrieved from centralbank.ie: https://www.centralbank.ie/docs/default-source/publications/quarterly-bulletins/quarterly-bulletin-signed-articles/where-are-ireland-s-construction-workers-(conefrey-and-mcindoe-calder).pdf

Dáil Debates. (2019, April 9). Private Rented Accommodation. Retrieved from https://www.oireachtas.ie/en/debates/question/2019-04-09/73/

Dáil Éireann. (2016, May 18). Dáil debates. *909(2)*. Retrieved from https://www.oireachtas.ie/en/debates/question/2016-05-18/10/#pq_10

Daly, F. (2016, May 12). Opening Statement by Mr. Frank Daly, Chairman of NAMA, to the Oireachtas Committee on Housing and Homelessness, Thursday 12th May 2016.

Retrieved from https://nama.ie/uploads/documents/Opening_Statement_by_Frank_Daly_to_the_Housing_and_Homelessness_Committee.pdf

Daly, P. (2022, August). Institutional Investment in Housing: Financialisation 2.0 in the Case of Ireland. Retrieved from https://www.ssisi.ie/_files/ugd/463248_4308c7b041b6472ba4c413e3cbbbb9a5.pdf

De Chaisemartin, C., & D'Haultfoeuille, X. (2020). *Difference-in-Differences Estimators of Intertemporal Treatment Effects.* Retrieved from SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3731856

Department of Finance. (2019). *INSTITUTIONAL INVESTMENT IN THE HOUSING MARKET*. Dublin: Department of Finance. Retrieved from https://assets.gov.ie/233991/210a2791-fb63-42e2-bb38-ffb81d4b82e4.pdf

Diamond, R., McQuade, T., & Qian, F. (2019). The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco. *American Economic Review*, *109*(9), 3365-94. Retrieved from https://www.aeaweb.org/articles?id=10.1257/aer.20181289

Dombay, D. (2021, November 14). *Spain takes on private equity landlords as cost of housing soars.* Retrieved from ft.com: https://www.ft.com/content/9ef1eb29-04a5-441f-ac77-f6a0fb7d2d85

DW. (2021, September 27). *Berliners vote 'yes' on property expropriation, but what happens now?* Retrieved from DW.com: https://www.dw.com/en/berliners-vote-yes-on-property-expropriation-but-what-happens-now/a-59070328

Fagan, J. (2013, May 15). *€35.5m for 338 apartments in four Dublin schemes.* Retrieved from Irishtimes.com: https://www.irishtimes.com/business/commercial-property/35-5m-for-338-apartments-in-four-dublin-schemes-1.1392957

Gabor, D., & Kohl, S. (2022). *My Home Is An Asset Class.* Brussels: The Greens/EFA.

Gabor, D., & Kohl, S. (2022). *My Home Is An Asset Class.* Brussels: The Greens/EFA.

Gallin, J., & Verbrugge, R. J. (2019). A theory of sticky rents: Search and bargaining with incomplete information. *Journal of Economic Theory*, *183*(C), 478-519.

Garriga, C., Gete, P., & Tsouderou, A. (2021, February). Investors and Housing Affordability. Retrieved from chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/https://www.ie.edu/faculty/pedro-gete/wp-content/uploads/sites/251/2021/02/2021-02-Garriga-Gete-Tsouderou.pdf

Gleeson, J. (2017, May 23). *A breakdown of Housing Vacancy figures in Ireland*. Retrieved from AIRO: https://airo.maynoothuniversity.ie/news/breakdown-housing-vacancy-figures-ireland

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254-277. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S0304407621001445

Gurun, U. G., Wu, J., Xiao, S. C., & Xiao, S. W. (2022, January 1). Do Wall Street Landlords Undermine Renters' Welfare?

Horgan-Jones, J. (2021, May 18). *Stamp duty will rise to 10% for multiple house purchases under Government plans.* Retrieved from irishtimes.com: https://www.irishtimes.com/news/politics/stamp-duty-will-rise-to-10-for-multiple-house-purchases-under-government-plans-1.4568346

Hyland, M., Lyons, R. C., & Lyons, S. (2013). The value of domestic building energy efficiency — evidence from Ireland. *Energy Economics*, 943-952.

