Buy-to-Live vs. Buy-to-Let: The Impact of Real Estate Investors on Housing Costs and Neighborhoods*

Marc Francke Lianne Hans Matthijs Korevaar Sjoerd van Bekkum

June 15, 2023

Abstract

How do buy-to-let investors impact local housing markets and the composition of neighborhoods? We investigate this question by examining a Dutch legal ban on buy-to-let investments, exploiting quasi-experimental variation in its coverage. The ban effectively reduced investor purchases and increased the share of first-time home-buyers, but did not have a discernible impact on house prices or the likelihood of property sales. The ban did increase rental prices, consistent with reduced rental housing supply. Furthermore, the policy caused a change in neighborhood composition as tenants of investor-purchased properties tend to be younger, have lower incomes, and are more likely to have a migration background. Our results suggest rental investors influence local housing conditions primarily through changing the residential composition of neighborhoods rather than direct house price effects.

*Francke: Amsterdam Business School, University of Amsterdam; Ortec Finance, Amsterdam (m.k.francke@uva.nl). Hans: Kadaster (Lianne.Hans@kadaster.nl). Korevaar: Erasmus School of Economics, Erasmus University Rotterdam (corresponding author: korevaar@ese.eur.nl). Van Bekkum: Erasmus School of Economics, Erasmus University Rotterdam (vanbekkum@ese.eur.nl).
Notes: This working paper evaluates the effect of a housing policy introduced in The Netherlands on January 1, 2022. This paper is evolving and will be updated as additional data becomes available; this is the first public version of the paper. Results are based on calculations by the authors using non-public microdata from Statistics Netherlands. We especially thank the Municipality of Rotterdam for providing funding for access to this data and for valuable discussions on the details of its policy. We also thank seminar participants at the NYC Real Estate Conference and the Erasmus School of Economics for valuable comments. Korevaar has been funded by a Marie-Sklodowksa Curie Fellowship from the European Commission (grant 101.028.821).

1 Introduction

Real estate property is typically acquired by either owner-occupiers intending to live in them or by investors intending to resell or rent them out. The market share of investors relative to owner-occupiers may vary significantly over time and across regions. Notably, since the 2007 mortgage crisis, buy-to-let investor activity has surged in most developed economies. Previous work suggests that investor activity supported local housing markets after the bust of 2008 (Garriga et al., 2021; Lambie-Hanson et al., 2022). However, as economic conditions improved and the housing market tightened, concerns grew about investors driving up home prices, pricing out first-time home-buyers, and decreasing neighborhood livability.

Despite a wide-spread, global articulation of this idea,[1] little well-identified empirical evidence exists regarding the extent to which homeownership versus investor ownership affects local housing costs, housing market transactions, and neighborhood change. In this paper, we aim to fill this gap using new microdata on all individual housing transactions as well as the individual owners and residents of these properties. The data contains information on the property holdings of any natural person and entity in the Netherlands, as well as tax valuations of each property. It can identify for each transaction whether the buyer is an owner-occupier or a buy-to-let investor. Furthermore, we can match property transactions to the individual occupants of these properties and their socioeconomic characteristics and incomes. We also gather data on rental prices.

We apply the data to a recent policy change in the Netherlands that enables municipalities to ban investors from buying properties for rental purposes if these are below a predetermined tax value (in Dutch: *Opkoopbescherming*, or “purchase protection"). The policy ban led to a reduction in properties bought by investors while increasing the

[1]For instance, in the U.S.: "Large investor purchases of single-family homes and conversion into rental properties speed the transition of neighborhoods from homeownership to rental and drive up home prices for lower-cost homes, making it harder for aspiring first-time and first-generation home buyers, among others, to buy a home. At the same time, these purchases are unlikely to meaningfully boost supply in the lower-cost portions of the rental market, as investors charge more for rent to recoup higher purchase costs." (https://www.whitehouse.gov/briefing-room/statements-releases/2021/09/01/fact-sheet-biden-harris-administration-announces-immediate-steps-to-increase-affordable-housing-supply/)

number of first-time homebuyers. In a country with housing in short supply, removing investors from the market did not significantly impact house prices or the likelihood of selling property. At the same time, our evidence does point towards a positive impact on rent prices, consistent with the effect of the policy on rental supply. More importantly, banning buy-to-let investors contributes to gentrifying treated neighborhoods: average income increased as owner-occupying households have significantly higher incomes than tenants in investor-purchased property; are significantly older; and are less likely to have a migration background. Thus, the ban has successfully increased middle-income households' access to homeownership, at the expense of buy-to-let investors. However, the policy also drove up rents in affected neighborhoods, thereby damaging housing affordability for individuals reliant on private rental housing, undermining some of the intentions of the law.

Although legislation for the investor ban was drawn up at the national level, municipalities can decide whether to implement the law, and if so, select the area and predetermined tax value below which buy-to-let activity is banned. We initially focus on the city of Rotterdam, the second-largest city of the Netherlands that was the first to introduce the policy and also the only major city to apply the ban only to specific neighborhoods. This provides us with discrete, predetermined, and detailed differences in investor activity between treated and untreated neighborhoods. Hence, we can compare observationally equivalent properties below the tax value cap in treated neighborhoods to those in adjacent neighborhoods, before and after the introduction of the policy and adjusting for detailed property characteristics, including the tax value of the property. Indeed, we find parallel trends in investor purchases of regulated and unregulated properties before the introduction of the policy and a meaningful decrease in purchases thereafter.

We then extend our analysis to the country level by exploiting substantial differences in implementation dates and coverage across municipalities that are exogenous to local housing market trends. We obtain difference-in-difference estimates at the national level by comparing the price of properties in regulated areas just above and below the locally binding price cap, as well as comparing the price evolution in cities that did or did not introduce the law during 2022. These results are consistent with the Rotterdam case.

A key challenge in the literature on this subject is that changes in homeownership tend to be endogenous to developments in the local housing market, and investor activity is hard to identify. Furthermore, measuring the impact of homeownership on the residential composition of neighborhoods is difficult because transitions from ownership to rentals within neighborhoods are rarely exogenous. A key advantage of our setting is that the ban shifted the composition of buyers, implying it also led to an exogenous shift in residents from renters to owner-occupiers.

To the best of our knowledge, our paper is the first to explicitly study how an exogenous shift in homeownership affects the composition of neighborhoods. Our experiment demonstrates how residents change if the marginal property available for sale is bought by an investor versus an owner-occupier, while holding the housing stock fixed. Such transitions are the main way through which local homeownership rates change over time. In this paper, we focus on establishing short-term effects. Over longer horizons, the effects of accumulated rent-own transitions over time on neighborhood composition might endogenously affect both house prices and neighborhood composition, giving rise to more complicated dynamics. We abstain from this with our short-term setting: over the one-year horizon we can study so far, only a small fraction of the housing stock trades hands.[2]

The investor ban reduced investor purchases on regulated properties by about 75%, equivalent to a 23 percentage point reduction in the share of properties bought by investors. We document a similar effect at the national level, although smaller in absolute magnitude (10 percentage points) as Rotterdam regulates properties with comparatively high ex-ante investor activity. In line with the decline in investor purchases, we find a significant increase in the share of purchases by first-time buyers, indicating that a large share of properties that would have become rental properties without the ban are now occupied by first-time home-buyers. We find no clear evidence for spillovers of the

[2] In Rotterdam, a back-of-the-envelope calculation suggests the policy prevented about 1.5% of the treated owner-occupied housing stock from being bought by investors in 2022.

policy: non-regulated properties nearby regulated properties do not see economically or statistically meaningful increases in investor purchases activity after a ban is instated.

While the ban intends to shut out private investors, our results offer guidance in understanding the impact of investors on property prices. If investor presence always causes higher property prices, as suggested by commentators and politicians, the removal of investors from the regulated property market should have a negative impact on house prices. However, for properties under the price cap, we find an insignificant treatment effect of 0.1% in Rotterdam (95% confidence interval: [-1.5,+1.7]), with confidence intervals at the national level even closer to zero (95%: [-0.4,+0.8]). We identify slightly lower sale rates for properties subject to the policy post-implementation, but this effect seems largely driven by anticipation effects pre-policy. By contrast, we find that the policy significantly inflated rental prices in Rotterdam with an estimated treatment effect of 4%, in line with a reduction in rental housing supply. Collectively, these findings suggest that rental investors influenced local housing conditions primarily through other mechanisms than through inflating housing prices, at least in the first year around the implementation of the investor ban.

One possible explanation for the absence of a price effect could be that the (unobserved) reservation prices of investors and regular home-buyers are similar. In such a case, investors might still successfully secure properties by leveraging superior negotiation skills, unconstrained borrowing capacity, or alternative strategies, for example by bidding without resolutive conditions, which minimizes risks for sellers (e.g. Cohen et al., 2022). Another explanation is that the decrease in buy-to-let investment due to the ban led to increased demand from owner-occupiers. The policy was primarily introduced due to the pervasive perception that buy-to-let purchases adversely affected neighborhood livability, which would likely impact prices too. By purchasing property in an area with a ban, owner-occupiers could insure themselves against these trends.

To study such effects, we examine how the shift in buy-to-let investment affects the composition of treated neighborhoods, where properties that would have been sold to investors without the ban are now acquired by owner-occupiers. This alters neighborhood composition to the extent that renters differ from owner-occupiers in comparable properties. Our difference-in-differences results reveal that following the policy's introduction, households moving into transacted properties subject to the ban belonged to income brackets that were 3 percentiles higher than those of control properties. Nationally, this impact is 2 percentiles. We demonstrate that this effect is entirely due to households in investor-bought properties having substantially lower incomes than those in owner-occupied properties. In Rotterdam, renter households end up 19 percentiles to the left in the country's income distribution compared to owner-occupied households in observationally equivalent properties. The national level difference is 15 percentiles. Consequently, while the ban-induced shift in purchases from investors to owner-occupants has enhanced homeownership access for middle-income households, it has curtailed housing opportunities for lower-income renter households.

The policy also influenced neighborhood composition on other fronts. Compared to observationally equivalent owner-occupied properties, we found that residents of regulated investor-owned properties are, on average, three years younger than owner-occupiers and about 20 percentage points more likely to not hold a non-Dutch nationality. Investor-owned properties also accommodate a higher number of residents, and these residents are 30 percentage points more likely to relocate within two years of moving in, more than three times the rate of owner-occupiers. This suggests that over extended periods, shifts in the rates of investor purchases relative to owner-occupiers can substantially influence neighborhood composition, even when the housing stock remains static.

1.1 Related literature

This paper most directly links to an emerging literature in both finance and real estate that studies the impact of investors on housing markets. Recent literature has suggested that investors played a role in the post-crisis recovery of the housing market (e.g. Garriga et al., 2021; Lambie-Hanson et al., 2022). The paper most closely related to ours is Ater et al. (2021), who study the impact of a temporary reduction in the capital gains tax for certain tax-exempted investors on house prices and find that the induced sales by these investors reduced house prices.

While some of the differences in our results vis-a-vis these studies might be explained by differences in identification, it is also possible that the price impact of investors could vary depending on the nature of the shock and the state of the housing market. At the time of the policy, house prices were very high and there was a large number of potential buyers looking for properties, implying there remained many non-investor buyers after the introduction of the policy. On the other hand, the Israeli policy induced investors to quickly sell before the exemption would end, while the U.S. based studies focus on the impact of investors in a period where there were many properties for sale and few other buyers. Investors might have more price impact if they step into the market when no one else is willing to buy. Thus, while we do not rule out that Dutch investors had a significant impact on prices when they entered during the downturn in the 2010s, our causal evidence shows this at least did not persist until the 2020s.

