Buy-to-Live vs. Buy-to-Let: Homeownership and Residential Sorting in Housing Markets*

Marc Francke Lianne Hans Matthijs Korevaar Sjoerd Van Bekkum

July 11, 2025

Abstract

How does homeownership affect local housing markets and neighborhood composition? Exploiting the Netherlands' 2022 buy-to-let ban, which prevents investors from buying homes to rent out, we find the policy significantly reduces investor purchases and increases owner-occupancy. Despite removing investor demand, the ban does not negatively impact prices or transaction volumes as neighborhoods become more desirable to owner-occupiers. This desirability emerges through residential sorting: New homeowners have higher incomes, are more often Dutch-born, move shorter distances, and stay longer than the renters whom they replace. Thus, promoting homeownership may improve neighborhood quality but reduce mobility and housing opportunities for low-income renters.

*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 (korevaar@ese.eur.nl). Van Bekkum: Erasmus School of Economics, Erasmus University Rotterdam (corresponding author: vanbekkum@ese.eur.nl).

Notes: Results are based on calculations by the authors using non-public microdata from Statistics Netherlands. The authors thank the Municipality of Rotterdam for providing funding for access to these data and for valuable discussions on the details of its policy. We also thank the participants at the NYC Real Estate Conference, AREUEA-ASSA conference, and seminars at Aalto University, CPB, DNB, the European Commission, the Dutch Ministry of Finance, the Dutch Ministry of Internal Affairs, Tilburg University, University of Sydney, UTS, and Zillow for valuable comments. Korevaar has been supported by a Marie-Sklodowksa Curie Fellowship from the European Commission (grant 101.028.821).

1 Introduction

What happens if an investor, rather than an owner-occupier, buys a property? This question lies at the center of many housing policy debates around the globe. Since the 2007 mortgage crisis, rental investor activity has surged in most developed economies, raising concerns that buy-to-let investments may inflate local house prices, limit access to homeownership, and alter the social composition of neighborhoods.[1] Yet, despite a widespread, global articulation of this idea, there is little well-identified empirical evidence regarding the extent to which ownership structure—who buys and who lives in a home—affects housing markets and the sorting of residents across neighborhoods.

Identification challenges exist since the transitions from ownership to rentals and vice versa are rarely exogenous to property prices and neighborhood quality. Changing macroeconomic and financial conditions (such as credit cycles, tax changes, and interest rate changes) could cause adjustments in both ownership and the housing market and neighborhoods. The causal impact of buy-to-let on neighborhoods may be reversed when landlords target units where neighborhood changes are expected to raise rents and prices. Owner-occupied properties differ from rental properties in terms of size, location, quality, etc., and owners may differ from renters in preferences and characteristics too.

This paper overcomes these challenges by studying the 2022 *Opkoopbescherming* (purchase protection) which allows Dutch municipalities to prohibit investors from purchasing owner-occupied or vacant homes below a predetermined tax value. This investor ban causes plausibly exogenous shifts in demand for, and supply of, owner-occupied and rental housing in neighborhoods. This enables us to measure the effect of homeownership along a previously unidentifiable margin we call residential sorting. Neighborhoods gradually change because different people arrive, bringing their own preferences and resources. This

[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/)

margin is important for all housing policies causing non-negligible turnover in residents (e.g. investor caps, down-payment grants, first-buyer subsidies.).

The investor ban serves as a quasi-experiment that provides us with a discrete and unanticipated change in investor activity in "protected" neighborhoods. Because housing remains the same while the type of occupant changes, the ban allows us to isolate the impact of properties' renting versus buying tenure status. Our main focus is on Rotterdam, the second-largest city of the Netherlands that was the first to introduce the policy and the only major city to apply the ban to specific neighborhoods rather than the entire city. Property eligibility hinges on predetermined neighborhood-level investor activity and a hard tax-value cap of €355,000 in 2022. Hence, we can compare geographically close and observationally equivalent properties below the tax value cap in treated neighborhoods to those in similar neighborhoods.[2] A staggered DiD across 23 later-adopting cities confirms external validity by employing the Deb et al. (2024) estimator for repeated cross-sections that absorbs heterogeneity in treatment effects.

We analyze the ban's impact by linking individual housing transactions from the Dutch Land Registry to administrative data on residents and owners of these properties. The housing transactions data contain information on market prices, tax valuations of each property, and the property holdings of any natural person and entity in the Netherlands. The administrative data contain socioeconomic characteristics and incomes, at the individual level, for the occupants of these properties and for every other resident in the country. Hence, we have the universe of households who "buy-to-live", private investors who "buy-to-let", and the renters and homeowners who do not buy or move.

Despite largely eliminating buy-to-let demand for housing, we do not detect a statistically significant decline in property prices, which in fact shortly *rise* 2 to 3% for the

[2]Remaining unobserved heterogeneity is absorbed by detailed (and potentially time varying) property-level controls and property valuations that explain over 90% of price variation. Concerns about selection are attenuated by property control variables, neighborhood-time fixed effects, and neighborhood matching on pre-ban investor history: We match treated neighborhoods to control neighborhoods that just failed the official investor-activity criteria, producing treatment and control groups that differ only by the ban. A difference-in-differences (DiD) analysis (three quarters pre, four post) differences out time-invariant unit attributes, parallel pre-policy trends reduce reverse-causality concerns, and quarter dummies bracketing two stamp-duty hikes purge any remaining macro shocks.

next two quarters for regulated properties. At first sight, this seems surprising as the ban reduces the number of potential buyers for properties. However, buying owner-occupiers might perceive the removal of investor buyers as an amenity if it improves neighborhood quality. This is in fact a main motivation for the Dutch government to implement the ban.

We illustrate this channel with a simple conceptual framework in which supply from value-maximizing investors and demand from utility-maximizing homeowners are shaped by how agents incorporate expected improvements in neighborhood quality. Opposing forces are at play: While the ban shuts down demand from investors, putting downward pressure on house prices, removing buy-to-let purchases changes the occupants of the neighborhood. If owner-occupiers perceive that higher homeownership increases neighborhood desirability, the ban increases their willingness to pay for properties there, putting upward pressure on prices. On the supply side, other opposing forces exist as an investor will only sell a property and increase supply if the proceeds from selling exceed any expected increase in rental income. Thus, the ban's impact on local house prices is ambiguous.

Our empirical findings align with the conceptual framework. Relative to unaffected properties, the share of investor purchases on regulated properties decreases by 73% or 19 percentage points, this is offset by an increase of 13 percentage points in first-time home purchases and 6 percentage points in 2nd to n-th time home purchases. The ban increases the stock of owner-occupied housing in treated areas by 1% in 2022, a meaningful effect as ownership rates tend to be slow-moving. On the supply side, we observe similar offsetting in net aggregate supply from owner-occupiers, whereas supply from investors to investors decreases. We find no evidence of spillover effects: In adjacent, unregulated neighborhoods, there has been no discernible increase in investor buying activity since the introduction of the ban. Collectively, these results point towards the framework's underlying assumption that neighborhoods become more desirable if landlords cannot buy.

To validate this assumption, we first demonstrate the ban significantly changes the characteristics of occupants because "to-live" buyers are wealthier, have higher incomes, lower turnover rates, and are incumbents on the Dutch housing market. Next, we verify these effects arise solely because the occupant changes from a renter to a homeowner; once buy-to-let vs buy-to-rent is controlled for, no residual treatment effect remains. Thus, banning investors changes neighborhoods through residential sorting: As characteristics of residents change when the marginal property for sale is bought by an owner-occupier, so does the neighborhood.

These results have important social implications. Tenure-shifting interventions can simultaneously raise housing prices and contribute to gentrification, even in the absence of physical upgrading. However, the ban also restricts housing opportunities for lower-income renters. Affected properties house fewer individuals and are less likely to alleviate housing stress following major life events—whether positive (e.g., relocation for study or employment) or negative (e.g., job loss or divorce). Furthermore, the ban reduces residential turnover, as private investors typically offer flexible rental arrangements that facilitate mobility in response to such events. In this way, the policy unintentionally diminishes household mobility.

This article contributes to several strands of literature. First, we contribute to the literature examining the causal effects of homeownership on economic and social outcomes. Most closely related, Sodini et al. (2023) and Hausman et al. (2022) exploit variation in right-to-buy schemes, where incumbent tenants get the opportunity to buy their home at a highly discounted price. These studies find that homeownership affects wealth accumulation, consumption, and ability to climb the housing ladder (Sodini et al., 2023), and leads to higher house prices through increased employment and income, suggesting rising housing quality (Hausman et al., 2022). By keeping both occupants and property equal, these papers pin down the effect of ownership when resident composition is frozen.

In many other settings, changes in home-ownership also lead to changes in residents. Our paper highlights that home-ownership affects local housing markets and neighbor-

hoods because owner-occupancy changes who moves in, attracting residents with different socioeconomic characteristics and thereby changing neighborhood desirability. In other words, while these studies isolate the within-resident behavioral channel, we isolate the between-resident composition channel.

Second, our paper informs the rapidly growing literature on how activity from rental investors influences housing markets. The literature on this subject largely emerged after the financial crisis, when real estate investor activity increased significantly and thereby supported the growth in housing prices (Allen et al., 2018; Lambie-Hanson et al., 2022; Garriga et al., 2023). Much of this literature focuses on the effects of institutional buy-to-let (US: buy-to-rent) investors (e.g. Mills et al., 2019; Gurun et al., 2023; Austin, 2022; Coven, 2023; Barbieri and Dobbels, 2025; Chang, 2025). Institutional investors provided liquidity in housing by buying distressed properties during the housing crisis and often made significant improvements to these properties (e.g. Goodman and Golding, 2021; Austin, 2022). Relatedly, LaPoint (2022) shows that purchases of foreclosed properties by private equity investors can contribute to gentrification.

Most of these studies indicate a positive effect of entry of large single-family rental investors on house prices, whereas we find no – or even a negative – association between reduction in investor activity and house prices. Besides differences in identification, we should note that the context of institutional buy-to-let investors emerging in crisis times might differ from the typical buy-to-let investment. Different than "value investors" rescuing abandoned homes by offering liquidity or investing in renovations, the typical buy-to-let investors in our context more straightforwardly shift supply and demand by competing directly with first-time buyers for the same properties, and typically do not add new physical supply or quality upgrades. Additionally, the difference in price effects might be related to the fact that institutional investors cater to different and often richer residents than typical buy-to-let investors (which we observe in our Dutch data). The key contribution of our paper is to show that the price effect of investors might be tightly linked to how they change residential composition: investors house residents who

are significantly different from owner-occupants of similar properties. As this might differ across investor types, so might their price effects.

The fact that the policy shapes residential composition also relates to broader work on the effects of housing policies. Prior work on rent controls (Diamond et al., 2019; Autor et al., 2014), construction of affordable- or market-rate housing (Diamond and McQuade, 2019; Pennington, 2021) and rent voucher programs (Collinson and Ganong, 2018) in the US suggest that some of the effects are due to these interventions affecting where tenants and owners live. However, lack of experimental variation and individual-level data makes identifying such effects difficult. In this paper, the policy and data allow to precisely quantify how owner-occupiers differ from renters while keeping housing the same.

Finally, our paper links to literature that aims to study how home-ownership affects local neighborhoods (e.g. Coulson and Li, 2013; Ihlanfeldt and Yang, 2021; Hausman et al., 2022). One argumentation is that homeowners contribute more positively to the neighborhood than renters (DiPasquale and Glaeser, 1999). More broadly, property investments can generate positive externalities (Rossi-Hansberg et al., 2010), and homeowners might be more willing to make them. Our evidence shows that the marginal home-owner looks very different compared to the marginal renter, which might explain some of these dynamics.

This paper is structured as follows. Section 2 elaborates on the ban and the institutional background of housing in the Netherlands, and develops hypotheses based on the ban's assumption that homeownership improves neighborhoods, and the resulting impact on demand, supply, and prices. Section 3 describes the data and explains the empirical methodology. Section 4 examines the effectiveness of the ban by documenting its effect on demand in terms of investor activity and homeownership rates, supply in terms of changes in the housing stock, and property prices. Section 5 studies the assumption that homeownership improves neighborhoods by analyzing the effect of the ban on neighborhood composition. Section 6 extends our analysis from the city of Rotterdam to the national level to increase the external validity of our findings. Section 7 concludes.