Irish Concrete. (2009, April). *The A-Rated Energy Efficient Concrete Home.* Irish Concrete. Retrieved from Irish Concrete: https://www.irishconcrete.ie/wp-content/uploads/2017/04/2009-A-rated-Homes.pdf

Joint Committee on Housing, Local Government and Heritage. (2023, October 10). Analysis of Private Rental Sector Discrepancies: Discussion. Retrieved from https://www.oireachtas.ie/ga/debates/debate/joint_committee_on_housing_local_government_and_heritage/2023-10-10/2/

Kahn-Lang, A., & Lang, K. (2020). The Promise and Pitfalls of Differences-in-Differences: Reflections on 16 and Pregnant and Other Applications. *Journal of Business & Economic Statistics*, *38*(3), 613-620. Retrieved from https://www.tandfonline.com/doi/full/10.1080/07350015.2018.1546591

Kollewe, J. (2017). *Build-to-rent – the solution to Britain's housing crisis?* Retrieved from theguardian.com: theguardian.com/business/2017/sep/14/build-to-rent-the-solution-to-britains-housing-crisis

Lambie-Hanson, L., Li, W., & Slonkosky, M. (2019, November 27). Institutional Investors and the U.S. Housing Recovery. *FRB of Philadelphia Working Paper No. 19-45*.

Li, X. (2021). Do new housing units in your backyard raise your rents. *Journal of Economic Geography*, 1309-1352.

Lyons, R. C. (2021). *Institutional Investment and the Private Rental Sector in Ireland.* Irish Institutional Property .

Lyons, R., & Muellbauer, J. (2013). Explaining the bubble: House prices, user-cost and credit conditions in Ireland, 1975-2012. *Working Paper*.

Mulligan, J. (2022, February 12). *Revealed: Ten biggest landlords now own 17,000 homes.* Retrieved from Independent.ie: https://www.independent.ie/irish-news/revealed-ten-biggest-landlords-now-own-17000-homes-41339550.html

O'Donovan, D. (2012, June 7). *Firm linked to the Kennedys snaps up 'Gasworks' building for €40m.* Retrieved from Independent.ie: https://www.independent.ie/business/irish/firm-linked-to-the-kennedys-snaps-up-gasworks-building-for-40m-26861939.html

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

Pennington, K. (2021, June 15). *Does Building New Housing Cause Displacement?: The Supply and Demand Effects of Construction in San Francisco.* Retrieved from SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3867764

Raymond, E., Duckworth, R., Miller, B., Lucas, M., & Pokharel, S. (2016, December 12). Corporate Landlords, Institutional Investors, and Displacement: Eviction Rates in Single Family Rentals. *Community & Economic Development Discussion Paper 04-16*.

Rebuilding Ireland. (2016). *Strategy For The Rental Sector.* Dublin: Government of Ireland.

Residential Tenancies Board. (2021). *Rent Index Q1 2021.* Dublin: Residential Tenancies Board.

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175-199. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S030440762030378X

Walsh, D., & NicAongusa, B. (2016). *Ensuring Delivery of Build-To-Rent Housing Projects.* Retrieved from opr.ie: https://www.opr.ie/wp-content/uploads/2019/08/2016-Build-To-Rent-Housing-Projects-1.pdf

Wooldridge, J. M. (2021, August). *Two-Way Fixed Effects, the Two-Way Mundlak Regression, and Difference-in-Differences Estimators.* Retrieved from https://papers.ssrn.com/: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3906345

Appendix A

Table 4: Summary of Sample Restrictions and Impact

	N	Difference
Initial Sample	1,281,491	
Remove missing and inaccurate Eircodes	1,114,969	166,522
Remove years outside sample range	1,074,177	11,265
Remove observations with rents beyond outlier range	1,062,913	27,830
Remove student accommodation	1,035,083	64,861
Removing units that appear only once and duplicates	839,204	195,879
Removing units that don't work for the model (e.g. not yet treated units in year of treatment, treated units where treatment has switched off, etc.)	834,860	4,344