Following the crisis, various papers have also examined the role of speculative investors in driving the housing boom and bust, e.g. Chinco and Mayer (2015); Gao et al. (2020); Bayer et al. (2020, 2021). We focus here on a different type of investment: speculative investment is typically signaled by individuals with short-holding periods, and those are not frequent in the Netherlands, with most investors making purchases as long-term investments. Such rental investments have also become increasingly prominent in various other countries in recent years, with a growing literature attempting to document and measure the impact of growing institutional and private investment in residential real estate (e.g. Allen et al., 2018; Mills et al., 2019; Bracke, 2019; Lambie-Hanson et al., 2022; Austin, 2022; Gargano and Giacoletti, 2022; Garriga et al., 2023; Gurun et al., 2023). Our key advantage relative to these papers is that we can exploit quasi-experimental variation in investor activity, as well as highly-detailed data on properties and their owners and residents. This allows us to precisely pin down the effects of investor activity on housing markets.

Our findings also relate to the literature measuring the impact of homeownership on neighborhoods. Two recent papers study the impact of right-to-buy schemes, focusing on household-level outcomes (Sodini et al., 2016) and on neighborhood outcomes (Hausman et al., 2022). In these programs, incumbent tenants have the opportunity to buy their property at a highly discounted price, implying that the effect of homeownership is measured while holding the property and resident constant. In our paper, we also hold the property constant: rent-to-own transitions based on market transactions result in changes in neighborhood composition as long as the marginal renter differs from the marginal homeowner. Our estimates are thus calculating the marginal effect of one rent-to-own conversion on tenant composition, and we demonstrate that a policy that stimulated homeownership led to a significantly larger entry of high-income residents. More broadly, this compositional effect can be a key channel to understand the price impacts and neighborhood effects of housing policies that change the bargaining position of renters relative to owners. Various literature has documented that rent controls (Glaeser, 2003; Autor et al., 2014, 2019; Diamond et al., 2019) or changing homeownership rates (e.g. Coulson and Li, 2013; Ihlanfeldt and Yang, 2021) can have substantial effects on prices and neighborhoods. Our paper suggests that such changes might occur because these policies alter the composition of residents. If lower-income residents move elsewhere if rent controls are abolished or rental investment is limited, as suggested in our paper, policies might primarily lead to a redistribution of individuals with more limited effects on general welfare.

Finally, the impact that we find of investors on neighborhood composition links to the emerging literature on gentrification, both internationally (e.g. Couture et al., 2019; Couture and Handbury, 2023) and within the Netherlands (e.g. Hochstenbach and Musterd, 2018). We specifically examine the role of buy-to-let investors in this process, and find that their activity increases the share of low-income households within neighborhoods. LaPoint (2022) suggests investors can contribute to gentrification by purchasing delinquent property and converting these to more luxurious properties. Both findings could be rationalized by the effect of investor activity on neighborhoods and their residents: we find that on average buy-to-let activity leads to a higher share of households with lower incomes. Although such data is not available for the US, it is likely that delinquent homeowners with under-maintained properties are likely poorer than the residents who replace them in rebuilt, more luxurious properties.

2 Data

To assess the impact of the policy on transaction prices and the number of sales, we use data from the Dutch Land Registry (DLR; *Kadaster*), which records all housing transactions in the Netherlands. The data cover the period from 2009 to 2022, although our analysis focuses on sales from 2021 and 2022. The Land Registry only covers sales of existing properties, thus excluding newly constructed properties, which are unaffected by the law change.

At the settlement of each property transaction, the DLR categorizes the buyer. We focus on owner-occupants and home-buyers. A transaction is classified as an investor purchase if the buyer is a non-natural person or if the buyer is a natural person who owns multiple properties and does not reside in the property.[3] In addition to the transaction price and date, the DLR data also contain hedonic property characteristics including the year of construction, the number of square feet of the property, the type of property, and neighborhood. We use these characteristics as control variables.

We combine the DLR data with administrative data from Statistics Netherlands, which provides us with information on household income, person-level characteristics including residence, and the tax values of 2022 and 2021 for all individual properties in

[3]For individuals who own two homes at the end of the sample, we cannot perfectly identify whether a property is an investment property or owner-occupied because residential data is only available until March 31, 2023. If an owner-occupier wants to move and takes ownership of the new property before selling the old property, this transaction is classified as an investment property until the owner-occupier registers at the new property. We verify that excluding these sales does not qualitatively affect any of our results.

the Netherlands. The tax value of 2022 (2021) reflects the property's market value on January 1, 2021 (2020) that is determined by municipalities based on detailed hedonic property characteristics and realized sales prices up to three quarters after the valuation date. These valuations are generally highly accurate, with tax values being able to explain over 90% of the country-wide variation in house prices.[4]

We link each housing transaction to its tax value and determine for each transaction whether a property is subject to the reform. We remove sales where more than one property or parcel is being sold (3.5% of sales), as we only observe the sale price for the entire transaction. In total, this leaves us with 2.07 million sales between 2009–2022.

To link to the income and other characteristics of residents of these properties, we use data on the residences and moves of the entire Dutch population, available until March 31, 2023. We only include individuals who actively registered at a property after it was sold. We also only include data for residents that are over 18 years old on December 31, 2022 (born before 2005).

3 The investor ban

The 'opkoopbescherming' law was announced in the summer of 2021, but details about how the policy would be implemented locally only materialized from November onwards. The law became effective on January 1, 2022, allowing municipalities to implement regulations requiring investors to obtain a permit for leasing property that was not leased or leased for less than 6 months at the time of purchase. This requirement applies for four years after the purchase and only covers properties bought after the law was enforced within a municipality. According to the law, a permit system can only be introduced if the municipality can justify that it is "necessary and suitable for combating the scarcity of cheap and mid-priced owner-occupied housing[,] or for the livability of the local envi-

[4] Tax values reflect market values if the full and unencumbered ownership is transferred and the buyer can take possession of that immediately and completely. The first condition means that homes with leaseholds are valued as if they were fully owned. The latter condition implies that rental properties are valued as if they were sold in the owner-occupier market.

ronment."

Because there is no clear definition of what is "necessary and suitable," municipalities have considerable freedom in how they implement the law. They can deny every investor a permit unless it falls under one of the three exceptions mandated by the national law: leasing to first- or second-degree relatives, rentals where the owner returns to live in the property within 12 months, or a house tied to a business. All municipalities that have implemented the law so far do not issue any renting permits to investors for properties covered under the law, although some grant exceptions for non-profit investments that serve an explicit societal purpose. The law only covers existing properties, as municipalities already had the legal tools to prevent investors from buying newly-constructed property designed for owner-occupancy.

Municipalities that have introduced the policy or plan to do so in the future have gradually rolled out their proposals for implementation since November 2021. Since January 1, 2022, these policies have gradually been put in place. All cities determine whether the law applies at the property level based on the assessed tax value that is determined annually. After the introduction of the regulation locally, any property with a value below the cap cannot be leased for four years unless the owner falls under an exception. This effectively bans buy-to-let housing for affected properties: investors are unlikely to be willing to buy property and leave it vacant for four years when they can buy unaffected housing elsewhere.

Table 2 contains a list of 23 large or mid-sized cities that have introduced a policy in 2022. The second-largest city in the country, Rotterdam, was the first to implement it and opted for a policy that covers only certain areas in the city, leading to strong spatial discontinuities in policy coverage with comparable regulated and unregulated properties just next to each other. Additionally, Rotterdam was the first and only city to immediately implement the law when it came into effect on January 1, 2022, and did so only six weeks after announcing it. This implies that anticipation effects are more limited compared to other cities. Therefore, we focus our analysis on Rotterdam but present country-level results for reasons of external validity and to obtain more statistical power.

3.1 Investor Ban: City of Rotterdam

Of the 71 neighborhoods in the city of Rotterdam, 16 are designated as regulated, containing about 30% of Rotterdam's housing stock. Rotterdam designated these neighborhoods based on the number and fraction of investment rental properties below the price cap of 355,000 euros, implying they generally contain a larger fraction of investment rental properties. In designated neighborhoods, over 90% of properties fall under the price cap. We focus solely on sales of properties with a tax value below the cap, comparing regulated neighborhoods and neighborhoods immediately adjacent to these neighborhoods.

Figure 1 shows that the neighborhoods are spread throughout the Rotterdam urban area, with strong spatial discontinuities in the policy at neighborhood borders.[5] Most neighborhoods covered by the policy are located just outside the city center but away from the limits of the Rotterdam municipality. There is considerable heterogeneity in the prevalence of investor purchases amongst purchase-protected areas. We will use neighborhood and month-level fixed effects to control for any unobserved intertemporal and geographic differences that potentially explain such heterogeneity.

In Figure 2, we plot the trends in properties purchased by investors and sold by owner-occupiers for both regulated (in red) and unregulated properties (in blue, dashed) in Rotterdam, up to the first twelve months after the introduction (shaded in grey). Importantly, while regulated properties tend to have a higher ex-ante likelihood of being bought by an investor, the trends in investor purchases over time do not vary across treated and untreated properties. Before the introduction of the reform, trends in the number of purchases follow perfect parallel trends. This holds also if we look at a more granular level, and it also holds nationally. In Appendix Figure 6, we show monthly counts of regulated and unregulated property in the entire Netherlands.

While buy-to-let activity was growing in the 2010s and was accommodated by rent

[5] Figure 1 uses all purchases in 2021; the Appendix provides similar figures for 2020 and 2022.

Figure 1: Policy Coverage and 2021 investor purchases, Rotterdam

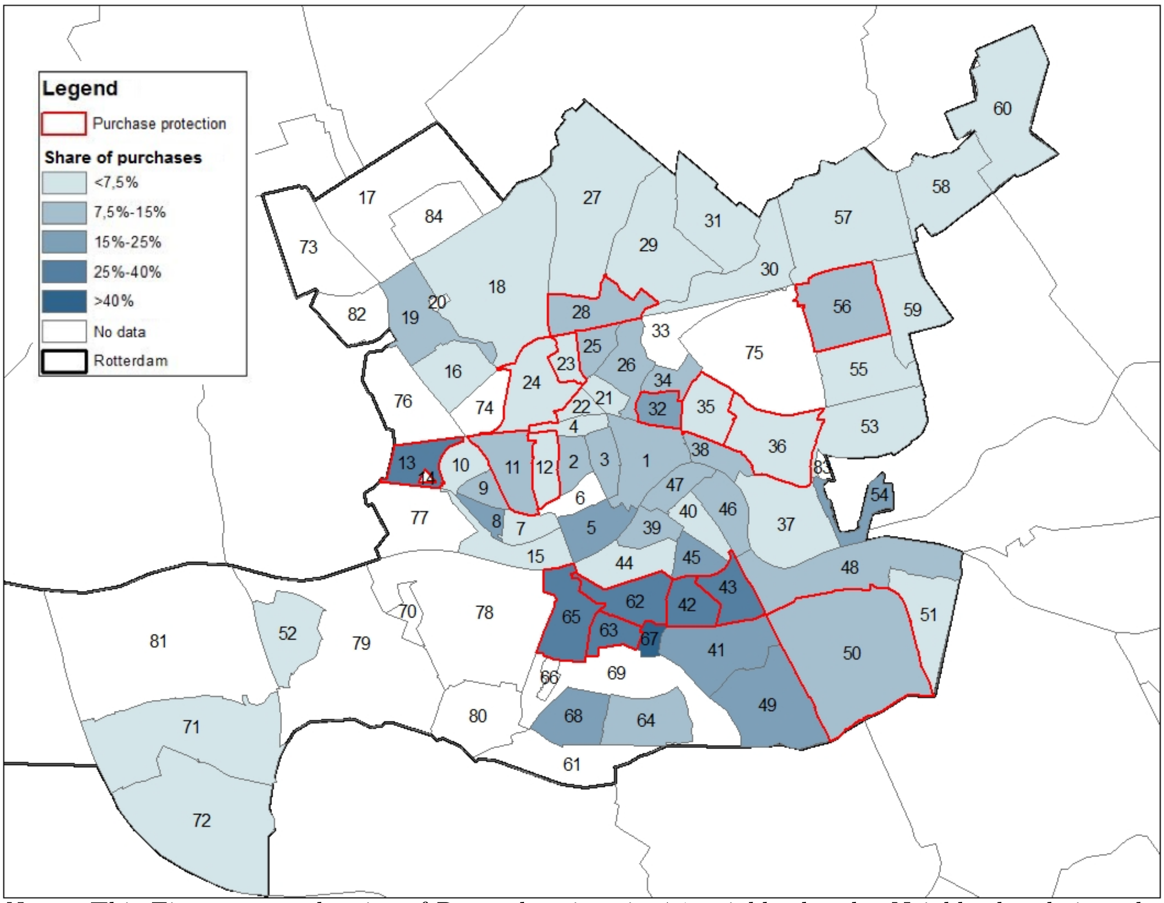

Notes: This Figure maps the city of Rotterdam into its 71 neighborhoods. Neighborhoods in red are affected by the policy. For each neighborhood, shades of blue indicate the fraction of investor purchases (excluding second homes).

deregulation, the policy stance towards investors changed substantially in the 2020s (see Hochstenbach, 2023). Three spikes in Figure 2 stand out, and these relate to policy reforms specifically aimed at buy-to-let purchases: the 2022 investor ban and the 2021 and 2023 increase in stamp duty. The spike in December 2020 coincides with an increase in stamp duty from 2% to 8% for any property that would not be occupied by the owner, and a reduction from 2% to zero for owner-occupiers under the age of 35. This increase was specifically targeted at long-term investors: there are exceptions for properties that are resold within six months. There is also a smaller spike in December 2022, which relates to anticipation of further increases in the stamp duty for investors to 10.4% by January 1, 2023.