2 The Buy-to-Let Ban: Context and Expected Impact

2.1 Homeownership and Rental Investment in the Netherlands

In The Netherlands, growth in buy-to-let and private rental investment is a fairly recent phenomenon, as most rental housing is provided by non-profit rental housing associations. Following the 2007 mortgage and 2011 euro crises, house prices dropped by about 20% in 2012 and the share of private rentals increased from 10% in 2012 to 13% in 2023, particularly in big cities (Amsterdam, from 24% to 30%; Rotterdam, from 17% to 22%). Roughly one-tenth of today's private rentals were converted from owner-occupied dwellings, which are classic buy-to-let conversions.

Various factors contributed to this rise: (i) Social housing construction was curtailed significantly after the crisis, making it more difficult to enter social housing, (ii) while generous to international standards, credit constraints were tightened in the 2010s through lower required loan-to-value ratios and scaled-down deductibility of mortgage interest expenses, (iii) new-build rates were low in the 2010s despite significant population growth, and (iv) regulations for private rental investment have been relaxed. Since investors tend to be less credit-constrained than (especially first-time) home buyers, competition in the housing market became more fierce in the 2010s, and younger households became less likely to enter homeownership (Van Bekkum et al., 2022; Rouwendal et al., 2023).

Columns 1 to 3 in the top panel of Table 1 compares residents and properties in the three main segments of the Dutch housing market as of January 1, 2022. Owner-occupied homes are worth €484k on average and house the richest households (€107k gross). Social rentals from mostly housing associations serve mostly lower incomes (€43k), but involve waiting lists of several years on average, to more than 17 years in desirable areas. Private rental incomes are in between with tenants earning €59k, still above social-housing incomes. Rental turnover is triple that of owner housing (25% vs 7%). Column 4 provides statistics for the buy-to-let subset of private rental properties. They are cheaper (€311k) and accommodate younger, lower-income, migrant-heavy households.

Table 1: Overview of Property & Resident Characteristics, Dutch Housing Market

	Housing Segments			
	Owner-occupier	**Social rental**	**Private rental**	***Buy-to-let***
Property tax value (mean)	€483,908	€250,271	€323,001	€310,835
Share of properties	58.1%	29.0%	13.0%	
Mean gross household income	€107,222	€43,113	€59,137	€55,086
Mean personal income (adult)	€45,012	€25,183	€35,548	€32,794
Share moving within 1 year	7.6%	10.5%	25.5%	31.1%
Share Netherlands-born	90.4%	72.6%	70.4%	63.4%
Average age (adults)	51.2	52.6	41.9	36.8
Share in major cities (G44)	37.3%	52.9%	63.9%	60.8%
	Private Rental Market by Landlord Type			
	Large	**Small/Medium**	**Retail**	**Second home**
Property tax value (mean)	€343,018	€320,706	€295,329	€361,752
Share of rental properties	15.1%	12.9%	26.9%	21.9%
Share buy-to-let (since 2009)	0.1%	11.3%	23.5%	22.5%
Size of landlord (median tenant)	5575	44	7	2
Mean gross household income	€74,168	€57,646	€51,699	€60,770
Mean personal income (adult)	€41,582	€34,851	€32,655	€35,921
Share moving within 1 year	17.3%	26.1%	29.9%	28.4%
Share Netherlands-born	77.1%	71.5%	64.0%	68.1%
Average age (adults)	48.5	40.9	36.5	40.6
Share in major cities (G44)	72.2%	62.6%	69.2%	55.8%

Notes: This table presets summary statistics of occupied dwellings and their residents on 1 Jan 2022, using Dutch Land Registry and administrative data. The top panel compares the three main tenure types with units converted from owner-occupation to rental since 2009 ("buy-to-let"), the homes covered by the 2022 investor ban. The bottom panel zooms in on the private-rental segment, distinguishing (i) large corporate landlords ($\geq$ 500 units), (ii) small/medium corporates, (iii) retail owners with $>$ 2 properties, and (iv) "second-home" landlords with $\leq$ 2 properties. Income figures are gross calendar-year 2022.

In the second panel of Table 1, we look at the distribution of properties in the private rental market, excluding non-profit and unknown owners. We distinguish commercial investors who are large (500+ properties) (Column 1) vs. small/medium-sized (Column 2), both of which own property as a business entity, and retail investors owning more than 2 properties (Column 3) vs. 2 properties or less (Column 4). Retail investors tend to invest in low-value units and about a quarter of their property ownership was converted from owner-occupied to rental since 2009. They rent more frequently to relatively low income tenants, young adults or immigrants. Second homes are most expensive on average as many of these are former owner-occupied property.

The properties of commercial investors, especially large investors, target higher incomes. Large investors are not active in buying owner-occupied or vacant homes to rent them out, and small and medium investors only to a lesser extent. In the US, some institutional investors have started buy-to-let activities, which has received considerable attention, but this development is not visible in the Netherlands.

In public opinion, the growing importance of private rental investors was seen as undesirable. Private rental housing is more expensive than social rental housing from housing associations, leading to complaints about affordability. In addition, concerns were raised that investors in rental housing were outbidding first-time home-buyers and hindering their transition to homeownership, as well as negatively affecting neighborhoods. This led to a series of policy measures that intended to restructure the Dutch rental market and improve the position of both renters and first-time buyers:

- On January 1, 2021, a year before the ban, the transfer tax for investor purchases was increased from 2% to 8% while it was reduced from 2% to 0% for owner-occupants under the age of 35. From April 1, 2021, the 0% rate was restricted to property purchases below 400,000 euros.
- On January 1, 2022, legislation went into force that provide municipalities the tools to ban buy-to-let investors, the intervention studied in this paper.
- On January 1, 2023, a second increase in transfer tax was implemented from 8% to 10.4%, aimed exclusively at investors. In July 2024, tenant protection laws were tightened by expanding rent controls and a law went into force that limited the room for temporary rental contracts.

Appendix B, Figure A.1, illustrates these events by plotting national trends in the number of investor purchases for treated (by the buy-to-let ban) and untreated properties. Immediately preceding the two transfer tax changes and the introduction of the investor ban, we note three December spikes in investor purchases. To ensure a valid estimation window based on these events, we use the three quarters before the ban's introduction

and the year afterwards to measure the effect of the ban, i.e., in between the two increases in transfer tax.

2.2 The buy-to-let ban

The so-called purchase protection ("*Opkoopbescherming*") act passed the parliament in mid-2021, was first operationalized by Rotterdam in November 2021, and became effective on January 1, 2022. Formally, the law provides municipalities with the possibility to implement regulations that forbid renting out any purchased property that was not leased on the day of the property transfer (or leased for less than six months). The restriction lasts four years and applies only to purchases made after the local start date of the ban.[3] Because investors rarely buy homes they must keep empty for four years, the rule operates as a de-facto ban on new buy-to-let conversions of the existing stock and is known as such. Two reasons were provided for putting this policy in place:[4]

- Investor purchases might increase the scarcity of affordable owner-occupied housing and corresponding house prices.
- The government argued that investor purchases can reduce neighborhood quality because owner-occupiers are assumed to be more connected to the neighborhood with much lower rates of turnover compared to private renters.

The law stipulates that a municipality can introduce a buy-to-let ban if it can justify that it is "necessary and suitable for combating the scarcity of cheap and mid-priced owner-occupied housing[,] or for the quality of the local environment." Because the statute does not offer an unambiguous definition of what is "necessary and appropriate," nor one of what is meant by neighborhood quality, municipalities therefore choose which areas and what price ceiling to cover, creating sharp cross-sectional and within-city variation. The national law only requires municipalities to provide three exemptions: letting to close

[3]The law affects only existing properties. For new construction, municipalities already had the legal means to prevent investors from buying properties built for owner-occupancy.

[4]See Memorie van Toelichting. ("Explanatory Memorandum")

family, a temporary owner absence, or a business use. These exemptions are relatively rare and do not blur this treatment definition.

This paper focuses on the city of Rotterdam, the first city to implement the ban on January 1, 2022, and the only major city that implemented the ban in specific neighborhoods rather than throughout the city. Treated neighborhoods are scattered throughout the city, but most are in a ring around the city center, see Appendix B, Figure A.1. In Rotterdam, the average neighborhood contains in 2022 around 8,000 residents, with a maximum of 28,890 residents. Rotterdam selected 16 out of its 71 neighborhoods for treatment, representing about 30% of its housing stock). Within these neighborhoods, Rotterdam applies the ban to all properties below a tax value[5] of €355,000, so that over 90% of properties fall under the price cap.

2.3 Expected impact of the ban

As the 2022 buy-to-let ban restricts investor purchases of homes below a statutory tax value cap, one would expect property prices to drop as demand decreases. However, the ban also aims to improve neighborhood quality by reducing renter turnover and enhancing social cohesion. In the legal motivation for the ban, buy-to-let activity was perceived to affect neighborhood quality ('livability') negatively.[6]

Expected improvements in neighborhood quality could lead to countervailing shifts in supply and demand, so that it is unclear how the ban affects the housing market. Specifically, if the expected quality of the neighborhood indeed increases, negative price pressure from eliminated investor demand may be offset by increased willingness to pay from owner-occupiers if renters and owner-occupiers have different preferences and socio-economic characteristics. Moreover, expected changes in neighborhood quality may also

[5]Based on the yearly assessed property tax value (in Dutch, WOZ-value), the estimated market value of the property on January 1 of the previous year, assuming the property is not rented and fully owned.

[6]Prior academic work finds a positive relation between neighborhood quality and homeownership. The main reasons are that owner-occupiers (1) have incentives to invest in their homes and neighborhood, because it increases their property values, and (2) are less likely to move due to higher transaction costs, reinforcing incentives to make long-run investment, see Hausman et al. (2022) for an overview of the literature.

affect investors' decision to sell the property now or charge higher rents in the future.

We illustrate these channels using a simple conceptual framework. To fix ideas, as neighborhood quality is multifaceted and lacks a precise definition, we simply assume in line with the policymakers that for a fixed neighborhood, neighborhood quality θ_t is an increasing function of owner-occupied properties:

$$\theta_t = f(O_t), f'(O_t) > 0 \tag{1}$$

where O_t is the number of owner-occupied properties at time t among all of the neighborhood's properties N with regularity conditions $f(0) = \theta_{min}$ and $f(N) = \theta_{max}$. We also assume that θ_t changes exogenously due to the ban.

The ban targets "buy-to-let" investors I who maximize expected discounted profits, to benefit "buy-to-live" owner-occupiers O who maximize utility derived from housing consumption h and neighborhood quality θ.

2.3.1 Homeowners' Utility

Homeowners maximize utility U subject to their budget constraint:

$$U(h_t, C_t, \theta_t) = \alpha\theta_t \ln(h_t) + \ln(C_t) \text{ s.t. } Y = P_t h_t + C_t, \tag{2}$$

where h_t captures housing consumption in a neighborhood of quality θ_t and C_t captures other consumption with preference scaled to one.

Since θ_t is exogenous, we can find the optimal choice of h_t by setting up the Lagrangian; taking the first-order derivatives with respect to h_t, C_t, and the budget constraint; and solving the first-order conditions for h_t:

$$h_t = \frac{\alpha\theta_t Y}{P_t(1 + \alpha\theta_t)} \tag{3}$$

As the ban improves neighborhood quality θ_t, it increases homeowners' housing utility and, therefore, homeowners' demand $D_{O,t}$. The ban also eliminates demand from

investors so after the ban, we assume $D_{I,t} = 0$. Thus, with N identical homeowners, we can define aggregate demand $D_t = Nh_t$. In Appendix C.1, we linearize D_t and formulate a linear, first-differenced demand function

$$\Delta D = e\Delta P + f\Delta\theta, \tag{4}$$

with the sensitivity of aggregate demand to price, e, always negative (see Appendix C.2) and the sensitivity to neighborhood quality f always positive (see Appendix C.3). Hence, aggregate owner demand is downward-sloping in price but upward-sloping in neighborhood quality. If neighborhood quality θ_t is indeed a positive function of owner-occupied properties, the impact of the ban on demand may be ambiguous. By contrast, if we assume $\theta_t < 0$, the ban pushes demand downward through both channels.