Table 5: Summary Statistics for Institutional and Non-Institutional Landlords in matched sample

	Institutional	Non-Institutional	Difference
Share in Dublin	0.82	0.73	0.09 ***
No. Bed	1.95	1.98	0.03 ***
Share RPZ	0.18	0.13	0.05 ***
BER B3 or Above	0.42	0.10	0.32 ***
Decade of Construction >= 2000	0.81	0.50	0.31 ***
N	25,576	116,171	

Table 6: Regressions Exploring Causal Mechanisms

Variable	(13)	(14)	(15)	(16)
Treatment	0.01 (0.01)	0.01 (0.00)*	30 (2.8)***	-0.12 (0.03)***
Year Dummies	Yes	Yes	Yes	No
Dublin Dummy	Yes	Yes	Yes	No
Year – County – Property Characteristic FE	No	No	No	Yes
N	1,233,489	257,487	12,418	582,209

Notes: Table 6 reports estimates of various regressions underpinning results in section 5.3. Specification (25)-(27) are cross-sectional regressions that take their data from daft.ie online rental ads. Specification (25) is a linear probability model with a dependent variable that equals 1 if an ad for a property has a different final rent to its initial rent, and 0 otherwise. Specification (26) focuses on a sample of rent decreases, and regresses institutional investor ownership status on the % decrease between initial and final rents. Specification (27) focuses on a sample of rent decreases with data on the length of time it takes to decrease rents, and regresses institutional investor ownership status on the number of days it takes for the rent on an ad to first be decreased. Specification (28) is a linear probability difference-in-differences model, regressing institutional investor ownership on nominal rigidity. It uses Wooldridge's two-way Mundlak (2021) difference-in-difference estimator, with individual group-year effects aggregated into a single treatment effect. The dependent variable is equal to 1 if rents have not changed since the previous contract and 0 otherwise. Standard errors are in parentheses and significance level displayed as * p<0.1; **p<0.05; ***p<0.01

Table 7: Robustness to BER improvements

Dependent Variable: Log of Monthly rent

Variable	(17)	(18)	(19)	(22)
Treatment	0.047 (0.014)***	0.048 (0.014)***	0.055 (0.014)***	0.057 (0.014)***
Rent Pressure Zone	Yes	Yes	Yes	Yes
Relative rent	Yes	Yes	Yes	Yes
Year – County – Property Characteristic FE	Yes	Yes	Yes	Yes
BER Upgrade	No	Yes	No	Yes
N	128,074	128,074	21,691	21,691

Notes: Table 2 reports estimates of the impact of purchase by an institutional landlord on a residential property's rent. The dependent variable is the log of monthly rent (€ per month). All regressions are performed with Wooldridge's two-way Mundlak (2021) difference-in-difference estimator, with individual group-year effects aggregated into a single treatment effect. Specifications (1)-(2) use the full sample, while specifications (3)-(4) use the sample matched on property characteristics and distance (described in section 4.2). Rent Pressure Zone is a control for the introduction of rent pressure zone legislation at the local electoral area level, and is also a difference-in-difference estimator. Lagged relative rent measures the log of the previous rent of a property relative to the average of similar properties in its local electoral area. BER upgrade measures whether a units BER has been upgraded. Standard errors are in parentheses and significance level displayed as * p<0.1; **p<0.05; ***p<0.01.