The third and final spike in December 2021 most likely reflects the anticipation of the

Figure 2: Investor purchases, Rotterdam, regulated and rest-of-city

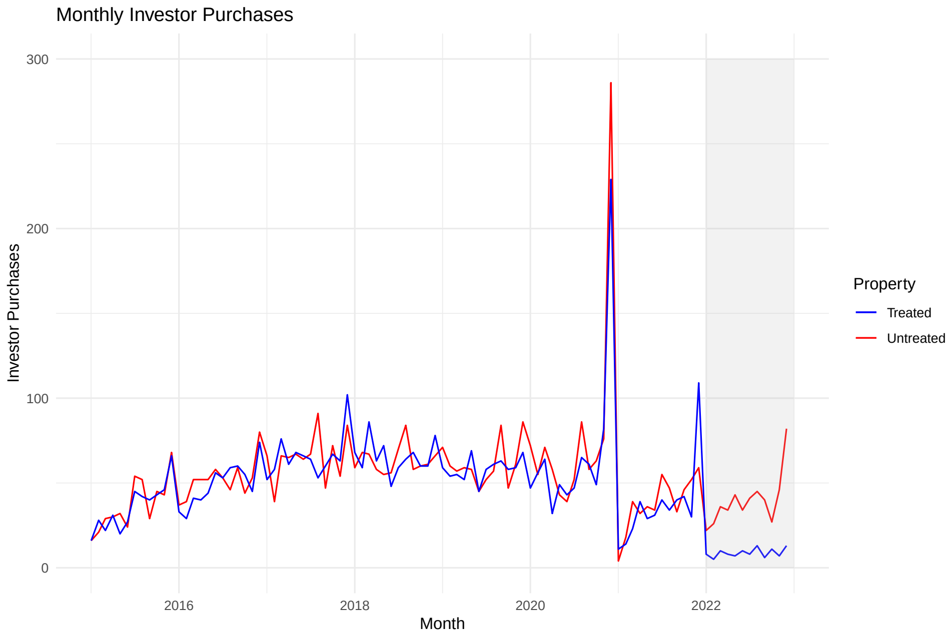

Notes: This Figure plots investor purchases over time in Rotterdam of properties purchased by investors and sold by owner-occupiers for both regulated (in red) and unregulated properties (in blue, dashed).

investor ban by investors ensuring the property could be leased if it were to fall under the regulation. The January 1, 2022 implementation date of the law was announced in September 2021, but investors did not know which properties would be covered until mid-November. To avoid contaminating our treatment effect measurements, we exclude December 2021 from the sample in the main analysis, and exclude the first three months of 2021 when investor activity was impacted by regulatory changes in the stamp duty.

Table 1 provides a univariate sneak peak at the results to come. Sale prices are significantly higher in both regulated and adjacent areas in 2022 due to high price growth throughout 2021 and early 2022. On average, sale prices and tax values are about 10% higher in adjacent areas. Investors tend to purchase cheaper properties, and areas with cheaper properties consequently have higher proportions of investors. Prior to the regulation, 35% of properties in treated areas sold by owner-occupiers with tax values below the price cap were bought by investors. In adjacent areas, this percentage is only 18%. The share in treated areas drops to 10% after the regulation.

The decrease is smaller for commercial investors and owners of second properties, as these are more likely to be fix-and-flip properties or temporary second homes due to moves. For private investors with more than two properties, the share drops to 3.3% after the regulation. First-time buyers purchase approximately 45% of properties in both treated and adjacent areas before the regulation, with the share increasing in 2022, particularly in regulated areas.

The next set of observations reports the position of households in the 2021 income distribution that moved to a property after it was sold. In regulated areas, this is slightly lower in both 2021 and 2022 compared to adjacent areas, in line with the difference in transaction prices. Residents of investor-bought properties have substantially lower incomes, tend to be younger, and are less likely to hold the Dutch nationality. Finally, properties on average contain about two adult residents.[6]

[6] Note that buy-to-let property here means that a property bought by an investor has one or more registered tenants, who by definition cannot be the owner. There are still some tenant households in regulated buy-to-let property sold after the law was introduced. It is likely that many of these properties are bought by parents for their children, which are exempt from the law: relative to regulated property sold in 2021, these tenants are substantially younger, and more likely to hold the Dutch nationality.

Table 1: Descriptive statistics of sold property in Rotterdam

	Regulated						*Adjacent*					
	Before			After			Before			After		
	Mean	σ	Obs.	Mean	σ	Obs.	Mean	σ	Obs.	Mean	σ	Obs.
Transaction data:												
Sale price, all	266,783	*74,552*	1,130	307,782	*89,411*	1,149	290,227	*71,800*	1,048	325,327	*99,138*	1,246
Sale price, investor	232,783	*68,696*	394	260,917	*103,205*	106	250,582	*69,645*	189	298,808	*148,201*	228
Tax value (2022), all	231,565	*63,633*	1,130	234,810	*66,752*	1,149	258,083	*57,295*	1,048	256,359	*58,490*	1,246
Tax value (2022), investor	201,267	*56,754*	394	208,749	*75,649*	106	225,657	*56,559*	189	238,443	*57,907*	228
Investor share	0.349		1,130	0.092		1,149	0.18		1,048	0.183		1,246
Private investors, share	0.189		1,130	0.033		1,149	0.098		1,048	0.083		1,246
First-time buyer share	0.426		1,130	0.624		1,149	0.439		1,048	0.474		1,246
Household-level data:												
Income percentile (2021), all	42.44	*27.28*	1,083	46.98	*25.32*	764	48.09	*27*	1,094	50.41	*25.83*	763
Income percentile (2021), buy-to-let	24	*25.11*	267	34.9	*29.94*	30	24.67	*26.34*	128	26.1	*22.67*	58
Resident-level data												
Age, all	31.43	*9.58*	2,435	32.28	*10.15*	1,507	34.68	*12.48*	2,309	34.13	*11.83*	1,624
Age, buy-to-let	29.43	*9.23*	847	27.24	*8.95*	89	30.25	*11.35*	371	30.71	*11.52*	184
Non-Dutch share, all	0.48		2,435	0.37		1,507	0.44		2,309	0.39		1,624
Non-Dutch share, buy-to-let	0.68		847	0.39		89	0.63		371	0.53		184
Property-level data												
Adult residents, all	2.02	*1.2*	888	1.81	*0.91*	647	1.85	*1.01*	933	1.81	*0.895*	741
Adult residents, buy-to-let	2.33	*1.44*	263	1.67	*0.9*	42	1.88	*1.1*	125	1.71	*0.97*	86

Notes: This table presents descriptive statistics for property transactions and matched residents in Rotterdam, which will be used in the empirical analysis. The data covers April 2021 to December 2022 and is divided into two periods: one year before and after an investor ban implemented on January 1, 2022. The 'Regulated' neighborhoods are the treated group, while 'Adjacent' neighborhoods serve as the control group (see Figure 1). Investor purchases refer to transactions by non-natural persons or individuals with multiple non-residential properties. We exclude non-investor buyers (e.g., non-profit housing corporations) and transactions involving multiple properties. Newly constructed and non-owner-occupied properties are also excluded as they are not subject to the policy.

3.2 Investor Ban: Rest of the Netherlands

Municipalities that have introduced the policy or plan to do so have gradually rolled out their proposals for implementation since November 2021. In 2022, 33 municipalities introduced the policy, all applying the law at the property level based on the annually determined assessed tax value. Other cities are still considering the policy but did not implement it in 2022 either due to a lack of bureaucratic capacity or because local politicians were not yet convinced of the necessity of introducing a policy swiftly.

As the policy focuses on cities with substantial buy-to-let activity, our sample excludes 10 smaller municipalities that introduced a policy but saw fewer than 30 investor purchases in the entirety of 2021. All cities with a population above 200,000 have introduced a policy during 2022. Most other cities that have introduced the policy are either mid-sized independent cities or part of the urban area of a larger city. This sample of 23 cities will be the main sample for our national analysis.

There are clear differences in the policy implementation. First, there is substantial variation in the price cap. While cities with higher sale prices tend to have higher caps and the mode is at 355,000 euros[7], sizable differences exist across cities. For example, The Hague and Tilburg have both opted for a price cap of 255,000 euros even though properties in The Hague on average sell for 20% more. Second, due to the bureaucratic burden to set up the policy, introduction and announcement dates vary. Third, while most cities have decided to cover the entire housing stock below the price cap, some have adopted a more targeted approach, implementing the law only in selected neighborhoods. For instance, Rotterdam's policy covers 27% of the housing stock, while Amsterdam's policy covers 60% of the housing stock.

[7] Many cities have opted for a price cap of 355,000 euros in assessed tax value. The Netherlands operates an attractive mortgage insurance scheme (Nationale Hypotheekgarantie) which in 2022 is available to buyers of properties with sales prices below 355,000 euros. Although purchase prices are typically higher than assessed tax values, many municipalities have used this limit as a price cap.

Table 2: Overview Policies *Opkoopbescherming*

City	*Population*	*Introduction*	*Limit*	*Avg. Price*	*Coverage*
>200,000 pop					
Amsterdam	882,633	01-Apr-22	512,000	569,890	Full
Almere	217,828	24-May-22	355,000	390,174	Full
Eindhoven	238,326	01-Apr-22	355,000	367,046	Full
Groningen	234,950	01-Mar-22	305,500	321,804	Near-Full
Rotterdam	655,468	01-Jan-22	355,000	364,018	Partial
The Hague	553,417	01-Mar-22	355,000	411,717	Full
Tilburg	224,459	01-Sep-22	355,000	342,612	Full
Utrecht	361,699	18-Mar-22	440,000	469,949	Full
Other					
Amersfoort	158,590	01-Apr-22	343,000	416,829	Full
Amstelveen	92,331	01-Jul-22	411,000	625,658	Full
Arnhem	163,888	01-Mar-22	325,000	344,738	Near-Full
Den Bosch	156,538	01-Jul-22	260,000	412,222	Full
Den Helder	56,334	28-Jan-22	250,000	225,902	Partial
Deventer	101,446	15-Oct-22	355,000	343,585	Partial
Dordrecht	119,537	14-Mar-22	355,000	299,241	Full
Gouda	74,095	01-Jul-22	355,000	327,961	Full
Haarlem	162,898	01-Feb-22	389,000	507,862	Full
Maastricht	121,151	01-Oct-22	355,000	353,620	Full
Nijmegen	179,100	07-Dec-22	350,000	364,849	Full
Oss	93,307	01-Nov-22	260,000	353,733	Full
Schiedam	79,644	01-Nov-22	355,000	279,455	Partial
Wageningen	39,939	15-Feb-22	355,000	403,316	Partial
Zeist	65,987	01-Jun-22	512,000	571,853	Near-Full

Notes: This table presents an overview of Dutch municipalities that implemented a policy in 2022 and that we follow in the sample. It includes information on introduction dates, limits, and coverage based on the proposals. Average prices reflect sales prices in 2021, and population data is as of January 1, 2022. In certain cities, the limit is not based on the 2022 tax value but on a previous year. Utrecht raised its limit to 487,000 on July 1, 2022.