2.3.2 Investors' Profit Maximization

Investors choose between keeping the property and earn rental income net of operational costs (e.g., maintenance) R versus selling the property now at price P net of transaction costs ϕ,

$$V_t = \max_{d_t \in \{0,1\}} \begin{cases} (1 - d_t)\left[R + \delta\mathbb{E}_t[V_{t+1}]\right], & \text{(Hold)} \\ d_t\left[P_t - \phi\right], & \text{(Sell)} \end{cases} \tag{5}$$

so an investor sells a property if its transaction price exceeds her reservation price $\tilde{P}, P_t \geq \tilde{P} = R + \delta\mathbb{E}_t[V_{t+1}] + \phi$. Expected profits are discounted to their present value using a discount rate of δ.

If the ban increases neighborhood quality θ, it has a positive effect on property value, and this is factored into the investor's expectation of future profit $\mathbb{E}_t[V_{t+1}]$:

$$R = R_0, \quad \mathbb{E}_t[V_{t+1}] = V_0 + \beta_V\theta_t, \tag{6}$$

where β_V represents how much the expected future value changes in response to a one-unit change in neighborhood quality. We can re-write the reservation price as

$$\tilde{P} = R_0 + \delta V_0 + \phi + \delta \beta_V \theta_t, \tag{7}$$

where the last term captures the present value impact of neighborhood quality on investors' reservation prices. The term is positive indicating that investors require a higher price for selling properties in better-quality neighborhoods.

Assuming N_I identical investors, aggregate supply is expressed as the fraction of investors willing to sell,

$$S_t = N_I \mathcal{G}(P_t - \delta \beta_V \theta_t), \tag{8}$$

where $\mathcal{G}_{\tilde{P}}(\cdot)$ is the cumulative distribution function of reservation prices. In Appendix C.4, we linearize S_t around the pre-ban equilibrium (P_0, θ_0) and formulate a linear, first-differenced supply function

$$\Delta S = b\Delta P + c\Delta\theta, \tag{9}$$

with the price elasticity of supply b the probability density function $g_{\tilde{P}}$ and always positive (see Appendix C.5), and the sign of coefficient c is always negative (see Appendix C.6). So while a higher price encourages selling, higher neighborhood quality discourages it by raising future expected gains. Thus, if owner-occupied properties positively affects neighborhood quality, the impact of the ban on supply is also theoretically ambiguous. By contrast, if we assume $\theta_t < 0$, the ban pushes supply upward through both channels.

2.3.3 Pricing

Market clearing occurs when total supply $S_t = S_{O,t} + S_{I,t}$ equals total demand $D_t = D_{O,t} + D_{I,t}$. After small changes in P_t and θ_t, the new demand and supply are $D_t = D_0 + \Delta D$ and $S_t = S_0 + \Delta S$, respectively. The market clearing condition becomes:

$$D_t = S_t \quad \Rightarrow \quad D_0 + \Delta D = S_0 + \Delta S \quad \Rightarrow \quad \Delta D = \Delta S, \tag{10}$$

since $D_0 = S_0$. We then write the linearized expressions for ΔD and ΔS in terms of ΔP:

$$-e\Delta P + f\Delta\theta = b\Delta P + c\Delta\theta \quad \Rightarrow \quad \Delta P = \frac{f-c}{e+b}\Delta\theta. \tag{11}$$

The last equation demonstrates that, similar to its effect on supply and demand, the net effect effect of the investor ban on property prices is also ambiguous if owner-occupied properties positively affect neighborhood quality. The numerator $(f - c)$ captures the owners' greater willingness to pay, net of the investors' reluctance to sell, due to changes in neighborhood quality.[7] So, if we assume that $\Delta\theta > 0$ and the ban and the subsequent rise in homeownership improve neighborhood quality, then prices may rise or fall depending on whether demand responds more strongly (via f) than supply contracts (via c).

3 Data and Methodology

3.1 Data

To assess the impact of the policy, we combine data from the Dutch Land Registry (DLR; *Kadaster*), which records all housing transactions and property ownership in the Netherlands, with administrative data of Statistics Netherlands on income, person-level characteristics including residence, and the property tax values of all individual properties in the Netherlands. Neighborhoods are also defined by the Central Bureau of Statistics.

We focus predominantly on outcomes and transactions in 2021 and 2022 in the main analyses, but sometimes use data from 2019 to 2024. The Land Registry data only cover sales of existing properties, thus excluding newly constructed properties, and these are unaffected by the law change. In addition to the transaction price and date, the DLR housing transactions data also contain hedonic property characteristics including the year of construction, the number of square meters of the property, the type of property, energy label, and neighborhood. We use these characteristics as control variables.

Our data also includes property tax appraisals from municipalities based on detailed

[7] The denominator $(e + b)$ captures the price elasticity of the housing market and is always positive.

hedonic property characteristics and up to three quarters of realized sales prices. These valuations are generally highly accurate, with tax values being able to explain over 90% of the country-wide variation in house prices. We use tax appraisal values to supplement the set of control variables above.

In our analysis of housing transactions, we focus on transactions of those properties that could be eligible for the ban: properties which are not rented out or rented out for less than six months at time of the sale. Since there is no administrative data on rental contracts, we match data on all resident spells to property transactions. If a resident started living in a property before the sale date and moved out after the sale date we consider a property tenant-occupied on the sale date and ineligible for the ban. Because there can be minor differences in resident registrations and transfer dates for owner-occupiers, we still consider properties eligible if a property was occupied because the seller had not yet deregistered by the transfer date or if the buyer already registered on the property before the sale and did within six months of the transaction. We also exclude sales by non-profit housing associations or other non-profit buyers and sales of more than three units, since these are in practice in most cases exempt. Note that being eligible for the ban does not mean a transaction is subject to it: This depends on whether a ban is actually in place for that property on the moment of sale.

3.2 Selection into treatment

Our key empirical challenge is that municipalities self-select whether and how they implement the ban: Municipalities observing more investor activity are more likely to implement a ban, which will correlate with other factors and policy changes. This is also the case for Rotterdam.

Figure 1 plots the raw average fraction of investor purchases for ban-affected and unaffected properties. Investor activity is time-varying: we observe a large spike and drop in investor purchases around the changes in the stamp duty in January 2021, with investor activity remaining at lower levels throughout 2021 and afterward.

After the introduction of the ban, we observe a big drop in the rate of investor purchases on treated property, but not on untreated property. Investor purchases on treated property do not drop to zero as some may not be affected, due to exemptions or not being rented out (and thus unaffected by the policy). There is also some evidence for a small spike in investor purchases just prior to the introduction. In 2023, the purchase of untreated property by investors also declines, as the effect becomes confounded with other regulations that also reduce investor activity (Francke et al., 2024).

Figure 1: Selection-into-treatment: investor purchases shares in Rotterdam

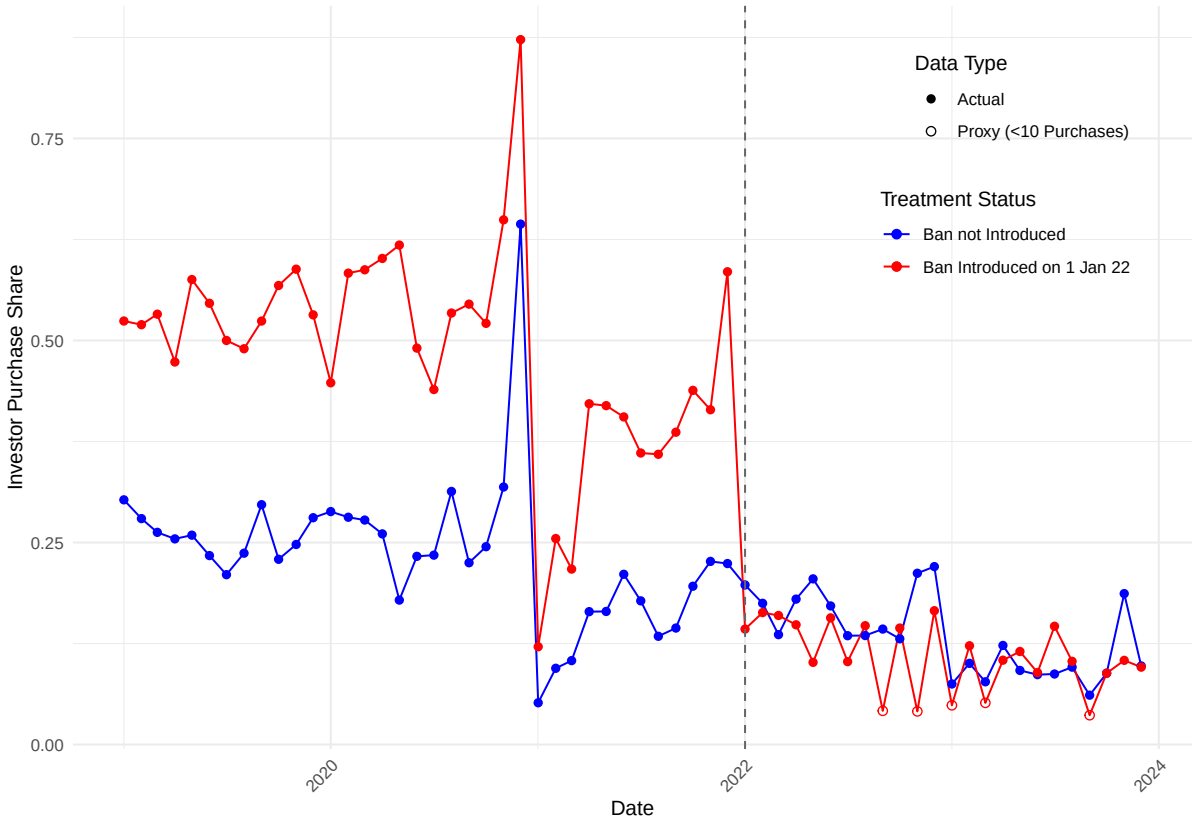

Notes: This Figure plots investor purchases over time in Rotterdam. The red line indicates properties in neighborhoods that became subject to the ban and have appraisal values below EUR 355,000. The blue line indicates properties that have appraisal values above EUR 355,000 or lie in neighborhoods that did not become subject to the ban.

There are two implications. First, considering Figure 1 and the series of policy measures discussed previously, we focus on comparing transactions in 2022 with transactions in the final three quarters of 2021, when investor activity for non-treated properties is fairly stable. Second, we know that Rotterdam selects neighborhoods for treatment based

on investor activity, so that the general decline in investor purchases in the Figure may have a larger effect on neighborhoods with high investor activity (as opposed to being driven by the policy itself).

To decide where to apply the ban, Rotterdam targeted high investor activity neighborhoods, but could only use approximate criteria to identify these. A neighborhood was designated as treated if it met at least two of the following three criteria:[8]

- The number of private rental properties should be above-average (1.000 properties);
- The fraction of private rental properties should be above-average (24%);
- The growth of the private rental stock should be above-average (+20%, 2015-2021).

Given that the city of Rotterdam applied these criteria, it is possible to match treated neighborhoods to untreated neighborhoods with similar expected investor activity. For example, due to selection on the absolute number of rental properties, smaller neighborhoods are less likely to be regulated. Thus, we can break the possible correlation between treatment and outcome variables by matching treatment neighborhoods to control neighborhoods that have similar pre-ban investor activity but were not selected for the ban.

To do so, we first estimate the likelihood that a property with a tax value below the price threshold is bought by an investor, conditional on observable property characteristics and neighborhood fixed effects. We estimate this model on sales in the three years preceding the ban introduction, the period from 2019 to 2021, and explicitly exclude time fixed effects. We also exclude small neighborhoods with 10 or fewer sales in 2022 and 2023. The estimation results of this model are provided in Appendix B, Table A.2. Investor purchases are largely determined by hedonic property characteristics.