Chart 12: Specification (1) in event study form

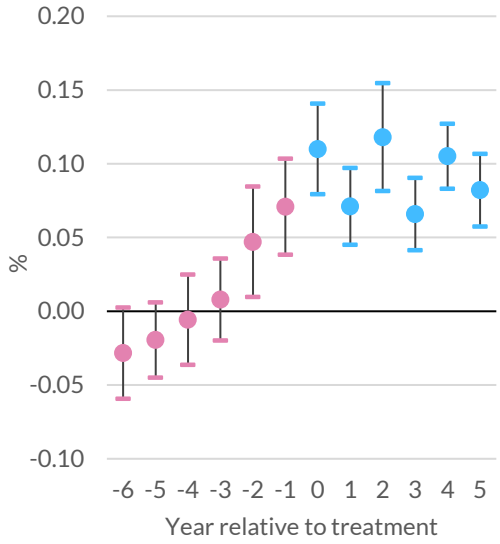

Source: Author's Calculations

Chart 13: Specification (2) in event study form

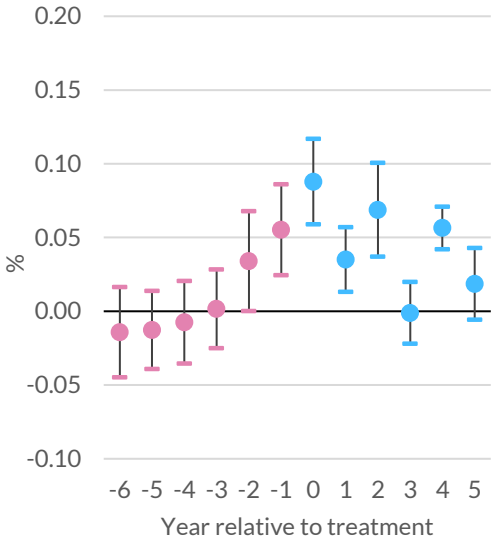

Source: Author's calculations

Chart 14: Specification (3) in event study form

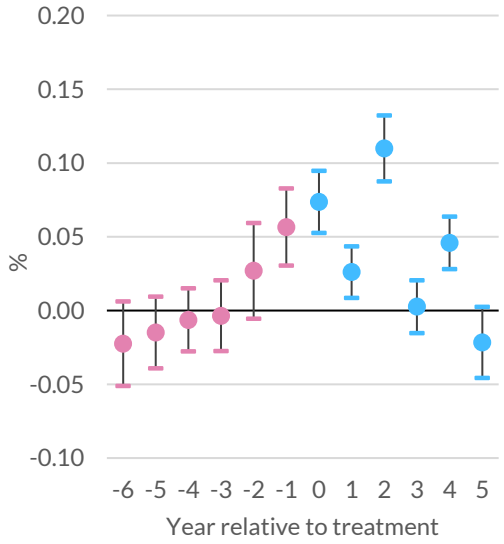

Source: Author's Calculations

Chart 15: Specification (4) in event study form

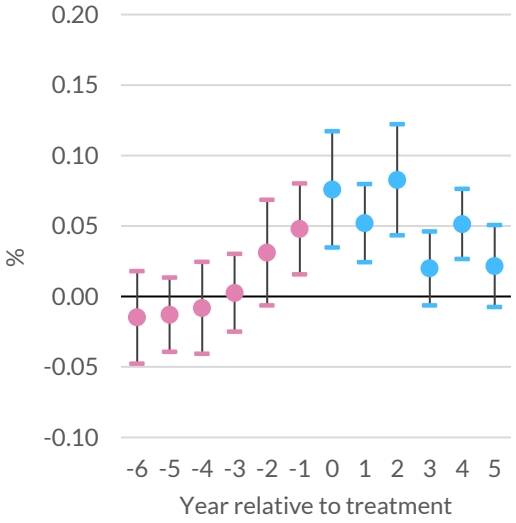

Source: Author's calculations

Chart 16: Specification (5) in event study form

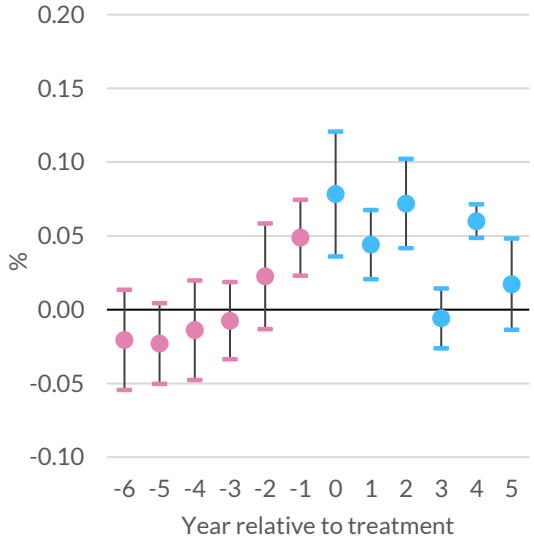

Source: Author's Calculations