4 Analysis

To estimate the impact of the policy, we compare treated properties to untreated properties both before and after the introduction of the policy. With y_i as the dependent variable at the transaction level, we estimate the following main specification:

$$y_i = \alpha_0 + \alpha_1 Post_i + \beta_1 Treated_i + \beta_2 Treated_i \times Post_i + \chi' z_i + \epsilon_i \tag{1}$$

In the baseline specification, $Post_i$ is a dummy variable that takes the value of one when the transaction date of property i is after the introduction of the local policy, and zero beforehand. For instance, in Rotterdam, $Post_i$ is one for every sale taking place from January 1, 2022, onward.

Next, $Treated_i$ is a dummy variable that assumes the value of one when a property is subject to the law. This implies that the property has a tax value below the local tax value cap and is located in an area affected by the policy. Tax value caps vary across cities.

In regressions using the Rotterdam sample, $Treated_i$ compares all properties below the cap in both treated ($Treated_i$ =1) and adjacent ($Treated_i$ =0) neighborhoods. The identifying assumption is that, conditional on controls, properties below the cap in treated and adjacent neighborhoods follow parallel trends prior to the regulation. We already observed this for investor activity, and we will verify it later for prices. In the national sample, our baseline sample will consist of transactions from cities that introduced a policy during 2022. Similar to the Rotterdam sample, we only include properties below the locally introduced price cap. This implies that the identification of the treatment primarily comes from cross-sectional variation in the introduction of the policy, as well as from geographic variation from the few cities that do not apply the ban to the entire city.

Our key coefficient β_2 measures the interaction between the two dummy variables. We focus on several outcome variables y_i: the probability that a property is purchased by investors, the probability that a property is purchased by first-time homebuyers, and the property's transaction price.

The set of control variables z_i includes most importantly the 2022 tax value of a property (valuated on January 2021), property type (apartment, row house, (semi)-detached, etc.), the number of square meters, building year, plus location and year-times-month fixed effects. In specifications that include time-fixed effects or neighborhood fixed effects, α_1 and β_1 will be absorbed by the fixed effects. We deviate from this specification using various other specifications and subsets of data.

The 2022 property tax value serves as a catch-all variable for a bundle of (potentially excluded) hedonic characteristics, as municipalities determine this tax value based on detailed hedonic property characteristics and realized sales prices up to three quarters after the valuation date. These valuations are generally highly accurate, with tax values explaining over 90% of the country-wide variation in house prices. The 2022 property tax value is the market value on January 1, 2021, and the valuation thus does not reflect any information about the investor ban.

4.1 Impact on the composition of buyers

We start by testing whether the law was effective in changing the composition of buyers: the policy can be expected to only have an effect on prices and the residents of these neighborhoods if it has meaningful economic effects on who buys properties available for sale. We initially estimate the model in Eq. (1) with an investor purchase dummy variable ($Investor_i$) as the dependent variable. Table 3 reports the main results for the impact on investor purchases in Rotterdam (Columns 1–2) and all major cities that introduced the law (Columns 3–4). In a final specification, we also include data from the entire Netherlands (Column 5). In the sample used to estimate the effects, we always remove data from transactions taking place one month before the introduction of a local policy to avoid anticipation effects contaminating our estimates (see Figure 2). For Rotterdam, we thus exclude December 2021. If we include this month, our treatment effects would increase. Standard errors are based on heteroskedasticity-adjusted errors.[8]

Before the start of the policy, 18% of non-treated properties in Rotterdam were purchased by investors and 31% of treated properties. Here, we focus only on properties sold by owner-occupiers in 2021 and exclude properties sold in December 2021. The

[8]We have also examined clustered standard errors, either by clustering at the buyer-level (as investors buy repeatedly) or at geographic level. Standard errors do not increase substantially if we cluster at neighborhood-level (Rotterdam) or municipality-level (national analysis), and the significance of our results does not depend on it.

Table 3: Impact on Buyer Composition

	Dependent variable:									
	Investor purchase					First-time buyer				
	Rotterdam		Cities with policy		All	Rotterdam		Cities with policy		All
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treated × Post	-0.226***	-0.230***	-0.099***	-0.095***	-0.109***	0.138***	0.125***	0.045***	0.063***	0.066***
	(0.024)	(0.023)	(0.009)	(0.006)	(0.004)	(0.031)	(0.031)	(0.013)	(0.010)	(0.007)
Post	0.006		0.012		0.001	0.047*		0.028*		0.010
	(0.017)		(0.009)		(0.003)	(0.022)		(0.013)		(0.005)
Treated	0.135***		0.036***		0.020***	0.012		0.093***		0.045***
	(0.020)		(0.006)		(0.004)	(0.023)		(0.009)		(0.006)
log(Tax Value, 2022)		-0.203***		-0.224***	-0.151***		0.069		-0.085***	-0.275***
		(0.050)		(0.015)	(0.004)		(0.056)		(0.019)	(0.005)
Square meters		0.002***		0.001***	0.001***		-0.003***		-0.002***	-0.001***
		(0.0004)		(0.0001)	(0.00002)		(0.0005)		(0.0002)	(0.00003)
Constant	0.177***		0.128***			0.427***		0.418***		
	(0.013)		(0.006)			(0.017)		(0.008)		
Property type FE	No	Yes	No	Yes	Yes	No	Yes	No	Yes	Yes
Building year FE	No	Yes	No	Yes	Yes	No	Yes	No	Yes	Yes
Neighborhood FE	No	Yes	No	Yes	Yes	No	Yes	No	Yes	Yes
Year–Month FE	No	Yes	No	Yes	Yes	No	Yes	No	Yes	Yes
Observations	4,202	4,202	46,689	46,689	287,162	4,202	4,202	46,689	46,689	287,162
R-squared	0.039	0.151	0.014	0.133	0.132	0.025	0.089	0.010	0.107	0.185

Notes: This table indicates shifts in property purchases by investors and owner-occupants around the 2022 investor ban. It uses transaction data from April 2021 to December 2022 on properties impacted by the ban, excluding data from the month prior to the local introduction of a policy. 'Post' marks the year after the ban and 'Treated' marks properties in designated or nationwide areas. Investor purchases refer to transactions by non-natural entities or owners of multiple properties who don't reside in them. Other variables are detailed in Appendix 9. Standard errors adjust for heteroskedasticity. Significance is marked by *, **, and *** for the 10%, 5%, and 1% levels respectively.

difference-in-differences estimate points to a decline in investor purchases in treated areas of 23 percentage points.[9] This suggests that the policy was effective in reducing investor purchases to a small fraction of its pre-policy level, in line with the plot in Figure 2. Note that we should not expect the effect to diminish entirely, as residences tied to businesses could still be purchased by rental investors, as well as properties for immediate family members or properties for non-rental investment (e.g. fix-and-flip activity).

To corroborate that our experimental setup allows for a reasonable claim of exogeneity, Column 2 includes an extensive set of property controls, such as tax values, square meters, and fixed effects for building years, property types, and neighborhoods. Importantly, our estimates hardly change after including observable controls and unobservables (in terms of fixed effects) that dramatically raise the models' explanatory power from 3.9% to 15.1%. This verifies that the results are not driven by unobservables (Altonji et al., 2005). In addition, Figure 2 demonstrates that before the investor ban, regulated and unregulated properties are on parallel trajectories until the final month before the ban becomes effective, and that the number of purchases falls steeply once the ban is implemented. These results clearly indicate a causal effect of the ban on investor purchases.

In Columns 3 and 4, we repeat the specifications of Columns 1 and 2 except that we now use the national sample. In the national sample, on average, 13 percent of properties are bought by investors pre-reform. For properties in regulated areas below the local price cap, this is 16 percent. Because treated properties in the national sample have less ex-ante investor activity, the treatment effect is also smaller and about 10 percentage points at the national level. Beyond this level-difference in investment activity, the overall pattern is very similar. The effect does not change substantially when including a wide set of controls, although the absolute move in the coefficient is slightly larger compared to Rotterdam, mostly because treatment and control properties differ more both geographically and in terms of property value in the national sample.

One could be concerned that our treatment effect on treated areas is upward biased

[9]Our point estimates would increase if we included data from December 2021 or from 2020, but activity in these periods was non-representatively high.

if investors that initially bought treated properties move to buying untreated after the introduction of a local policy, implying that the ban also increases transaction activity in control areas. Such 'waterbed' effects were indeed discussed heavily in Dutch politics. If investors do not have a strong preference for local real estate investment, such effects will arguably be limited as the bans are a negligible reduction in the total available investment opportunities for investors. However, real estate investments tend to be highly local with 50% of private investors' real estate located in the same municipality as where they live (Hochstenbach, 2022). Additionally, if buy-to-let investments accommodate local rising rental demand, investors might increasingly buy non-treated properties to accommodate this demand. In Column 5, we use sale data from the entire Netherlands to address whether this is an economically important concern. If spillovers are significant, we would expect investment activity on non-treated property in cities with a policy to rise after introduction. Because Column 5 also includes data from municipalities that did not have a policy, the $Post_i$ coefficient would measure such effects. However, it is insignificant and close to zero. This indicates that investment purchases of non-treated local properties did not increase significantly after a municipality introduced a ban. Note that we find a similar null result on the $Post_i$ dummy in Columns 1 and 3, but unlike Column 5 they do not control time fixed effects, implying it could also be driven by other time-varying factors.[10]

In Columns 6 to 10, we change the dependent variable to a dummy indicating whether a buyer of a property is a first-time buyer. Except for the change in dependent variable, the specifications are exactly the same as in Columns 6-10. Starting with the Rotterdam sample in Column 6, we find that pre-reform around 44% of properties below the cap were sold to first-time homebuyers. After the reform, treated properties observe a significant increase in the share of first-time buyers by about 14 percentage points. Again, the effect is similar when adding an extensive set of control variables (Column 7). At the national

[10] In unreported estimates, we also extended the Rotterdam sample to the entire city and instead included a dummy for non-treated neighborhoods adjacent to treated neighborhoods. We do not find higher investor activity in adjacent neighborhoods relative to neighborhoods farther away after introduction of a policy.

level, we again find a slightly smaller effect of about 4.4 percentage points (Column 8). This effect increases slightly to around 6 percentage points when including control variables (Column 9) and when adding all Dutch housing sales (Column 10). The share of first-time buyers pre-reform appears to be similar in the national and Rotterdam sample except for a higher rate among treated properties. Comparing the coefficients in the first five columns with those in the final five columns indicates that around half of the reduction in investor activity due to the policy resulted in an increase in the share of first-time buyers.

In sum, the number of investor purchases drops dramatically for regulated properties, confirming that the law removed most investors from the market. The remaining purchases either fall under one of the three exceptions, are not used for rental investment, or are leased illegally. Thus, the law was effective in reducing investor activity and increasing the chances of first-time homebuyers. Unsurprisingly, these shifts are largest in areas with high ex-ante activity, such as the regulated areas in Rotterdam. Including observable controls and unobservables dramatically raises the models' explanatory power but hardly changes our estimates, indicating that the results are not driven by unobservables.