We use the model coefficients to predict for each neighborhood the average rate of investor purchases, after the introduction of the ban, based on the observable characteristics

[8]The government also incorporated neighborhood quality into the selection decision, but this was not a binding constraint as all neighborhoods eligible based on neighborhood quality were also eligible based on the three criteria presented here. For the discussion on neighborhood selection that involves the municipal council, see https://gemeenteraad.rotterdam.nl/Document/View/2fa7572c-3863-4026-a464-2a35fda5e010.

of the properties sold in 2022 and 2023. This gives us a prediction of buy-to-let activity in the main sample period if the ban had not been implemented, assuming conditions in 2019-2021. These predictions can be found in Figure A.1, top panel.

We then rank neighborhoods by predicted investor activity and remove: (1) all treated neighborhoods whose predicted investor activity is higher than the untreated neighborhood with the highest predicted investor activity, and (2) all untreated neighborhoods whose predicted investor activity is lower than the treated neighborhood with the lowest predicted investor activity. This leaves us with a set of 12 treated neighborhoods and 29 control neighborhoods in Rotterdam. Figure A.1, bottom panel, plots these neighborhoods, as well as the neighborhoods dropped from the analysis.

Table 2 shows the result of the procedure. Column 1 demonstrates the existence of treatment selection, as the investor ban indicator strongly predicts investor ownership. In Column 2, this coefficient is no longer significant and close to zero. The adjusted R^2 drops from 27% in Column 1 to below zero in Column 2. Thus, after matching, the relation between treatment assignment and investor activity is effectively random.

Table 2: Relation between treatment and investor activity

	Dep. var.: Predicted Investor Share	
	All	Matched
Constant	0.2619***	0.4019***
	(0.0222)	(0.0200)
Investor ban	0.1853***	-0.0107
	(0.0269)	(0.0275)
Observations	127	41
Adjusted R^2	0.2690	-0.0216

Notes: This table presents the results of a weighted OLS regression on the neighborhood-level predicted level of investor activity for Rotterdam. Weights are based on the number of transactions in each neighborhood. Column 1 shows results for unmatched neighborhoods; Column 2 excludes high-investor activity treated neighborhoods and low-investor activity untreated neighborhoods. ***p<0.01.

3.3 Methodology

With y_i as the dependent variable at the transaction level, we estimate a standard difference-in-differences (DiD) equation for Rotterdam:

$$y_i = \alpha_0 + \alpha_1 Post_i + \beta_1 Treated_i + \beta_2 Treated_i \times Post_i + \chi' z_i + \epsilon_i \tag{12}$$

In the baseline specification, $Post_i$ is a dummy variable equal to one if the transaction date of property i is after the introduction of the local policy. $Treated_i$ is a dummy equal to one when a property is subject to the law, so a tax value below the local tax value cap and located in an area selected for the ban.

As mentioned in Section 2.1, our estimates are based on the time window in between the two increases in transfer tax, i.e., the three quarters before the ban's introduction on January 1, 2022 and the year afterwards. There is little room for anticipation effects as the announcement which neighborhoods would be treated came just a few weeks before the actual introduction of the law.

Outcome variables y_i include an indicator equal to one if a property is purchased by investors, an indicator equal to one if a property is purchased by first-time home buyers, an indicator equal to one if a property was sold in a given quarter (based on whether it changed ownership on the first day of the next), the property's transaction price, and various characteristics of the residents.

In most analyses, the set of control variables z_i includes the 2022 property tax value, property type (apartment, row house, (semi)-detached, etc.), the number of square meters, label of energy efficiency, construction 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. As we do not have the full set of control variables available for the entire housing stock, our supply regressions control for tax value, quarter fixed effects and neighborhood fixed effects alone. Similar to our analysis on buyers, we distinguish between property sold by owner-occupiers and by investors.

Since property tax appraisal valuations explain over 90% of the country-wide variation in house prices, we use them as a catch-all variable for any hedonic characteristics not available in the data.

4 Impact of the ban on property market

4.1 Demand

We start by asking to what extent the law was relevant and effective, in the sense it reduced purchases by investors. For every eligible transaction in the sample, we classify buyers as owner-occupier if (s)he lives in the property at any point in time within 2 years of the purchase date (and, if applicable, before the property is sold).[9] We further classify owner-occupiers as first-time homeowner if this is his/her first purchase on record (in administrative data or the Land Registry). Any buyer that is not an owner-occupier is classified as an investor. For both types of dwellings, we determine who lives there and obtain their socio-economic characteristics. We focus on residents who registered at a property on or after the transaction date.

In Table 3, we estimate Eq. (12) with as dependent variable the indicator equal to one if the transaction is an investor purchase (Columns 1 and 2) or an first-time owner-buyer (Columns 3 and 4). The sample includes all transactions in matched neighborhoods of treatment and control properties potentially eligible for the ban, i.e., below the tax value threshold and sold by owner-occupiers or sold by investors without a long-term tenant. Standard errors are clustered at the level of treatment, so at the neighborhood level.

The difference-in-differences estimate in Column 1 indicates that the ban significantly reduces investor purchases in treated areas by 19 percentage points, in line with Figure 1. The Treated coefficient indicates that before the ban, treated and control dwellings were statistically indistinguishable in their likelihood of selling to an investor. Column 2

[9]This leads to some censoring for 2023 transactions as these data are available until December 31, 2024. Therefore, our main analysis focuses on 2021 and 2022 transactions.

Table 3: Impact of the Ban on Investor and First-time-Buyer Demand

	Investor buyer (0/1)		*1st-time owner-buyer*	
Treated × Post	-0.192***	-0.166***	0.127***	0.093***
	(0.034)	(0.028)	(0.030)	(0.031)
Treated	0.014		0.018	
	(0.0368)		(0.034)	
Post	-0.034		0.026	
	(0.0236)		(0.020)	
Log(tax appraisal value)	No	Yes	No	Yes
Property controls	No	Yes	No	Yes
Neighborhood FE	No	Yes	No	Yes
Year-Month-of-sale FE	No	Yes	No	Yes
Observations	4,182	4,182	4,182	4,182
R^2	0.043	0.184	0.015	0.094

Notes: This Table presents DiD estimates of the demand effect of the 2022 investor ban for all property transactions in matched neighborhoods, using transactions data from the Dutch Land Registry from 2021Q2-2022Q4. In Columns 1 and 2, the dependent variable is a dummy variable that equals 1 if the buyer is an investor. In Columns 3 and 4, the dependent variable is a dummy variable that equals 1 if the buyer is a first-time owner-buyer. Columns 1 and 3 show baseline regressions without controls; Columns 2 and 4 add appraisal value and property controls, and neighborhood and Year-Month-of-Sale fixed effects. Property controls include size in square meters, property type, building year, and energy label. All variables are detailed in Appendix Table A.1. Standard errors are clustered at the neighborhood level. Significance levels: * p <0.10, ** p <0.05, *** p <0.01.

indicates that the estimate only slightly changes to 17 percentage points after controlling for property tax values, and square meters, and after adding fixed effects for construction periods, property types, energy labels, and neighborhoods.

In Columns 3 and 4, we observe that the drop in investor demand is partially offset by an increase in demand from fist-time owner buyers of 13 percentage points. The coefficient decreases to 9 percentage points when adding control variables and fixed effects. Thus, the ban appears to effectively help first-time home buyers by removing investors from the market. Since transactions can only be done by investors or owner-occupiers, the latter of whom may be a first-time buyer, the difference of -6 to -7 percentage points can be ascribed to repeat-buyers who already own property at the time of purchase.

In Figure 2a, we plot difference-in-differences coefficients from 2019 to 2023 with the

Figure 2: Parallel Trends in Investor and First-time-Buyer Demand

(a) Effect on Investors

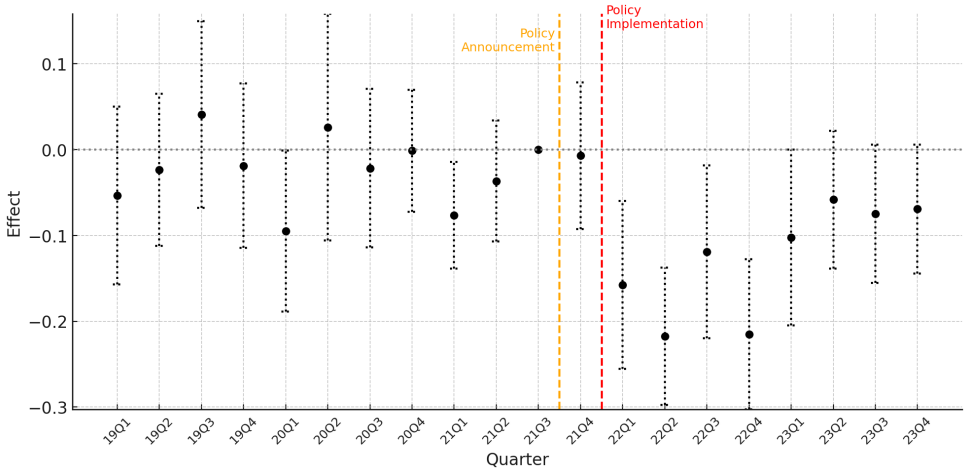

(b) Effect on First-Time Buyers

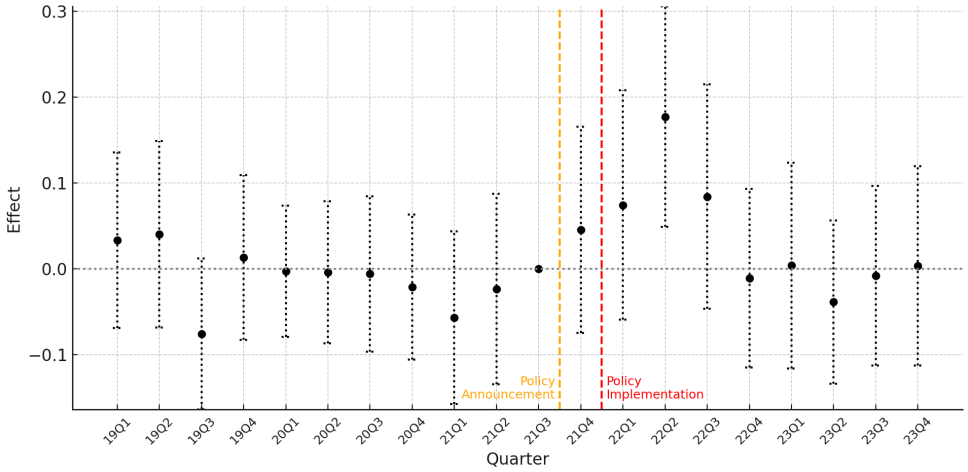

Notes: These Figures show the impact over time of the investor ban on the likelihood of investor purchases (Panel (a)) or first-time buyer purchases (Panel (b)) for all property transactions in matched neighborhoods, using Dutch Land Registry and administrative data. We plot the difference between treatment and control properties by estimating the treatment coefficients of Table 3, columns 2 and 4, separately for each quarter. Vertical bars indicate 95% confidence intervals. Standard errors are clustered at the neighborhood level.

policy announcement period as baseline. Despite large fluctuations in investor purchases, the coefficients are indistinguishable from zero in the quarters prior to the ban. This indicates parallel secular trends. We observe no statistically significant anticipation at the end of 2021. The actual neighborhoods where the ban was applied were disclosed only by mid-November: Any anticipatory purchases in treatment and control neighborhoods were made with equal probability.

Figure 2b shows that the reduction in investor purchases is significantly offset by the increase in purchases from first-time home buyers. Fluctuations in owner-occupier purchases fluctuate less, and are indistinguishable from zero in the quarters prior to the ban indicating parallel trends with limited evidence of anticipation. In line with Table 3, Figure 2b is the mirror image of Figure 2a. The remaining changes are caused by the share of first-time buyers relatively repeat-time buyers changing over time, for example due to rising interest rates.

4.2 Supply

The investor ban focuses exclusively on the demand side of the housing market, raising the question of how sellers respond. We study this in terms of the effect of policy on the probability of selling a property. Since any property can be sold at any time, this requires analyzing the entire housing stock from the Dutch Land Registry and Statistics Netherlands.

To this end, we start on January 1, 2021, and restrict the sample to properties with a property value below the tax limit of €355,000. We then determine who owns these properties and split the sample into residents' properties versus investors' properties, excluding property owned by housing associations and unknown owners. We then measure the probability of sales by quarter. For the stock of properties we do not have information on the full set of control variables; we can only control for the tax appraisal value and neighborhood and time fixed effects.