Chart 17: Specification (9) in event study form

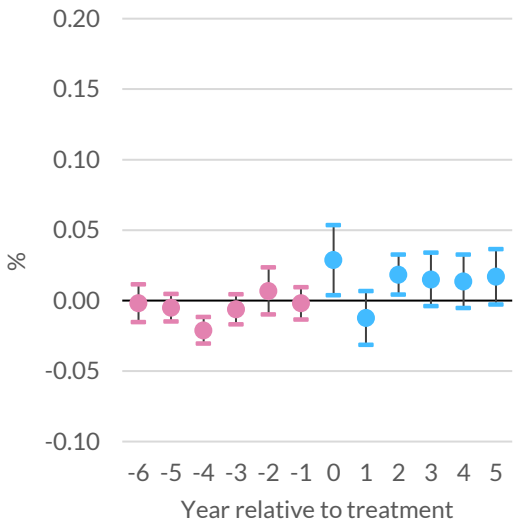

Source: Author's calculations

Chart 18: Specification (10) in event study form

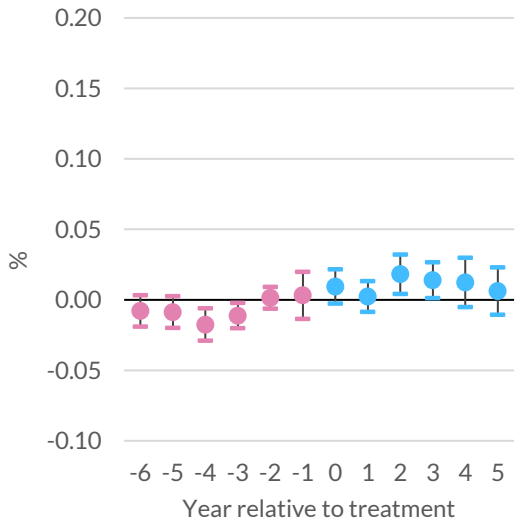

Source: Author's Calculations

Chart 19: Specification (11) in event study form

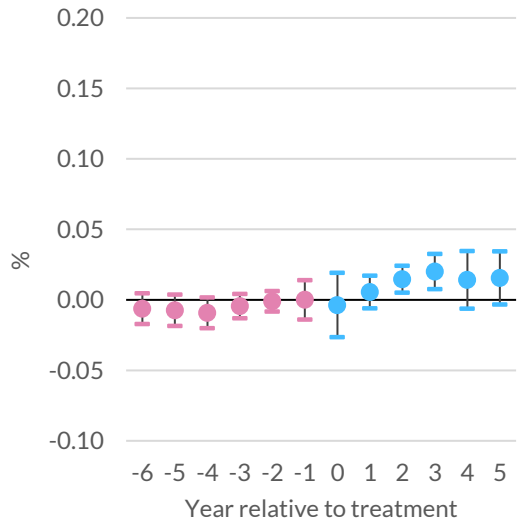

Source: Author's calculations

Chart 20: Specification (6) with a 250m matching radius in event study form

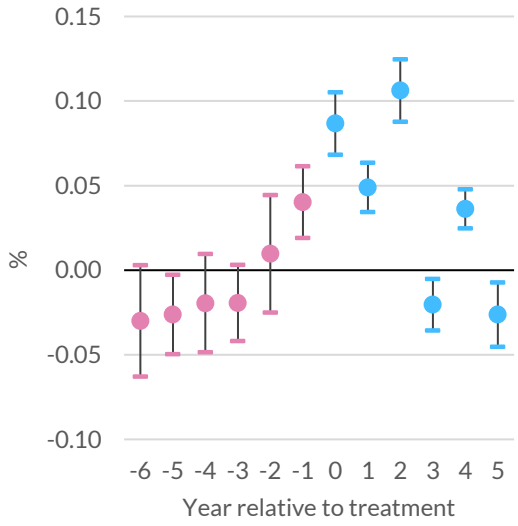

Source: Author's Calculations

Chart 21: Specification (6) with a 750m matching radius in event study form