4.2 Impact on house prices

To estimate the effect on house prices, we use log sale prices ($p_{i,t}$) as the dependent variable in Eq. (1) and three different difference-in-difference approaches (and test the parallel trends assumption for each of these in Figure 3). In addition to the control variables from Eq(1), we now control for tax value by each quarter since high- and lower-priced properties might be on different property price trends: Properties may have substantially different valuations in treatment and control neighborhoods, even below the price cap). We remove 0.5% of observations where the sale price deviates significantly from the tax value, which indicates (unobserved) renovations or updates to the property.[11] We also

[11] To define extreme outliers, we regress log sale price on the tax value and neighborhood-by-time fixed effects. We remove observations with residuals four standard deviations away.

again remove sales taking place within a month before the introduction of local legislation.

The results for Rotterdam are tabulated in Columns 1 and 2 of Table 4, comparing the evolution of properties below the cap in regulated and adjacent neighborhoods. We find a treatment effect of -0.5%, not significantly different from zero with a standard error of 0.9%. When we add more granular property controls and neighborhood fixed effects, we find the effect to be even closer to zero (0.1%) with a standard error of 0.8%, indicating that we can rule out that the policy had a large impact on the prices of treated properties in Rotterdam.

Table 4: Impact on house prices

	Dependent variable:					
	log(Sale Price)					
	Rotterdam (1-2)		National (3-6)			
Identifying variation:	Spatial		Close (10%)		Cross-section	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.019**		-0.018***	-0.064***	-0.0003	
	(0.009)		(0.004)	(0.004)	(0.002)	
Treated × Post	-0.005	0.001	0.001	-0.0003	0.004	0.002
	(0.009)	(0.008)	(0.005)	(0.004)	(0.003)	(0.003)
log(Tax Value)	Yes	Yes	Yes	Yes	Yes	Yes
log(Tax Value) × Quarter	No	Yes	No	No	No	Yes
Neighborhood FE	No	Yes	No	Yes	No	Yes
Property controls	No	Yes	No	Yes	No	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,076	4,076	11,408	11,408	49,802	49,802
R-squared	0.772	0.813	0.689	0.822	0.794	0.836

Notes: This table measures the house price effects of the investor ban. Using data from April 2021 to December 2022, excluding one month before the introduction of the policy, the analysis focuses on properties valued below the property tax cap in treated and adjacent areas in Rotterdam (Columns 1-2), and nationally on properties in treated areas within 0.1 log points of the local binding price cap (Columns 3-4) or on properties in mid-sized cities with and without regulation (Columns 5-6). 'Post' denotes the year following the ban, while 'Treated' represents properties subject to an investor ban. Further variables are outlined in Appendix 9. Standard errors adjust for heteroskedasticity. Significance is denoted by *, **, and *** for 10%, 5%, and 1% levels, respectively.

While the identification based on Rotterdam alone is most credible, we can get more precise and externally valid estimates by looking at all cities with a policy. At the national level, we have to use different identification strategies, though. While pre-reform trends in investor purchases follow parallel trends with treated property almost no matter how we define the control group, trends in house prices vary substantially across geographies and segments. As most cities regulated the entire housing stock below a threshold value, and prices in this segment evolved differently from prices in adjacent suburbs without a policy, we cannot use the local geographic identification we used in Rotterdam. We use two alternative identification strategies.

In Columns 3 and 4 of Table 4, we compare, in all cities with a policy, properties within treated areas with tax values within 0.1 log points of the local binding price cap. The identifying assumption is that properties just below and above the price cap are on parallel price trends conditional on controls. Analyzing close to the price cap thus provides the most credible identification. In Columns 3 and 4, we do not find any price effects, with estimates extremely close to zero and standard errors of just 0.5%. The results are also similar with and without the full set of controls. Note ex ante investor activity is somewhat less intense for properties close to the price cap, potentially resulting in a smaller treatment effect.[12]

Finally, in columns 5 and 6, we compare the price evolution in mid-sized cities that did or did not introduce a policy during 2022. In this sample, we remove the 8 largest cities since these have all introduced a policy and might be on different trends compared to mid-sized cities. From the remaining set of municipalities, we include all municipalities that are part of the so-called *G40* cities, a collaboration of all mid-sized cities in The Netherlands. For cities with a policy, we only include properties below the local price cap. In cities without a policy, we only include properties with a tax value below 389,000 euros, which matches the price cap of Haarlem, the mid-sized city with the highest price cap in the sample. Other than that, the treatment definition stays the same as before. Columns 5 and 6 show a point estimate of 0.4% with a standard error of 0.3% in the baseline specification, and an estimate of 0.2% with a similar standard error in the specification

[12] In 2021, investors bought around 9% of properties close to the cap relative to 14% among all properties below the cap.

using full controls. Hence, we find no effect on prices also when we compare cities with and without a policy.

Figure 3: Parallel Trends: Price Impact of Investors

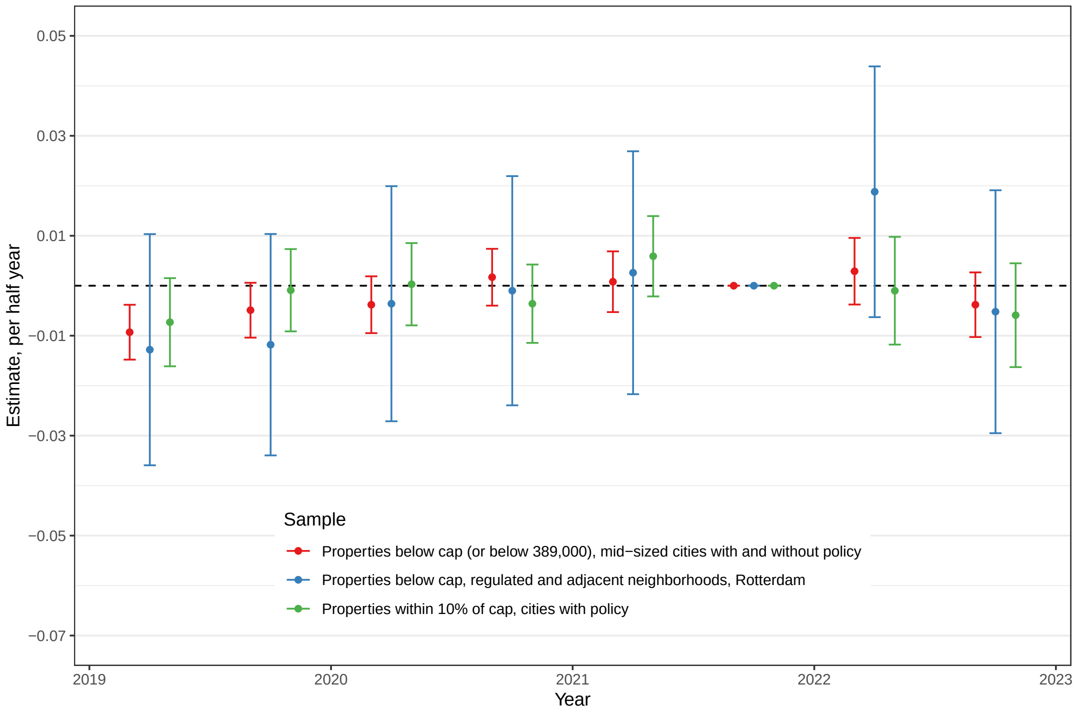

Notes: This Figure shows the impact of the investor ban on house prices for each of the three identification strategies in the text. The plot shows differences in house prices between treatment and control properties, corresponding to the main coefficient in Eq. (1) estimated separately at 6-month intervals. Vertical bars indicate 95% confidence intervals. Standard errors are adjusted for heteroskedasticity.

4.3 Potential channels

Our estimates for both Rotterdam and the wider panel of cities are conditional on the presence of parallel trends in the treatment and control groups. To verify whether our effects are not driven by alternative trends, we estimate the interaction terms per half-year since 2019 for the specifications including the more extensive set of controls (Columns 2, 4, and 6).

The results are presented in Figure 3, with the second half-year of 2021 as the baseline. In general, the economic magnitudes are small and precisely estimated, with standard errors around 1.2% for the Rotterdam analysis and 0.4-0.6% for the national analysis.

For all of the identification strategies, we observe almost completely parallel trends before the investor ban was implemented. So, conditional on our control variables, the absence of a causal effect of the policy on house prices does not seem to be offset by different price developments between treatment and control areas before the ban. In our plot of parallel trends in Figure 3, we also do not observe sizable effects when we look at the second part of 2022 alone, when prices have had more time to adjust. Overall, our analysis indicates that the removal of investors did not have economically sizable effects on house prices, at least until the end of the sample period.

The absence of a price effect suggests that, at least in the time period around the reform, the valuations of investors and regular homebuyers are sufficiently close to each other to substitute investors after the policy by owner-occupiers making very similar price bids. Investors might have been able to buy large amounts of properties if they were able to offer better secondary conditions, for example, by being able to buy a property quickly without resolutive conditions. Existing literature has found that growing private and institutional investment in the wake of the financial crisis supported the recovery of prices.[13] However, the effect of investors on market prices likely differs across different stages of the housing cycle. The investor ban was introduced in response to a short supply of residential properties and during a time when the housing market was booming, with properties generally sold above asking prices with multiple interested bidders. This implies there was a much larger pool of non-investor buyers in this period compared to crisis or recession periods.

Other channels could play a role too, and we explore some of these in the next sections. Most importantly, we look at the impact of investor presence on the residential composition. One motivation for introducing the policy was that buy-to-let investment was suspected to deteriorate neighborhoods. If that is the case, owner-occupiers might be willing to pay more for properties that insured them against such developments, off-

[13]Similarly, when investors increasingly started buying property in the early 2010s, both in The Netherlands and other developed countries, there was ample supply of housing for sale as many households faced underwater problems and it was hard to obtain credit for prospective home-buyers.

setting the reduction in demand coming from the removal of investors. In the final part of the paper, we will also study whether the lack of a price effect is due to the effect of investors on liquidity, for example if some investors bought properties that otherwise would not have been sold, potentially even at discounted prices.

4.4 Neighborhood effects of changing home-ownership

Legal documents underpinning the investor ban specifically stipulate that municipalities can introduce the policy to combat livability problems in neighborhoods. The implicit assumption is that buy-to-let activity might have negative externalities on neighborhoods, particularly because of high resident turnover compared to owner-occupancy. Homebuyers might perceive the policy as insurance against such developments, which could increase their willingness to pay. Sellers generally advertise whether a property is subject to the policy or not. Hence, the ban might have changed owner-occupiers' willingness to pay for property in affected neighborhoods relative to homeowners.

Many resident interest groups were also actively lobbying to increase the coverage of the policy to improve livability, explicitly stating that buy-to-let residents caused nuisances and limited social cohesion due to high levels of turnover.[14] In line with this, political parties in the Rotterdam City Council came up with several motions to expand the policy. One of these passed, and the city government committed to consider these proposals after the evaluation of the policy became public, which largely builds on the results of this study.[15]

While we cannot estimate this insurance value directly, we can test whether these concerns are legitimate by studying how investor ownership affected the composition of residents. We have documented above that after the introduction of the policy, properties initially bought by investors tend to be purchased by owner-occupiers. The policy

[14] Examples include the following letters to the City Council: Letter of February 21, 2022 and Letter of 23 December 2021.

[15] For example, see the minutes of the Rotterdam city council meeting of October 20, 2022 and of December 15, 2022

may also affect the composition of individuals residing in these neighborhoods if renting households have different characteristics from owner-occupying households, even if they reside in the same property. If such differences are large, policies affecting conversions from renting to owner-occupancy could have a sizable effect on tenant composition in the longer run.