We report on the ban's housing supply effects in Table 4. Column 1 tabulates the effect

of the policy on the probability of selling by owner-occupiers. We find no effect of the policy on the overall sale rate, but observe significant changes in purchase rates of homeowners and investors: In Column 2, we zoom in on the fraction of these properties sold to other owner-occupiers. Here, we find an increase of about 20 basis points per quarter. Compared to the population-average sales rate of 1%, this is an economically significant increase. In Column 3, the fraction of these properties sold to investors decreases by an almost identical amount of 23 basis points per quarter.

The 3 basis points difference between coefficients is indistinguishable from zero: since first-time buyers cannot supply property by definition, Columns 2 and 3 combined constitute net aggregate supply and the coefficients perfectly offset. This reiterates our result on the demand side from Table 3: After the ban, properties sold by owner-occupiers are less likely to be bought by investors, but the lower probability is offset by an increased likelihood of sales to owner-occupiers.

Table 4: Impact of the Ban on Investor and Owner-Buyer Supply

	Dependent variable:					
Sample:	*Owner-occupied homes*			*Investor-owned homes*		
Buyer	All	Owner-Occ.	Investor	All	Owner-Occ.	Investor
Treated × Post	-0.0011	0.0020**	-0.0023***	-0.0015	0.0006	-0.0021**
	(0.0012)	(0.0008)	(0.0007)	(0.0010)	(0.0005)	(0.0009)
Log(tax appraisal value)	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	No	No	No	No	No	No
Neighborhood FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month-of-sale FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	294,368	294,368	294,368	317,408	317,408	317,408
R^2	0.002	0.001	0.002	0.002	0.002	0.002

Notes: This Table presents DiD estimates of the supply effect of the 2022 investor ban for all property sales in matched neighborhoods, using the stock of owner-occupied housing from administrative data from 2021Q2-2022Q4. In Columns 1-3, the dependent variable is a dummy variable that equals 1 if the seller is an owner-occupier. In Columns 4-6, the dependent variable is a dummy variable that equals 1 if the seller is an investor. All columns control for property characteristics, and neighborhood and year-month-of-sale fixed effects. Property controls include square meters, property type, building year, and energy label. All variables are detailed in Appendix Table A.1. Standard errors are clustered at the neighborhood level. Significance levels: * p <0.10, ** p <0.05, *** p <0.01.

In Columns 4–6, we present the results for the probability of selling by investors.

The overall probability decreases by 15 bps, which is borderline insignificant at the 10% level. Column 5 shows that the probability does not change for investor-owned properties bought by owner-occupiers, whereas it decreases significantly in Column 6 for investor-owned properties bought by investors.

Thus, net aggregate change in supply from owner-occupiers does not change whereas net aggregate change in supply from investors decreases (in)significantly depending on the type of buyer. To better understand these results, we present all six parallel trend plots (one for each column) in Figure 3. The plots do not indicate any evidence of pre-trends, and confirm the results tabulated just above. When significant in the table, the plots suggest that the change due to the ban is discrete and persists until the change in transfer tax that marks the end of our estimation window. We also observe a statistically significant decrease in net aggregate investor supply in the first two quarters after the ban, but this effect disappears later on (potentially due to new rental market reforms).

4.3 Prices

The ban's zero net aggregate effect on demand and a marginal one on supply indicate an ambiguous impact that, according to the conceptual framework, could result from an expected increase in neighborhood quality, as the policy intended. Under this assumption, the impact of the ban on property prices is also ambiguous. In Columns 1 and 2 of Table 5, we compare the evolution of log property prices ($p_{i,t}$) below the cap in Rotterdam's treated and control neighborhoods with the usual set of controls.[10] In Column 1 we find a positive but statistically insignificant treatment effect, indicating that the investor ban does not significantly alter house prices over the three quarters after the introduction of the ban. Since tax values more accurately predict house prices closer to the valuation date on January 1, 2021, we also add Log(appraisal value) $\times$ Quarter interactions in Column

[10]We exclude properties for which the purchase price pertains to two or three properties from the sample. We also remove about 1% of transactions where sale prices deviate significantly from predictions based on tax value and property characteristics. These extremely over- and underpriced transactions reflect stale or incorrectly updated prices.

Figure 3: Parallel Trends in Owner-Occupier and Investor Supply

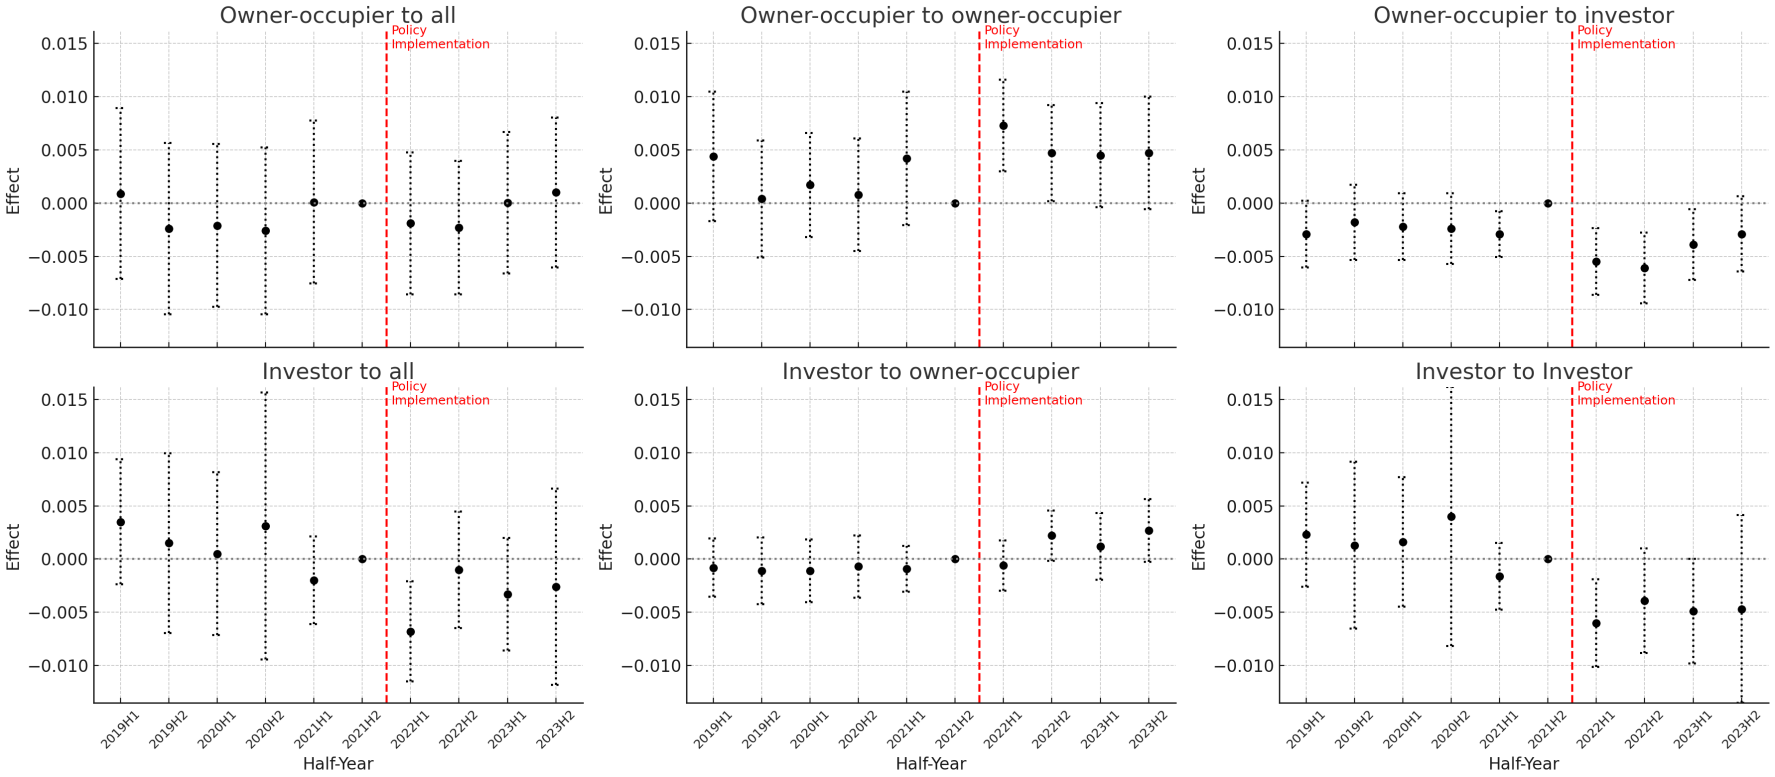

Notes: This Figure shows the impact over time of the investor ban on the likelihood of sales by owner-occupiers (top row) or by investors (bottom row), to owner-occupiers and investors (first column), owner-occupiers (second column) or investors (third column) on the matched sample of property purchases, using the stock of owner-occupied housing in matched neighborhoods from administrative data. We plot the difference between treatment and control properties by estimating the treatment coefficients of Table 4, columns 1-6, separately for each quarter. Vertical bars indicate 95% confidence intervals. Standard errors are clustered at the neighborhood level.

3. This does not significantly change the ban's price effect with a coefficient that is statistically significant only at the 14% level.

Table 5: Impact of the Ban on Property Prices

	(1)	(2)	(3)
Treated × Post	0.0125	0.0139	0.0143
	(0.0143)	(0.0100)	(0.0096)
Log(tax appraisal value)	No	Yes	Yes
Property controls	No	Yes	Yes
Neighborhood FE	Yes	Yes	Yes
Year-Month-of-sale FE	No	Yes	Yes
Log(appraisal value)×Quarter FE	No	No	Yes
Observations	3,967	3,967	3,967
R^2	0.36036	0.82954	0.83041

Notes:
This Table presents DiD estimates of the property pricing effect of the 2022 investor ban for all property transactions in matched neighborhoods, using transactions data from the Dutch Land Registry from 2021Q2-2022Q4. The dependent variable is log-transformed property price with different sets of controls across the columns, neighborhood fixed effects, and/or Year-Month-of-Sale fixed effects. Property controls include square meters, property type, building year, and energy label. All variables are detailed in Appendix Table A.1. Standard errors are clustered at the neighborhood level. Significance levels: * p <0.10, ** p <0.05, *** p <0.01.

Figure 4 provides more detail into the price impact dynamics and parallel trends. We find parallel, statistically insignificant trends in house prices in the period preceding the ban and a short-lived positive effect just after introduction, with an effect of around 3%. Note that the standard errors are fairly large because we observe relatively few transactions per neighborhood × quarter. Price effects seem to weaken over time and are no longer statistically discernible by late-2022. The short-lived price effect is consistent with the ban increasing the appetite of regular home-buyers for bidding on regulated neighborhoods. It fades possibly because of higher interest rates and the broader rental market reforms from the second half of 2022, both of which affect supply and demand in regulated and unregulated neighborhoods.

Figure 4: Parallel Trends in Property Prices

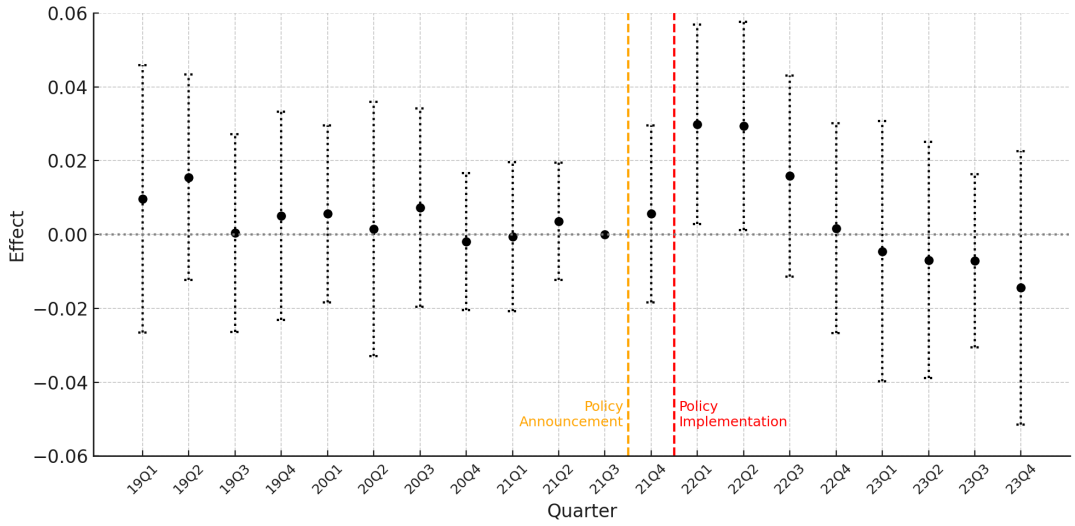

Notes:
This Figure shows the impact over time of the investor ban on property prices for all transactions in matched neighborhoods using Dutch Land Registry. We plot the difference between treatment and control properties by estimating the treatment coefficients of Table 5, column 3, separately for each quarter. Vertical bars indicate 95% confidence intervals. Standard errors are clustered at the neighborhood level.