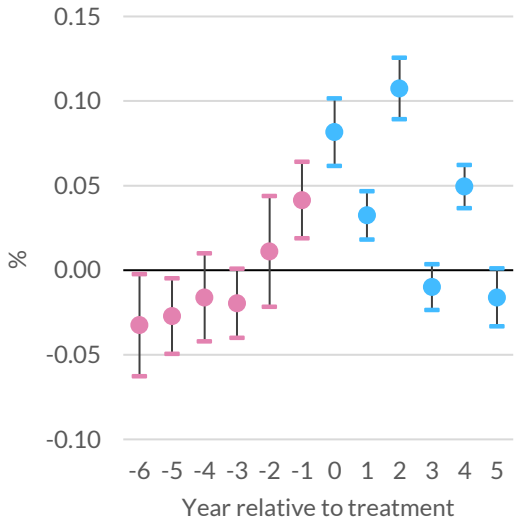

Source: Author's calculations

Chart 22: Specification (6) with a 500m matching radius on bedroom and small area decade of construction in event study form

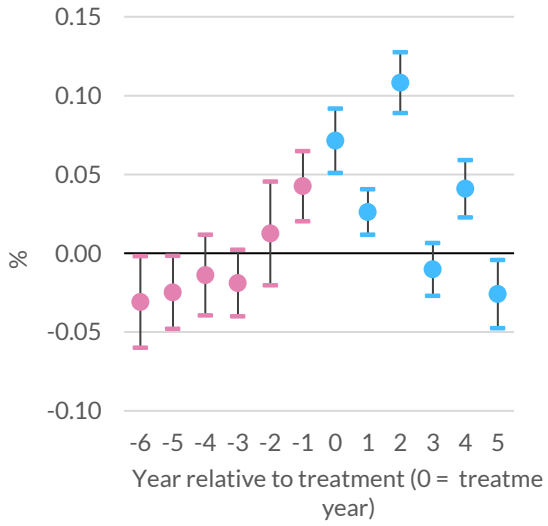

Source: Author's Calculations

Chart 23: Specification (6) with a 500m matching radius on bedroom and small area BER in event study form

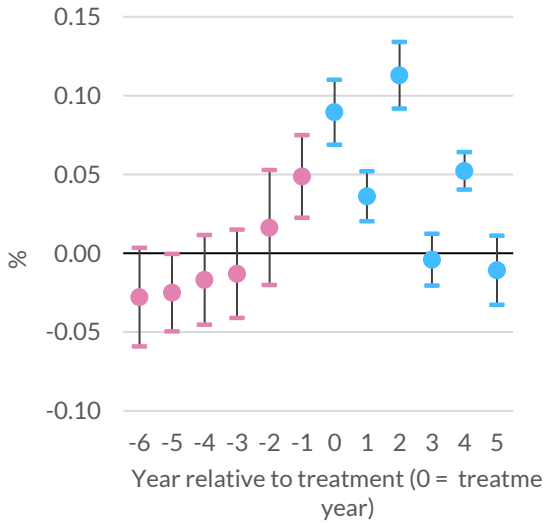

Source: Author's calculations

T: +353 (0)1 224 5800
E: publications@centralbank.ie
www.centralbank.ie

Banc Ceannais na hÉireann
Central Bank of Ireland
Eurosystem