4.4.1 Household income

To test the impact of the law on tenant composition, we conduct four different analyses. Our sample includes properties below the local tax limit in Rotterdam and in the cities mentioned in Table 2. We measure household income using the percentile in the 2021 household income distribution or log-transformed 2021 household income. As we will see shortly, the latter is sensitive to how we deal with the households that report income below what is reasonably required to live from. The results are in Table 5.

Table 5: Impact on Income Composition

	Dependent variable:							
	Household income, percentile						log(Household inc.)	
Sample area:	R'dam	Cities	R'dam	Cities	R'dam	Cities	R'dam	Cities
Transaction year:	2021	2021	2021/22	2021/22	2021/22	2021/22	2021/22	2021/22
Investor-owned	-19.36***	-15.09***			-19.51***	-16.58***	-0.219***	-0.228***
	(1.680)	(0.723)			(1.492)	(0.546)	(0.031)	(0.012)
Treated × Post			2.803*	1.894**	-0.407	0.697		
			(1.615)	(0.667)	(1.680)	(0.656)		
log(Tax Value)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Neighborhood FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year–Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,819	21,146	4,925	37,883	4,925	37,883	4,537	35,489
R-squared	0.199	0.232	0.125	0.165	0.169	0.190	0.125	0.167

Notes: This table displays income shifts in areas influenced by the 2022 investor ban, using household income percentiles and log household income for households earning at least €1,000/month. The sample covers properties sold below the local tax value limit. Except for Columns 1 and 2, the sample covers both sales in 2021 and 2022. 'Treated' denotes properties in affected areas for Rotterdam (odd-numbered Columns) and in cities adopting an investor ban by end of the sample period (even-numbered Columns). 'Post' marks the period after the ban, while 'Investor-owned' indicates properties owned by non-resident investors. Further details are available in Appendix 9. Standard errors are clustered at property-level. Statistical significance is marked at the 10%, 5%, and 1% levels.

First, we measure the difference in household income for residents of properties bought by owner-occupiers and investors in the period just before the law passed, holding the characteristics of the property fixed. We focus here on residents of properties bought in the entire 2021, as it seems unlikely that the income of a tenant would be affected by when an investor bought the property. Column 1 shows that for properties below the price cap in regulated and adjacent neighborhoods in Rotterdam, investor-owned properties have residents with incomes 19 percentiles lower in the income distribution. In Column 2, the national sample, this difference is 15 percentiles.

Next, we measure the impact of the reform on the income of residents of properties sold before and after the introduction of the reform. In the Rotterdam sample (Column 3), the average household income of residents of treated property was around 2.8 percentiles higher in the income distribution for properties sold after the policy introduction. This effect is only marginally significant, though, as the small sample results in large standard errors. At the national level (Column 4), we can establish this more precisely, and we find an effect of 1.9 percentiles. The national sample consists of all properties below the local price cap, and identification relies on the different introduction dates across areas.

Note that for properties bought in 2022, some residents might not have moved or registered on the property by March 31, 2023. If this correlates with income, that might lead to biased effects. In Columns 5 and 6, we test more directly how much of this effect is due to the removal of investors from the market by adding back the control for investor-ownership. After controlling for investor-ownership, the treatment effect on income disappears, and the effect of investor-owned property does not seem to differ based on the 2011 sample only. Thus, the increase in average household income of sold property appears entirely driven by the ban's substitution of low-income tenants in investor-bought property for middle-income owner-occupiers.

Finally, Columns 7 and 8 indicate that the average household in investor-owned property also has significantly lower log household income, of about -22% in both samples. For this analysis, we have excluded households with incomes below 1,000 euros per month;

the coefficient would be significantly larger if we include individuals with income closer to zero. However, these households likely earn unregistered income, either by not reporting it to the tax authorities or because they receive financial or in-kind support from their parents, which is generally not registered. The latter is very common for students.

4.4.2 Resident characteristics

In terms of other residential differences between investor-owned and owner-occupied properties, we look at the likelihood that a resident does not hold Dutch nationality, the age of adult residents, the number of adult residents in the property one year after purchase, and the likelihood that a tenant moves within two years of moving into the property. The latter two outcomes cannot yet be estimated for properties sold after the introduction of the law. Hence, for reasons of brevity, we focus on documenting ex ante differences in residential composition. For all analyses, we only include residents who moved into the property after it was sold.

Table 6 presents results and follows the specifications in Table 5, except for the change in dependent variable and sample. Columns 1 and 2 show that residents of investor-bought property are respectively 24 percentage points (Rotterdam) and 18 percentage points (national) more likely to be a citizen of a different country. Although most rental units in the Netherlands are owned by social housing associations, the private rental sector plays an important role in housing migrants, particularly because it is difficult to obtain social housing without spending years on waitlists, and buying property is not an option available to everyone.

Residents of investor-bought property are, on average, about 3.5 years in Rotterdam (Column 3) and 2.9 years in the national sample (Column 4). Young households with limited income have generally accumulated limited time on the waiting lists for social housing, and those who move to another city for a new job generally first rent in the private rental sector before buying property or moving into social housing.

Buy-to-let activity is also associated with more intense usage of housing. In Rot-

Table 6: Impact on Residential Composition

	Dependent variable:							
	Non-Dutch		Age		Adult Residents		% moved within 2y	
Sample area:	R'dam	Cities	R'dam	Cities	R'dam	Cities	R'dam	Cities
Transaction year:	2021	2021	2021	2021	2021	2021	2019/20	2019/20
Investor-owned	0.252***	0.179***	-3.553***	-2.976***	0.329***	0.250***	0.352***	0.297***
	(0.022)	(0.009)	(0.437)	(0.198)	(0.073)	(0.031)	(0.017)	(0.006)
log(Tax Value)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Neighborhood FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year–Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6,099	43,284	6,099	43,284	2,391	17,273	9,922	75,368
R-squared	0.135	0.236	0.132	0.164	0.094	0.141	0.222	0.188

Notes: This table compares personal characteristics of residents that moved to properties after they were sold, separating owner-occupants and renters in investor-owned properties. The dependent variables are: residents with non-Dutch nationality (dummy, Columns 1-2); resident age (Columns 3-4); number of adult residents per property (Columns 5-6); and residents moving out within two years (dummy, Columns 7-8). The sample is based on properties below the tax limit in Rotterdam (odd-numbered columns) and nationally (even-numbered columns). The "Investor-owned" indicator represents residents in properties owned by investors. Definitions for all variables are in Appendix 9. Standard errors are clustered at property-level. Symbols *, **, and *** denote 10%, 5%, and 1% significance levels, respectively.

terdam, the number of adult residents registered in buy-to-let property one year after purchase is 0.30 higher than non-buy-to-let property, an increase of around 15% relative to the average (Column 5). In the national sample, the effect is somewhat smaller with 0.25 additional residents.

Finally, buy-to-let property is characterized by very high rates of turnover. Based on properties sold in 2019 and 2020, we find that the likelihood of any given resident moving out within two years after moving in is 35 percentage points (30 percentage points) higher in Rotterdam (in the national sample). On average, around half of the residents in investor-bought property in Rotterdam move within two years. These high rates of turnover were one of the main reasons for introducing the law, as high turnover rates are perceived to reduce the livability of neighborhoods. High levels of mobility are likely an endogenous product of both market factors and institutional factors. Most importantly, the private rental sector generally caters to a class of residents who are more likely to move frequently. This effect might be exacerbated in the Dutch context, as fiscal policies are targeted to make it more attractive to live in owner-occupied housing or in social housing from housing associations with much lower rents.[16]

4.5 Other effects: impact on sales and rents

In the analysis so far, we have highlighted that by changing the composition of buyers, the buy-to-let ban also changed the composition of residents of these properties. These compositional changes might be an important reason why we find no discernible effect on house prices. However, the policy might also have affected housing markets and residents in other ways. First, investor demand might also have led to an increased number of sales of owner-occupied properties, for example, if investors were actively approaching individuals to sell their property or by purchasing property owner-occupiers deem unattractive. Second, the changes in buy-to-let conversions following the policy would also affect the rental market. This could lead to increased rental prices on incumbent rental property. Such higher prices could make treated neighborhoods even less accessible for low-income renter households. In this section, we aim to provide evidence for such effects. For reasons of data availability, this analysis is (currently) only based on more limited data from Rotterdam, and it should be treated as more suggestive. In the next iterations of the paper, extended data availability should allow for more precise identification.

4.5.1 Investor demand and sale rates

In Rotterdam, as well as in other Dutch cities and countries, investors often actively approach owner-occupiers to buy property that they deem attractive as buy-to-let investments, for example, with door-to-door flyering. Sellers might respond to these offers if it saves them time and brokerage fees or because they would like to sell property dis-

[16]Tenants in the private rental segment also increasingly rented on temporary rental contracts (up to 2 years), which were only legalized in 2016. However, these contracts do not seem to drive high turnover: turnover rates were already high prior to 2016. The Dutch parliament recently passed a law to ban these contracts again. The government is also proposing a new law to tighten rent price regulations.

creetly, while buyers use this strategy to obtain property at lower prices. The investor ban may have stopped investors from approaching property owners of properties subject to the regulation, with little to no effect on prices. Alternatively, investors might have been able to buy properties at discounted prices that otherwise would not have been sold in the regular residential market. We can test for these channels through the impact of the policy on the likelihood of a home being sold. This also allows us to give more context to the economic magnitude of investor purchases relative to the housing stock.

The dependent variable is a dummy variable that takes on the value of one if a property is sold in a given year, with the sample covering both 2021 and 2022. The sample consists of all properties in Rotterdam that are owner-occupied at the beginning of the year, have a tax value below the limit of 355,000, and are located in a treated neighborhood or an adjacent one. The results are in Table 7.

Table 7: Impact on likelihood of sale

	Dependent variable:		
	Sold	Sold (select.)	Sold (select.)
	(1)	(2)	(3)
Post	-0.005^{***}	0.003^{**}	0.003^{**}
	(0.001)	(0.002)	(0.002)
Treated	0.016^{***}	0.003^{*}	
	(0.002)	(0.002)	
Post × Treated	-0.007^{**}	-0.003	-0.003
	(0.003)	(0.002)	(0.002)
Square meters			-0.0001^{***}
			(0.00002)
Constant	0.057^{***}	0.041^{***}	
	(0.001)	(0.001)	
Neighborhood FE	No	No	Yes
Building Year FE	No	No	Yes
Property Type FE	No	No	Yes
Observations	115,162	115,162	115,162
R^2	0.0006	0.0003	0.0043
Note:	*p<0.1; **p<0.05; ***p<0.01; $^{\cdot}$p<0.1		

In the first column, we use sales in 2021 as a baseline, including transactions in

December 2021 and the first quarter of 2021. In 2021, around 7.3% of treated properties owned by owner-occupiers at the beginning of the year were sold. We find a treatment effect of -0.7%, indicating that transaction activity in treated areas went down by about 10 percent relative to the previous year and untreated property.

However, it is important to realize that this effect includes any anticipation effects of the policy. As we saw in Figure 2, there was a large increase in December 2021 of investor purchases of treated property. Investors likely tried to advance their purchases to before January 1, 2022. Hence some properties sold in 2021 would have been sold in 2022 without the policy. In Column 2, we therefore change the dependent variable to a dummy that takes the value of one if a property was sold between April 2021 and November 2021, or during 2022. This way, we exclude most of the anticipation effects. After imposing this restriction, the effect turns insignificant to around -0.3%.

In the months April 2021 to November 2021 alone, investors bought around 1.2% of the treated owner-occupied housing stock, so a point estimate of -0.3% would suggest that a quarter of these would not have been sold otherwise. Although this is an economically sizable effect, the large standard errors imply that in the end we cannot make claims on whether investor activity led to additional property being sold. Nonetheless, the fact that investors were able to buy over 1% of the treated owner-occupied housing stock between April and November 2021 alone (and 2% in the entire 2021) suggests that they had economically meaningful effects on neighborhoods over time, given that homeownership rates tend to be slow-moving.