4.4 Spillover effects

The measurements developed so far attempt to demonstrate the causal effect of the ban on treated neighborhoods relative to control neighborhoods. The difference between treated and control neighborhoods post-ban might not just come from a pure treatment effect on treated properties, but also from a spillover effect as investor reallocate their activity towards the control properties we use.

Spillovers are a reason for concern if investors who would have invested in treated properties are to allocate their money towards properties in the identified control groups instead of re-allocating over all other available assets. If there are spillover effects towards control areas, one would expect decreases for properties not in the control group, and increases for properties in the control group.

Table 3 shows that investor purchases of properties in Rotterdam decreased over 2022 for the control group by 5.6 percentage points. However, the share of investor purchases of properties in the same municipality that are not in the control group does not decrease

and actually increases from about 10% to 12%. These patterns suggest that spillover effects are limited. It appears that, while plenty of properties *not* subject to the policy can be bought, investors face important information asymmetries in unfamiliar markets and other frictions (e.g., maintenance over large distances).[11] Instead, investors may choose to reallocate their money towards other assets such as savings deposits or the stock market.

5 Impact of the ban on neighborhood change

Taking stock, we have documented that the investor ban reduces demand from private investors and increases demand from owner-occupiers resulting in a partially offset negative net change in aggregate demand. The supply of owner-occupiers to private investors pre-ban shifts to owner-occupiers post-ban, while the supply from investors slightly decreases without any offsetting. Yet, the resulting impact on prices is not negative, and in fact positive for two quarters.

Our conceptual framework rationalizes these results under the assumption that, as the ban increases the homeownership rate in treated neighborhoods, neighborhood quality is expected to increase. We call this mechanism residential sorting: Neighborhoods may gradually change as different people arrive, bringing their own preferences and resources.

Unfortunately, neighborhood quality lacks a precise definition, which makes it difficult to directly test the assumption. Instead, we track incremental changes in neighborhood composition that occur as the ban reallocates dwellings from investors to homeowners.[12] Our premise is that selling investor-owned units to resident buyers gradually alters the socio-economic composition of the neighborhood. As we will see, changes in the socio-

[11] As the market for buy-to-let investments is predominantly populated by retail investors, investors often own real estate in or close to the municipality where they live (Hochstenbach, 2022).

[12] In the long-term, accumulated transitions from renting to owning or vice versa could have general equilibrium effects on the resident composition in the neighborhood. For example, if homeowners are older than renters, a neighborhood with a large shifts towards owner-occupants might endogenously attract even more older residents if these prefer to live together. Over the nine-month horizon we look at in this paper, we can plausibly rule out such effects.

economic characteristics of new and counterfactual residents without the policy allow us to better understand possible effects on neighborhood quality.

Since neighborhood quality is multifaceted, we look at three margins through which a shift from investor- to owner-occupancy reshapes who enters the neighborhood, and the corresponding effect on neighborhood's socio-economic composition. We focus on (i) neighborhood affluence, as new residents may have different wealth and incomes; (ii) neighborhood homogeneity, as more transient renters may be replaced by more rooted owner-occupiers; and (iii) household mobility, noting that higher turnover typically signals lower neighborhood quality.

5.1 Wealth and Income

One key potential difference between renters and owner-occupiers residents is wealth. Since more affluent households could positively affect house prices and stimulate neighborhood amenities and services, neighborhood composition may change if the ban benefits buy-to-live residents with higher average wealth and income levels.

We first examine the effect of the ban on wealth and income, followed by determining how much of this effect can be attributed to the difference between buy-to-let vs. buy-to-live purchases. We use the specification in Eq. (12) except that our outcomes are now measured at the resident level (rather than the transaction level). The sample consists of residents in Rotterdam living in properties below the limit within our selection of treated and control neighborhoods.

Starting with the impact on wealth, we focus on wealth from sources other than real estate in order to compare renters with buyers. Table 6, Column 1 shows that for treated properties sold after the introduction of the ban, average non-real estate wealth increased by 0.28 log points relative to the control group. The result is significant at the 1% level. Column 1 indicates that banning investors increases the fraction of higher-wealth residents. In Column 2 we control for investor ownership. The Investor-owned coefficient indicates that on average, the wealth of residents of investor-owned properties

is about 0.67 log points lower compared to owner-occupiers who live in similar properties. The Investor-owned variable renders the treatment effect insignificant: Treated neighborhoods see a larger increase in average wealth after the ban was introduced, implying the treatment effect we find is directly explained by shifts in buy-to-live versus buy-to-let.

Table 6: Impact on Neighborhood Wealth and Income

	Dependent variable:			
	log(Non-real estate wealth)		Personal income, percentile	
Treated × Post	0.2757***	0.1699	4.986**	1.611
	(0.0996)	(0.0925)	(2.347)	(2.041)
Investor-owned		-0.6732***		-21.46***
		(0.0662)		(2.011)
Log(tax appraisal value)	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes
Neighborhood FE	Yes	Yes	Yes	Yes
Year-Month-of-sale FE	Yes	Yes	Yes	Yes
Observations	2,991	2,991	7,054	7,054
R^2	0.09380	0.12556	0.07974	0.16152
Within R^2	0.00876	0.04349	0.00368	0.09222

Notes:
This Table presents DiD estimates of wealth and income effects of the 2022 investor ban for all households involved in property transactions in matched neighborhoods, using administrative data from 2021Q2-2022Q4. In Columns 1 and 2, the dependent variable is log-transformed non-real estate wealth. In Columns 3 and 4, the dependent variable is the household's percentile in the national income distribution without the 1st and 100th percentiles. Columns 1 and 3 show baseline regressions without a dummy variable that equals one if the property is investor-owned, i.e., owned by non-residents; Columns 2 and 4 add in the dummy variable. Property controls include square meters, property type, building year, and energy label. All variables are detailed in Appendix Table A.1. Standard errors are clustered at the neighborhood level. Significance levels: * p <0.10, ** p <0.05, *** p <0.01.

The results in Columns 3 and 4 present similar findings for personal income percentiles.[13] We drop the top percentile, which consists of censored observations, and the bottom percentile, which contains many negative-income households. The effect of the ban on income is significant at the 5% level, and again fully explained by the fact that the ban removes purchases by investors who rent property to lower-income residents.

[13]We define income as individual income from salaried employment, income from self-employment or entrepreneurship, and income from retirement benefits or other social insurance schemes (e.g. unemployment benefit). Using log income levels for those reporting positive income generates similar results as for wealth.

The results are economically large. Throughout the sample, the average resident of an investor-bought property is in the 40th percentile of the income distribution. If the same home were purchased by an owner-occupier, the resident's predicted income would be in the 62nd percentile of the income distribution.

Our results are similar for household income in levels instead of percentiles, and indicate one margin along which residential sorting changes affected neighborhood and (in turn) improves neighborhood quality, i.e., the ban attracts new residents who are wealthier, with higher incomes.

5.2 Neighborhood Homogeneity

Besides the economic margin, another mechanism affecting neighborhood composition is homophily, i.e., the tendency of individuals to associate and bond with similar others. Assuming such a tendency, shifts in residential composition that lead to more similar group of residents may increase neighborhood desirability.

One potential measure for neighborhood homogeneity is whether the resident was born in the Netherlands or is a migrant. We examine this idea in Table 7, Column 1, showing that the ban significantly increases the likelihood that a resident is Dutch-born by about 7 percentage points. This is a sizable increase relative to the base rate of 12 percentage points. As before, this effect becomes insignificant after we control directly for investor purchases (Column 2), confirming that it is indeed the removal of investors that is driving this effect.

Another measure for neighborhood homogeneity is the distance between the current property of a resident and its previous property, with larger distance increasingly indicating an "outsider". Columns 3 and 4 indicate that the ban does not significantly reduce the moving distance to Rotterdam, but standard errors are wide. The coefficient becomes statistically significant in the country-level analysis which has a much larger number of observations (cf. Section 6). Residents of investor-owned properties on average do significantly move over about 0.3 log points longer distances.

Table 7: Impact on Neighborhood Homogeneity

	Dependent variable:					
Sample area:	Dutch-born (0/1)		Log moving distance		Nr. of Residents (> 0)	
Treated × Post	0.0713**	0.0275	-0.0168	0.0360	-0.1001*	-0.0545
	(0.0316)	(0.0302)	(0.1349)	(0.1354)	(0.0593)	(0.0597)
Investor-owned		-0.2044***		0.2876***		0.2849***
		(0.0164)		(0.0922)		(0.0319)
Log(tax appraisal value)	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes
Neighborhood FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month-of-sale FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	9,392	9,392	8,186	8,186	3,540	3,540
R^2	0.10167	0.13124	0.04007	0.04399	0.12978	0.14976
Within R^2	0.00237	0.03522	0.00014	0.00421	0.04811	0.06996

Notes:
This Table presents DiD estimates of neighborhood homogeneity effects of the 2022 investor ban for dwellings involved in property transactions in matched neighborhoods, using administrative data from 2021Q2-2022Q4. In Columns 1 and 2, the dependent variable is a dummy variable that equals one if the main resident is born in the Netherlands. In Columns 3 and 4, the dependent variable is the log-transformed distance between current and prior dwelling. In Columns 5 and 6, the dependent variable is the number of people that reside in the dwelling. Odd-numbered Columns show baseline regressions without a dummy variable that equals one if the property is investor-owned, i.e., owned by non-residents. Even-numbered Columns add in the dummy variable. Property controls include square meters, property type, building year, and energy label. All variables are detailed in Appendix Table A.1. Standard errors are clustered at the neighborhood level. Significance levels: * p <0.10, ** p <0.05, *** p <0.01.

A third measure for neighborhood homogeneity is the number of residents living in a property. Survey evidence finds that in neighborhoods with a higher concentration of room rentals, residents experience lower social cohesion.[14] The findings have been influential in that many municipalities have successfully prohibited the subdivision of dwellings into rooms for rent following legislative proposals for reasons of neighborhood quality.[15] We follow this measure for social cohesion by examining how the ban and buy-to-rent housing relate to the average number residents housed in affected properties.

For every property bought within the sample period we compute the number of residents registered at the property two years after the transaction date. We then analyze whether the policy affects the number of residents in a property, conditional on a property being occupied.[16]

The results in Columns 5 and 6 report our specification, this time at the property level rather than the resident level. In Column 5, we find that the ban leads to a reduction in the number of residents living together in a property, with about 0.1 fewer residents. The effect is significant at the 10% level and again disappears after we control directly for investor purchases in Column 6. Investor-owned properties in Rotterdam have 0.3 more residents, an increase of almost 15% relative to the typical property. Thus, properties are used more intensively after they are bought by investors, indicating that tenants in these properties more often share housing.

[14] Anecdotally, groups of residents in untreated neighborhoods in Rotterdam have been actively lobbying to have to policy extended towards their neighborhood. Their letters explicitly state that buy-to-let residents cause nuisances and limited social cohesion, e.g.: Letter of February 21, 2022 and Letter of 23 December 2021.

[15] The survey respondents associate property co-sharing with street littering, parking problems, illegal (over)occupation, and property neglect, all of which negatively impact neighborhood quality (Research Instituut Gebouwde Omgeving, 2022; Leefbaarometer 2022). See https://www.rigo.nl/wp-content/uploads/2022/03/Rapport-leefbaarheid-en-kamerbewoning-v3.pdf; https://open.overheid.nl/documenten/013bdf32-57d3-4a4a-8d45-17ca5d9c676a/file

[16] Due to under-registration, these estimates might underestimate the total number of residents in investor-owned properties. There are limits to the number of adults that can be registered legally on a property, which means that in some cases residents are living in properties while being registered on another address. It is also well-known that labor migrants often do not register their residency correctly, and they are not mandated to do so if they stay less than four months. This implies that underestimation is particularly common for private rental properties. Investor-owned properties by definition have higher vacancy rates because we define a property as owner-occupied when the buyer moves into the property, so we cannot determine the effect of the policy on vacancy rates.

These results suggest that neighborhood homogeneity is a second residential-sorting margin that changes neighborhood composition and, through that channel, could gradually raise neighborhood quality.

5.3 Housing Mobility

While improved neighborhood quality is evidently beneficial to incumbent residents, the nature of buy-to-let housing is such that it benefits "outsider" residents. People who move to a new city from further away usually start in private rental housing, which offers more flexibility and can also be arranged without a stable employment contract. Similarly, co-living is typical for students or labor immigrants, both of whom tend to move frequently, especially when coming from abroad, due to the temporary nature of their work or study. Thus, the private rental sector adds to household mobility. Given that housing is a scarce good, the ban might affect how these groups get access to affordable housing and can rely on buy-to-let for a place to live.

Housing mobility captures the degree to which people move between or within households due to housing instability, affordability, or transitional living. Thus, to test whether the private rental sector provides a temporary safety net for people who cannot (yet) access more stable forms of housing, a first proxy for household mobility is the occurrence of the life events that tend to initiate a move. We measure life events by a dummy variable that equals one for residents who form a new household in the Netherlands after moving into a sold property. These include individuals who move in from abroad, individuals who move out after a divorce, and individuals who move out after living with other family members or in institutional properties (e.g. prison or youth homes). We exclude children from this analysis who cannot yet form independent households.

Column 1 of Table 8 indicates the ban leads to a 5% lower fraction of residents who moved into the property after a life event. This result is just shy of significance at the 5% level. Column 2 demonstrates that residents of investor owned properties are different along this dimension too, and 17% more likely to cater to residents who form

new household. The difference between buy-to-let and buy-to-live again entirely explains the impact of the ban.

We also measure housing mobility ex post, i.e., by the extent people who rent from private investors also move away relatively quickly. We find that the ban significantly reduces the fraction of residents who move out of the neighborhood within two years by about 9 percent (Column 3). In Column 4, residents of investor-owned properties are 0.33 percent more likely to move away within this period, an effect that again entirely explains the impact of the ban.

These results suggest that household mobility is a third residential-sorting margin that changes neighborhood composition and, through that channel, *lowers* neighborhood quality. In addition, they suggest that improving neighborhood quality can be a double-edged sword. An increase in residents who move quickly implies increased housing mobility, but high residential turnover may result in lower neighborhood homogeneity and hence reduced neighborhood quality. A similar tradeoff applies to the "Nr. of Residents" measure from Table 7 discussed in the previous Subsection: While Rotterdam mentions limiting home sharing as one reason the ban was to improve home quality, limiting home sharing removes low cost housing and reduces supply by reducing the number of tenants per property.

6 External validity analyses

By the end of 2022, 34 Dutch municipalities varying from small towns to all major cities with more than 200,000 inhabitants had implemented a buy-to-let ban, see Appendix Table A.1.[17]

While all cities had a political majority to implement the ban and applied the ban city-wide to all properties below the ban's tax value threshold, they differ in the required

[17]We exclude the eight smallest, mainly rural municipalities with fewer than 400 treated sales around the introduction date. Most municipalities applied the ban *city-wide* and used the national mortgage-insurance limit (€355,000 in 2022) as the value cap; coverage therefore falls between the 40th–60th percentile of the local price distribution, capturing roughly 90% of pre-ban buy-to-let activity.

Table 8: Impact on Housing Mobility

	Dependent variable:			
Sample area:	Life event (>0)		Move within 2y (0/1)	
Treated × Post	-0.0508*	-0.0147	-0.0883**	-0.0205
	(0.0254)	(0.0264)	(0.0432)	(0.0420)
Investor-owned		0.1743***		0.3276***
		(0.0145)		(0.0221)
log(tax appraisal Value)	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes
Neighborhood FE	Yes	Yes	Yes	Yes
Year-Month-of-sale FE	Yes	Yes	Yes	Yes
Observations	8,463	8,463	7,084	7,084
R^2	0.06027	0.08624	0.11826	0.19438
Within R^2	0.00245	0.03002	0.00538	0.09124

Notes: This Table presents DiD estimates of neighborhood homogeneity effects of the 2022 investor ban for households involved in property transactions in matched neighborhoods, using administrative data from 2021Q2-2022Q4. In Columns 1 and 2, the dependent variable is a dummy variable that equals one if at least one resident has experienced a life event. In Columns 3 and 4, the dependent variable is a dummy variable that equals one if the household moves out within two years after moving in. Odd-numbered Columns show baseline regressions without a dummy variable that equals one if the property is investor-owned, i.e., owned by non-residents. Even-numbered Columns add in the dummy variable. Control variables include property controls, neighborhood fixed effects, Year-Month-of-Sale fixed effects, energy label, and housing type. All variables are detailed in Appendix Table A.1. Standard errors are clustered at the neighborhood level. Significance levels: * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

political and administrative procedures leading to staggered investor ban introduction dates that can produce biased estimates when treatment effects are heterogeneous over time or across municipalities (Goodman-Bacon, 2021). Therefore, we identify the treatment effect by employing the estimator developed by Deb et al. (2024) designed for such heterogeneous treatment effects in the context of our data, a repeated cross-section. The identifying assumption is that these differences are plausibly exogenous to any other shocks that affected the local activity of investors over the course of 2022.

The Deb et al. (2024) estimator saturates a single OLS regression with interaction terms between time periods and treatment groups, effectively absorbing all group- and time-specific heterogeneity. Similar to Wooldridge (2021), it compares the outcomes of "ever-treated" units before and after treatment with those of "never-treated" units over the same period, which avoids biases common in traditional two-way fixed effects models under heterogeneous treatment effects. The Deb et al. (2024) estimator differs from the one in Wooldridge (2021) in that it consistently estimates staggered treatment effects by using outcome means and co-variates from repeated cross-sections rather than requiring panel data.

Appendix Table A.2 shows that the ban reduced the probability that a treated dwelling was bought by an investor by 12–13 pp and increased first-time buyer purchases by about 9 percentage points (pp).[18] Relative to Rotterdam (Table 6, Columns 1–4), the national demand response is smaller in magnitude (Rotterdam: −19 pp and +13 pp). As a large metropolitan area, Rotterdam had the highest pre-ban investor share leaving more room for contraction.

Appendix Table A.3 and Figure A.2 confirm that treatment effects are roughly monotonic in pre-ban investor activity: cities in the top quartile of investor shares show reductions of 15–18 pp, whereas smaller markets show single-digit effects. Allowing the effect to vary by both cohort and time leaves the average treatment effect on the treated (ATET) virtually unchanged: Appendix Table A.3, Column 5 replaces the single DiD

[18] Coefficients come from the flexible TWFE column that allows for municipality-specific effects.

term with ten cohort-by-quarter interactions. Weighting those effects by the number of treated sales in each cohort–quarter yields an average ATET $\approx$ –0.119, i.e. an 11.9 percentage point fall in the probability that a treated home is bought by an investor. This is virtually identical to the –11.8 p.p. reported in the simpler group-time model in Column 4, suggesting that treatment heterogeneity over the nine-month roll-out is fairly small.

Appendix Table A.4 replicates the neighborhood composition results (Tables 6–8) at the national level. Similar patterns emerge: (i) Non-housing wealth rises by ≈ 0.22 ln-points nationally versus 0.28 ln-points in Rotterdam; income percentiles shift right by +5 versus +7; (ii) the share of Dutch-born residents increases by 3 pp nationally versus 7 pp in Rotterdam, while the decrease in average migration distance is significant in the national sample only; (iii) life-event entries and two-year exit rates fall in both samples, with the national coefficients (-3 pp and -3 pp) two-thirds of the Rotterdam estimates.

Collectively, the buy-to-let ban generates qualitatively similar shifts in buyer composition and resident characteristics at the national level, with coefficients in identical directions but smaller, and stronger statistical significance. The latter may be due to up to 10$\times$ more observations in the national analysis. Most importantly, the Rotterdam results extend to all of the country's municipalities with substantial buy-to-let activity.

7 Conclusion

Leveraging the 2022 Dutch buy-to-let ban as a plausibly exogenous shock to tenure composition, we aim to isolate the effect of buy-to-let investors on housing-market outcomes. Our matched difference-in-differences estimates at the property level suggest that removing investors shifts demand from investors to first-time buyers and slightly reduces investor supply, while property prices do not decrease and in fact shortly *increase* by 2–3 percent for two quarters. Individual-registry data reveal residential sorting as an important contributing mechanism: treated dwellings are re-occupied by households that are wealthier, more locally rooted, and significantly less transient than the renters they

displace. Our evidence suggests that the ban's neighborhood effects are primarily driven by changes in occupant identity; after netting out this sorting channel, we find little residual treatment effect.

The policy appears to have intended and unintended consequences. Our results imply that such interventions may simultaneously raise prices and contribute to gentrification, even in the absence of physical upgrading. However, by barring landlords without expanding aggregate supply, the ban is likely to limit access for groups dependent on flexible rental housing such as lower-income, foreign-born households in transient living conditions. Optimal housing policy may therefore need to weigh the welfare gains from higher owner-occupancy against reduced housing mobility.

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.

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.

Barbieri, F. and G. Dobbels (2025). Market power and the welfare effects of institutional landlords. Working Paper.

Chang, K. (2025). Diversifying the suburbs: Rental supply and spatial inequality. Working Paper.

Collinson, R. and P. Ganong (2018). How do changes in housing voucher design affect rent and neighborhood quality? *American Economic Journal: Economic Policy 10*(2), 62–89.

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

Coven, J. (2023). The impact of institutional investors on homeownership and neighborhood access. *Available at SSRN*.

Deb, P., E. C. Norton, J. M. Wooldridge, and J. E. Zabel (2024). A flexible, heterogeneous treatment effects difference-in-differences estimator for repeated cross-sections. Technical report, National Bureau of Economic Research.

Diamond, R. and T. McQuade (2019). Who wants affordable housing in their backyard? an equilibrium analysis of low-income property development. *Journal of Political Economy 127*(3), 1063–1117.

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.

DiPasquale, D. and E. L. Glaeser (1999). Incentives and social capital: Are homeowners better citizens? *Journal of urban Economics 45*(2), 354–384.

Francke, M., L. Hans, and M. Korevaar. (2024). Beleggers verkopen nu meer huizen aan eigenaarbewoners dan ze opkopen. *Economisch-Statistische Berichten 109*(4833), 218–221.

Garriga, C., P. Gete, and A. Tsouderou (2023). The economic effects of real estate investors. *Real Estate Economics 51*(3), 655–685.

Goodman, L. and E. Golding (2021). Institutional investors have a comparative advantage in purchasing homes that need repair. *Urban Institute, October 20*.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*(2), 254–277.

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.

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 50*(6), 1425–1461.

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 47*(2), 399–430.

Pennington, K. (2021). Does building new housing cause displacement?: The supply and demand effects of construction in San Francisco. *SSRN Electronic Journal*.

Rossi-Hansberg, E., P.-D. Sarte, and R. Owens III (2010). Housing externalities. *Journal of political Economy 118*(3), 485–535.

Rouwendal, J., F. Sniekers, and N. Jia (2023). Borrowing constraints, housing tenure

choice and buy-to-let investors: An assignment model. Technical report, Tinbergen Institute Discussion Paper.

Sodini, P., S. Van Nieuwerburgh, R. Vestman, and U. von Lilienfeld-Toal (2023). Identifying the benefits from homeownership: A swedish experiment. *American Economic Review 113*(12), 3173–3212.