4.5.2 Rental Market

We now explore the consequences of the shift in buyer composition for the rental market in affected neighborhoods. To test this, we use source listings from Spotzi, a private company that collects recent and historical data from the Netherlands' main real estate agents associations. These associations cover 70% of the Dutch housing market, with information on 4,967 rents in Rotterdam in the period 2019–2022.[17]

For these rental listings, we have information on the rental price, the number of rooms, the number of square meters, building year, energy label, monumental status, as well as location (neighborhood fixed effects) and time of listing (time fixed effects). These will all serve as dependent and control variables in our main regressions. As we (currently) cannot link this data to the restricted-access microdata from Statistics Netherlands, we only have an identifier for whether a property is subject to the policy or not (supplied by the city of Rotterdam), rather than the actual tax value. The lack of a tax value implies we can control for hedonic characteristics to a much more limited extent compared to the house prices.

There are a few notes to be made about the Rotterdam rental market and the coverage of the sample. First, the data covers private rentals only, and their rents are substantially higher compared to housing offered by housing associations.[18] Although these associations own most of the rental stock, long waitlists for association housing imply that movers from other Dutch regions or countries generally have to rely on the private rental market. Second, the listings include rentals from both large institutional investors and private investors or smaller commercial investors. The latter two groups are predominantly active in the buy-to-let market. We found no evidence that lease prices of recently bought buy-to-let property differ substantially from the overall private rental market.[19]

To assess the representativeness of the listings data, we compare list prices to contractual rents reported in the Dutch rent survey of Statistics Netherlands. Focusing on individuals who moved to their rental property in 2021 or later, the median 'base' rental

[17]We retain 2,677 rents after merging this data with the other data sources file from Rotterdam and filtering on properties below the price cap and trimming the bottom and top 2.5% (monthly rental price per square meter: > 11 and < 26 euros).

[18]Based on the Statistics Netherlands rental survey, we estimate that new tenants of housing association rentals in Rotterdam pay on average about 20% less rent compared to new tenants of observationally equivalent private rentals.

[19]To investigate whether rent-setting differs substantially between the buy-to-let investors in our sample and the general commercial rental market, we linked data from the annual Dutch rental survey of Statistics Netherlands to buy-to-let purchases in the sale records. Although the number of matched buy-to-let properties is small, as the survey covers few small rental investors, these properties do not rent for different prices compared to observationally equivalent private rental properties not acquired recently.

price in the 2022 survey for a private rental unit in Rotterdam with a tax value below the price limit is around 1080 euros. In the listings sample, the median asking rent for properties below the tax limit is around 1250 euros. The difference is likely mostly driven by costs for flooring, furnishing, or other service costs that are not incorporated into the 'base' rental price in the rent survey.

In the listings data, about 80% of the rentals in treated neighborhoods have tax values below the cap of 355,000, compared to 53% in non-treated neighborhoods. From 2021 to 2022, the total supply of rentals in properties below the price cap decreased by 20.6% in treated neighborhoods vs 16.5% in control neighborhoods. The difference is more stark when we look at the second half of 2022 only, as many properties bought by investors in late 2021 were likely only listed in the first half of 2022. Relative to the average half-yearly supply in 2021, the supply of treated properties declined by 25%, while the supply of rentals below the cap in other neighborhoods declined by only 5%. This reduction in the number of listings of treated properties is not surprising, given the impact of the policy on investor purchases we documented previously.

We use the standard difference-in-difference setup we used previously, but given the more limited sample both in terms of observations and property controls, we compare prices over the entire 2019-2022 period. The rental market was likely less directly affected by the changes in the stamp duty around January 1, 2021, implying there are fewer concerns about including earlier data. To stay consistent with earlier results, we will also provide half-yearly difference-in-difference plots, but these estimates will have wide confidence intervals. In Table 8, we regress (log-transformed) monthly rents against our usual set of explanatory variables, for all rental properties below the price cap.

Columns (1) and (2) indicate that rents have increased by around 4% for rentals in regulated neighborhoods, relative to those in non-regulated neighborhoods. For the full-sample average from Table 1, this amounts to just over €50/month, an economically meaningful amount. Note that after controlling for basic property characteristics, we do not find that rental prices in treated neighborhoods were much lower compared to

Table 8: Impact on rent prices

	Dependent variable:	
	log(Rent Price)	
	(1)	(2)
Post × Regulated	0.035**	0.043***
	(0.015)	(0.013)
Post	0.048***	
	(0.011)	
Regulated	-0.007	
	(0.008)	
Property controls	Yes	Yes
Neighborhood FE	No	Yes
Bi-annual FE	No	Yes
Observations	2,677	2,677
R^2	0.444	0.555
Adjusted R^2	0.440	0.548

Notes: In this table we regress (log-transformed) monthly rents against our usual set of treatment and time variables, and a limited set of property controls, for all rental properties below the price cap. Data are from Spotzi, who collect recent and historical data from the Netherlands' main real estate agents associations. We have information on 4967 2019–2022 rents in Rotterdam. We retain 4,312 rents after merging these data with the other data sources file from Rotterdam, and filtering on property tax availability. We trim the bottom and top 2.5% (monthly rental price per square meter: > 11 and < 26). 2,677 rents are below the price cap. Standard errors are adjusted for heteroskedasticity. *p<0.1; **p<0.05; ***p<0.01

untreated neighborhoods before the policy (Column 1).

Figure 4: Parallel Trends: Rent Price Impact of Investors

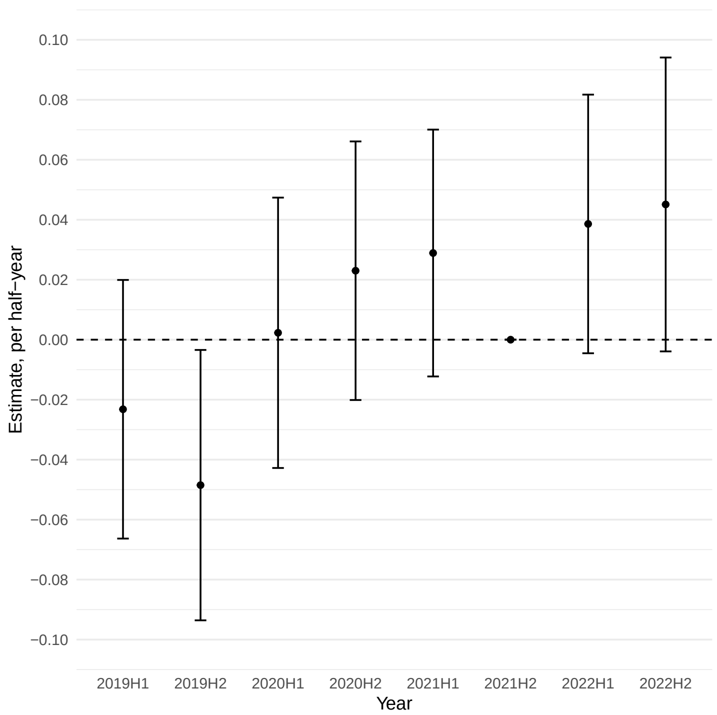

Notes: This Figure shows the impact of the investor ban on rent prices. The plot shows differences in rent prices between treated and non-treated neighborhoods, corresponding to the main coefficient in Eq. (1) estimated separately at 6-month intervals. Vertical bars indicate 95% confidence intervals. Standard errors are robust to heteroskedasticity. Given the low number of observations, confidence intervals are wide at 6-month intervals. The regression controls for property characteristics.

Figure 4 reports half-yearly estimates of the diff-in-diff coefficients as we have done previously. It should not be surprising that this results in very wide 95% confidence intervals, given the limited controls and lower number of observations compared to our main house price analysis. As a result, the coefficients on 2022 rental prices are individually not significant (although they are jointly significant). Their magnitude is comparable to the main effect in Table 8. While there does not seem to be a pronounced pre-trend in the evolution of rent prices prior to the introduction of the policy, the individual half-yearly coefficients move around starkly before the introduction of the policy, with lower prices in some years and higher prices in others. The lack of stability and corresponding wide confidence intervals might be related to the lack of precise controls, as well as more fundamental (temporal) divergence in rent prices across neighborhoods pre-policy. For this reason, the rent price effect should be treated more carefully.

The question is to what extent this positive effect might be driven by the fact that rent prices across neighborhoods did not follow parallel trends. Rather than estimating a single pre- and post-coefficient, Figure 4 reports half-yearly estimates of the diff-in-diff coefficients as we have done previously. It should not be surprising that this results in very wide 95% confidence intervals, given the limited controls and lower number of observations compared to our main house price analysis. As a result, the coefficients on 2022 rental prices are individually not significant (although they are jointly significant). Their magnitude is comparable to the main effect in Table 8.

Overall, there does not seem a systematic pre-trend in the evolution of rent prices prior to the introduction of the policy. However, the individual half-yearly coefficients move around starkly before the introduction of the policy, with lower prices in some years and higher prices in others. The lack of stability and corresponding wide confidence intervals might be related to the lack of precise controls, as well as more fundamental short-term divergence in rent prices across neighborhoods pre-policy. For this reason, the rent price effect should be treated somewhat more carefully than our results on house prices and residents.

5 Conclusion

We studied the impact of buy-to-let investment activity on the housing market and the composition of neighborhoods. While we find that the removal of investors from the housing market increases the share of first-time buyers, we find no effects on house prices and suggestive evidence for increases in rent prices. Residents in properties bought by investors have substantially lower incomes compared to residents of equivalent owner-occupied property. These differences explain the entire effect the policy had on the average income of residents in regulated properties. Residents of investor-owned property are also more likely to be young and foreign, move out of the property quickly. This shows that investor activity can have significant consequences on neighborhood composition, particularly over the long run, even when their direct price impact appears limited in the short run. Accordingly, the neighborhood effects of policies that affect local home-ownership rates might come predominantly from redistribution where lower-income tenants and higher-income home-owners live, rather than being a direct effect of home-ownership itself.

References

Allen, M. T., J. Rutherford, R. Rutherford, and A. Yavas (2018). Impact of investors in distressed housing markets. *Journal of Real Estate Finance and Economics 56*(4), 622–652.

Altonji, J. G., T. E. Elder, and C. R. Taber (2005). Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *Journal of political economy 113*(1), 151–184.

Ater, I., Y. Elster, and E. B. Hoffmann (2021). Real-estate investors, house prices and rents: Evidence from capital-gains tax changes. *Maurice Falk Institute for Economic Research in Israel. Discussion paper series.* (1), 1–33.

Austin, N. (2022). Keeping up with the blackstones: Institutional investors and gentrification. *Available at SSRN 4269561*.

Autor, D. H., C. J. Palmer, and P. A. Pathak (2014). Housing market spillovers: Evidence from the end of rent control in Cambridge, Massachusetts. *Journal of Political Economy 122*(3), 661–717.

Autor, D. H., C. J. Palmer, and P. A. Pathak (2019). Ending rent control reduced crime in cambridge. In *AEA Papers and Proceedings*, Volume 109, pp. 381–384. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Bayer, P., C. Geissler, K. Mangum, and J. W. Roberts (2020). Speculators and middlemen: The strategy and performance of investors in the housing market. *The Review of Financial Studies 33*(11), 5212–5247.

Bayer, P., K. Mangum, and J. W. Roberts (2021). Speculative fever: investor contagion in the housing bubble. *American Economic Review 111*(2), 609–51.

Bracke, P. (2019). How much do investors pay for houses? *Real Estate Economics forthcoming.*

Chinco, A. and C. Mayer (2015). Misinformed speculators and mispricing in the housing market. *Review of Financial Studies 29*(2), 486–522.

Cohen, J., R. Fuss, and C. Ghosh (2022). Do housing investors bargain differently across various racial and socioeconomic status neighborhoods? *Mimeo*.

Coulson, N. E. and H. Li (2013). Measuring the external benefits of homeownership. *Journal of Urban Economics 77*, 57–67.

Couture, V., C. Gaubert, J. Handbury, and E. Hurst (2019). Income growth and the distributional effects of urban spatial sorting. Technical report, National Bureau of Economic Research.

Couture, V. and J. Handbury (2023). Neighborhood change, gentrification, and the urbanization of college graduates. *The Journal of Economic Perspectives 37*(2), 29–52.

Diamond, R., T. McQuade, and F. Qian (2019). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco. *American Economic Review 109*(9), 3365–94.

Gao, Z., M. Sockin, and W. Xiong (2020). Economic consequences of housing speculation. *The Review of Financial Studies 33*(11), 5248–5287.

Gargano, A. and M. Giacoletti (2022). Individual investors' housing income and interest rates fluctuations. *Available at SSRN 4251621*.

Garriga, C., P. Gete, and A. Tsouderou (2021). Real estate investors and housing affordability. *Available at SSRN 3574001*.

Garriga, C., P. Gete, and A. Tsouderou (2023). The economic effects of real estate investors. *Real Estate Economics*.

Glaeser, E. L. (2003). Does rent control reduce segregation? *Swedish Economic Policy Review 10*, 179–202.

Gurun, U. G., J. Wu, S. C. Xiao, and S. W. Xiao (2023). Do wall street landlords undermine renters' welfare? *The Review of Financial Studies 36*(1), 70–121.

Hausman, N., T. Ramot-Nyska, and N. Zussman (2022). Homeownership, labor supply, and neighborhood quality. *American Economic Journal: Economic Policy 14*(2), 193–230.

Hochstenbach, C. (2022). Landlord elites on the dutch housing market: Private landlordism, class, and social inequality. *Economic Geography 98*(4), 327–354.

Hochstenbach, C. (2023). Balancing accumulation and affordability: How dutch housing politics moved from private-rental liberalization to regulation. *Housing, Theory and Society*, 1–27.

Hochstenbach, C. and S. Musterd (2018). Gentrification and the suburbanization of poverty: Changing urban geographies through boom and bust periods. *Urban geography 39*(1), 26–53.

Ihlanfeldt, K. and C. F. Yang (2021). Not in my neighborhood: The effects of single-family rentals on home values. *Journal of Housing Economics 54*, 101789.

Lambie-Hanson, L., W. Li, and M. Slonkosky (2022). Real estate investors and the us housing recovery. *Real Estate Economics*.

LaPoint, C. (2022). Property tax sales, private capital, and gentrification in the us. *Private Capital, and Gentrification in the US (September 15, 2022)*.

Mills, J., R. Molloy, and R. Zarutskie (2019). Large-scale buy-to-rent investors in the single-family housing market: The emergence of a new asset class. *Real Estate Economics*.

Sodini, P., S. Van Nieuwerburgh, R. Vestman, and U. von Lilienfeld-Toal (2016). Identifying the benefits from homeownership: A swedish experiment. Technical report, National Bureau of Economic Research.

6 Supplementary Figures and Tables

Figure 5: Policy Coverage and 2020 (top) / 2022 (bottom) investor purchases, Rotterdam

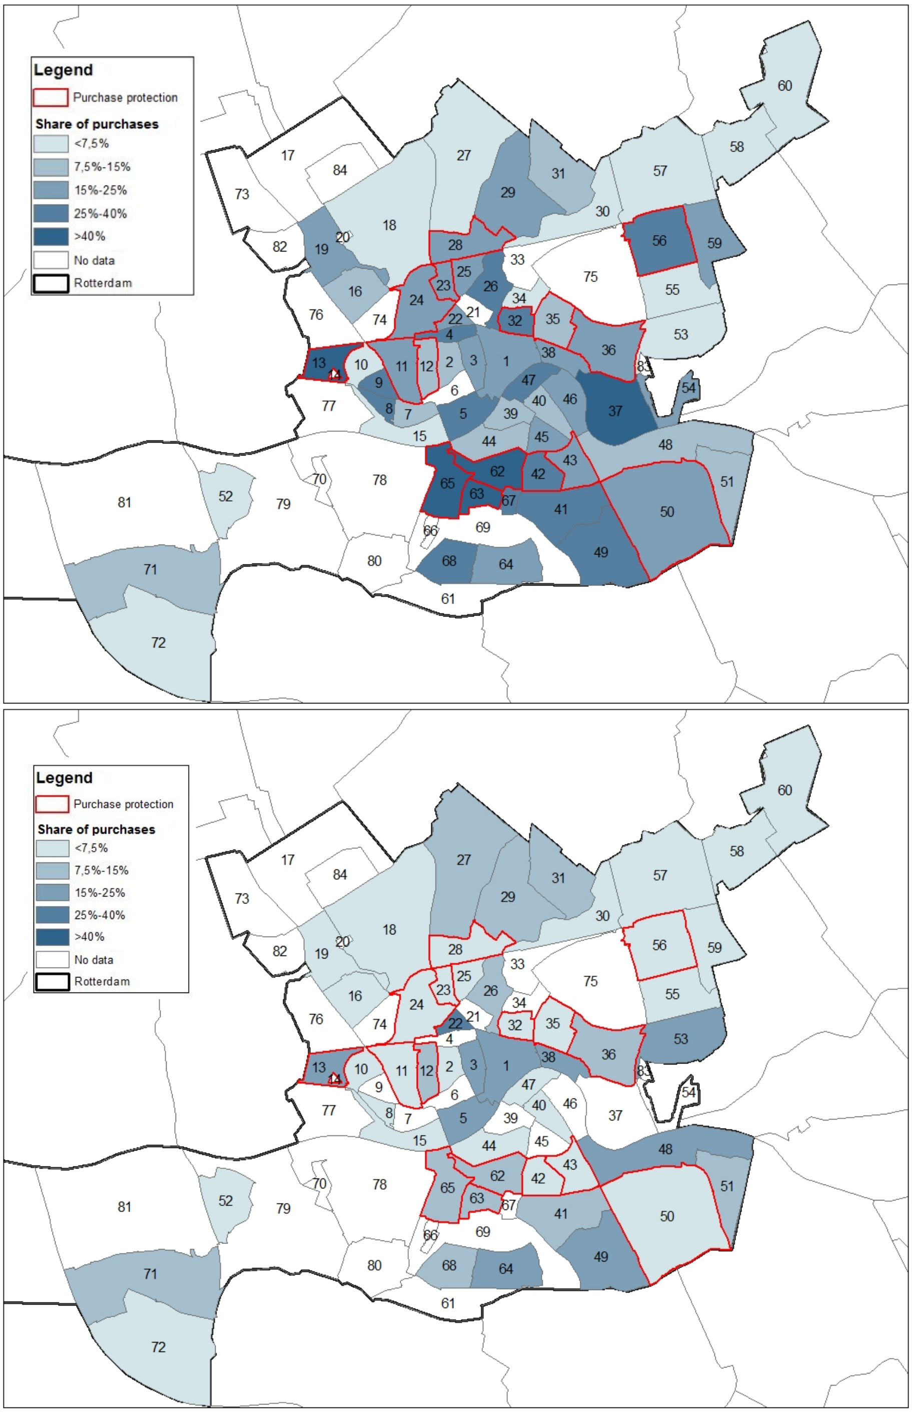

Notes: Figure 1 provides a plot of the Rotterdam urban area, indicating the regulated and non-regulated areas in the city and the share of investor purchases in 2020 (top) and 2022 (bottom) in each. This excludes purchases of second properties by private individuals.

Figure 6: Buy-to-let purchases in The Netherlands

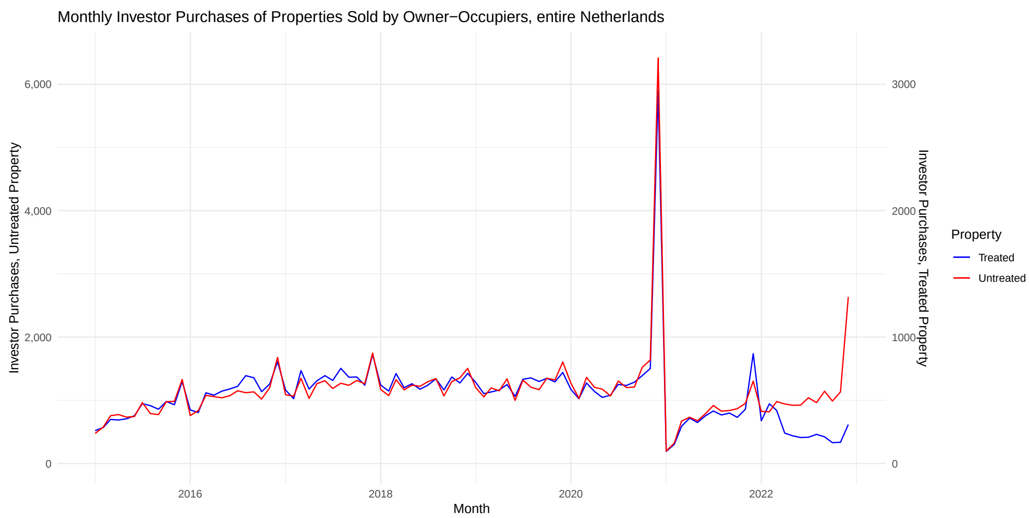

Notes: Figure 6 provides a plot of the monthly number of properties that is sold by owner-occupiers and bought by investors from January 2015 to December 2022. The plot separates this number for treated and untreated properties. Treated properties are properties that become subject to a buy-to-let ban during 2022; untreated properties are all other properties. The plot indicates that treated and untreated property follow perfect parallel trends in the years preceding the policy. There is some anticipation in December 2021. Afterwards, there is a gradual reduction in purchases of treated properties, as more and more municipalities start introducing buy-to-let bans. The spikes in December 2020 and December 2022 are related to anticipation of increases in the stamp duty for investment property (and a reduction for first-time buyers).

Table 9: Data Definitions

Variable	**Definition**
Tax Value	(Dutch: *WOZ-Waarde*). Estimated market value of the property on January 1 of the previous year. The market value is based on the assumption that the property is not rented and fully owned.
Square meters	The usable surface of a property in square meters.
Neighborhood	The neighborhoods we use are based on the *buurten* as defined by the Central Bureau of Statistics . In Rotterdam, the average neighborhood contains in 2022 around 8000 residents, with a maximum of 28,890 residents.
Property type	Classification whether a property is an apartment, a row house, corner house, semi-detached house or detached house.
Building year	Construction year of the property.
Investor	An owner of a property that is not a natural person or a natural person that owns at least one other property and is not living in the property under consideration.
Owner-occupier	A natural owner of a property that is (intending to) using the property as its main residency.
First-time homebuyer	A homebuyer that did not own any property in The Netherlands prior to buying the property under consideration.
Household income, percentile	Percentile of the disposable income of the household in the Dutch household income distribution. Household income is the gross income of the household minus paid income transfers such as taxes on income and wealth, and social or government insurance premiums (e.g. unemployment, aging, death). CBS name: INHP100HBEST
Household income, level	Disposable income of the household in euro's. Household income is the gross income of the household minus paid income transfers such as taxes on income and wealth, and social or government insurance premiums (e.g. unemployment, aging, death). CBS name: INHBESTINKH
Non-Dutch	Dummy variable for whether an individual resident holds a nationality that is different from the Dutch nationality.
Age	Age of the individual on December 31, 2022.
Adult Residents	Number of residents registered on a property one year after purchase.
% moved within 2 years	Dummy variable that takes the value of one if a resident moved out of a property within 2 years of moving in.