Van Bekkum, S., M. Gabarro, R. M. Irani, and J.-L. Peydró (2022). Take it to the limit? the effects of household leverage caps.

Wooldridge, J. M. (2021). Two-way fixed effects, the two-way Mundlak regression, and difference-in-differences estimators. *Available at SSRN 3906345*.

A Supplementary Tables and Figures

Table A.1:

Variable	**Definition**
Tax Appraisal Value	(Dutch: *WOZ-Waarde*). Estimated market value of the property on January 1 of the previous year, assuming the property is not rented and fully owned.
Square meters	The usable surface of a property in square meters.
Property type	Classifies property into apartment, row house, corner house, semi-detached house or detached house.
Building year	Construction year of the property.
Energy label	Energy efficiency score (letter from A to G) at moment of sale
Investor	An owner of a property that does not move to the property within 2 years of purchasing it
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.
Personal income, percentile	Percentiles calculated from the Dutch income distribution. Gross income of the person minus paid income transfers such as taxes on income and wealth, and social or government insurance premiums.
Dutch-born	Dummy variable for whether an individual resident holds a nationality that is not Dutch.
Age	Age of the individual on December 31, 2022.
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 1 years of moving in.
Distance to previous residence	Distance between current and previous residence in meters, as the crow flies.

Table A.2: Determinants of investor purchases

	Dependent variable:			
	Investor Purchases			
	Baseline	Location FE	Month-of-sale FE	Time+Location FE
log(Tax appraisal Value)	-0.2669***	-0.4409***	-0.2574***	-0.4282***
	(0.0256)	(0.0227)	(0.0285)	(0.0265)
Square meters	0.0007***	0.0017***	0.0006***	0.0016***
	(0.0001)	(0.0002)	(0.0001)	(0.0002)
Property controls	Yes	Yes	Yes	Yes
Neighborhood FE	No	Yes	No	Yes
Year-Month-of-sale FE	No	No	Yes	Yes
S.E. type	Heteroskedasticity-robust			
Observations	13,947	13,947	13,947	13,947
R^2	0.197	0.246	0.235	0.279

Table A.2 presents the determinants of whether a property is bought by investors in Rotterdam during 2019 and 2020 based on property characteristics, both with and without neighborhood fixed effects, and with or without time fixed effects. We find a sizable role both for hedonic property characteristics such as building year, property type and size as well as location characteristics.

Figure A.1: Treated and control neighborhoods in Rotterdam

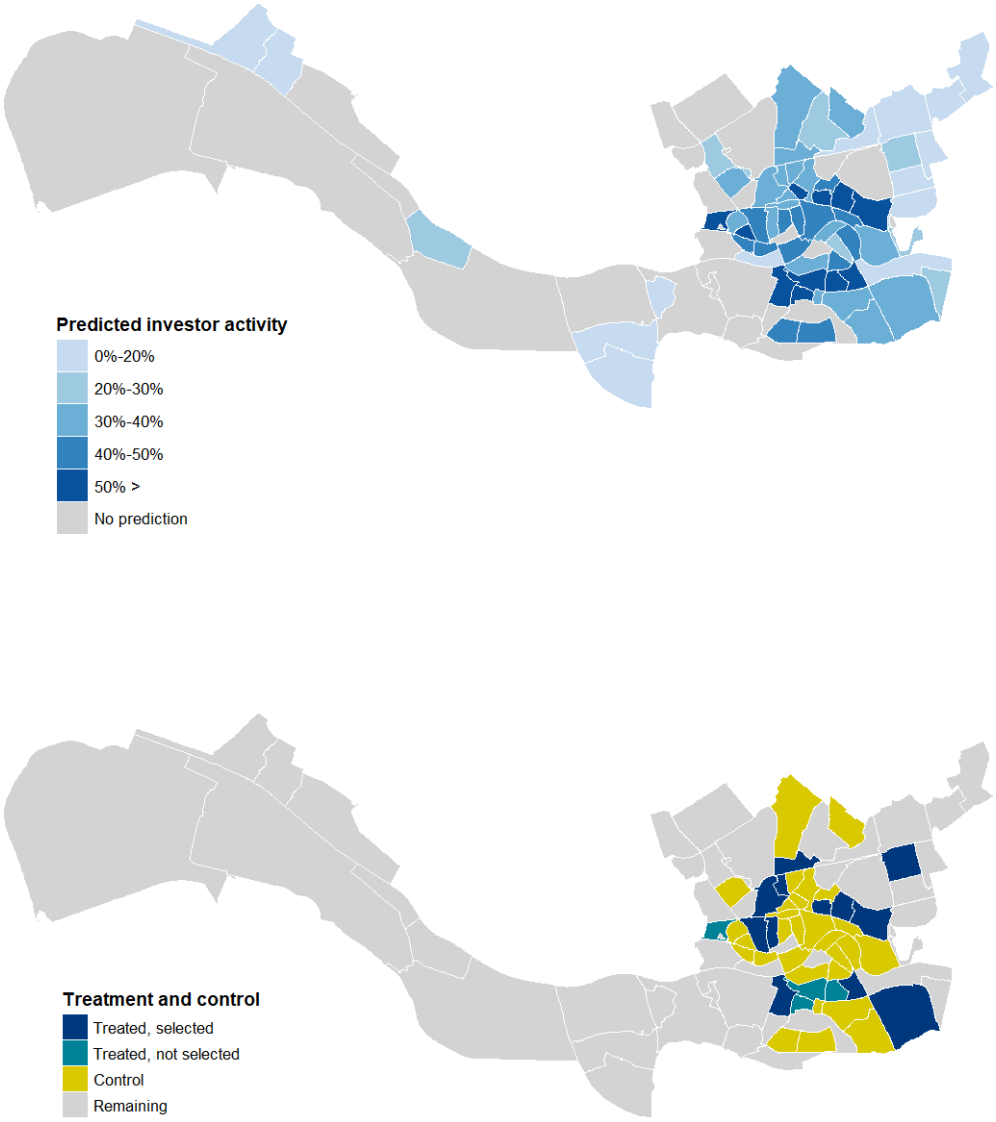

Figure A.1 plots all neighborhoods in the Rotterdam urban area to illustrate how we separate treatment and outcome variables. The top panel plots the share of predicted investor purchases in 2022, based on prediction model parameters estimated from 2019-2021, for different investor activity levels. Investor share is defined by the Dutch Land Registry. The bottom panel plots the resulting treated neighborhoods (those with ban introduced on 1 Jan 2022), split into neighborhoods selected for use in our analyses (dark blue) and neighborhoods not selected for use in our analyses (dark green). The yellow areas indicate control neighborhoods.

B Further Evidence: Estimation Window & Treatment

Figure A.1 plots national trends in the number of investor purchases defined as any purchase where the buyer does not end up residing in the property within two years. The buy-to-let ban prevents investors from renting out these properties after purchase. We separate two types of properties: "treated" properties are properties which became subject to the buy-to-let ban at some point after January 1, 2022, and "untreated" properties are all other properties.

Figure A.1: Investor purchases of properties, Netherlands

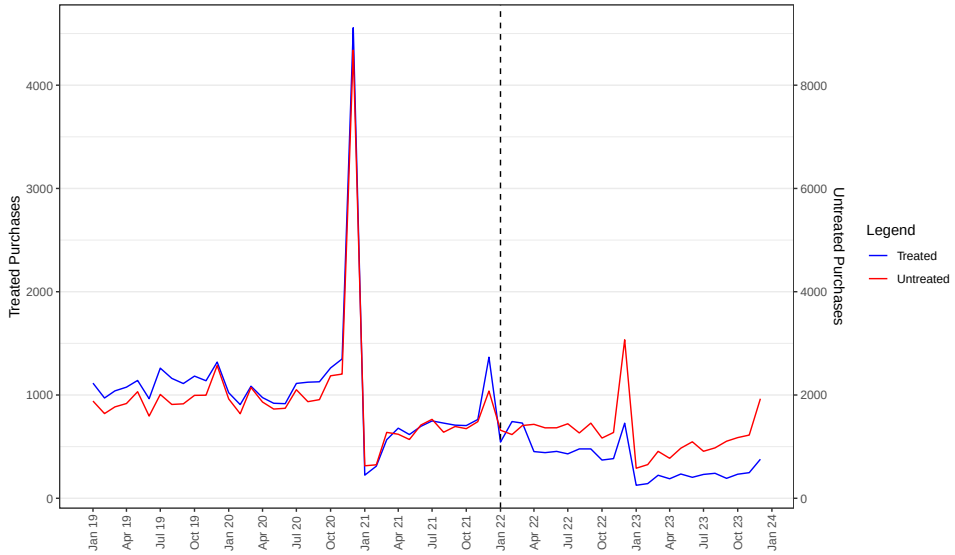

Notes: This Figure plots monthly number of investor purchases of properties sold by owner-occupiers or sold by investors when vacant. We separate properties that at some point in 2022 or 2023 fall under the new buy-to-let ban (in red) and untreated properties that do not (in blue). The dashed black line refers to the introduction of the policy that allows municipalities to start banning buy-to-let investments. We observe parallel trends between both groups up to this point.

We see a large spike in investor purchases in December 2020 and very few purchases in the first quarter of 2021, as investors were anticipating on the large increase in the stamp duty on January 1, 2021. Due to the higher tax, investor purchases do not recover to their pre-2021 levels.

In December 2021, we see again a spike in investor activity. The 2021 spike is partially seasonal and partially reflects anticipation of the buy-to-let ban: we find a larger increase in investor activity among treated properties. Throughout 2022 and 2023 we observe a significant decrease in investor purchases of treated properties that is not visible for untreated properties: the effect of the buy-to-let ban is directly visible from the raw data.

In December 2022, we see again a third and final spike in investor purchases and a decrease afterwards as investors anticipated the second increase in stamp duty by January 1, 2023. Investment activity remains lower through most of 2023. Considering the two

increases in stamp duty, we use the three quarters before the ban introduction and the year afterwards as the period over which we can cleanly measure the effect of the ban.

C Full model derivations

This appendix provides detailed derivations of the partial derivatives used in the main analysis.

C.1 Linearized demand

We approximate aggregate demand around the pre-ban equilibrium (P_0, θ_0) using a first-order Taylor expansion:

$$\Delta D \approx \left(\frac{\partial D_t}{\partial P_t}\right)_0 \Delta P + \left(\frac{\partial D_t}{\partial \theta_t}\right)_0 \Delta\theta, \tag{13}$$

where $\Delta D = D_t - D_0$, $\Delta P = P_t - P_0$, and $\Delta\theta = \theta_t - \theta_0$. We interpret the Taylor expansion as a linear demand function of the form $D_t = d + eP_t + f\theta_t$, or $\Delta D = e\Delta P + f\Delta\theta$ in first-differences.

C.2 Sign of e

Recall the aggregate demand function:

$$D_t = \frac{N\alpha\theta_t Y}{P_t(1+\alpha\theta_t)}. \tag{14}$$

Differentiating $e = \left(\frac{\partial D_t}{\partial P_t}\right)_0$ with respect to P_t:

$$\frac{\partial}{\partial P_t}\left(\frac{\alpha\theta Y}{P_t(1+\alpha\theta_t)}\right) = \frac{\alpha\theta_t Y}{(1+\alpha\theta_t)} \cdot \left(\frac{-1}{P_t^2}\right) = -\frac{\alpha\theta_t Y}{P_t^2(1+\alpha\theta_t)}$$

Thus,

$$\frac{\partial D_t}{\partial P_t} = -\frac{N\alpha\theta Y}{P_t^2(1+\alpha\theta_t)} = -\frac{D_0}{P_0}$$

reflecting an inverse relationship between price and demand with unitary price elasticity. Since both terms are positive, the partial derivative is negative.

C.3 Sign of f

From the same aggregate demand function we differentiate with respect to θ:

$$\frac{\partial D_t}{\partial \theta_t} = \frac{N\alpha Y}{P_t} \cdot \frac{d}{d\theta_t}\left(\frac{\theta_t}{1+\alpha\theta_t}\right) = \frac{N\alpha Y}{P_t}\left(\frac{1\cdot(1+\alpha\theta_t) - \theta_t\cdot\alpha}{(1+\alpha\theta_t)^2}\right) = \frac{N\alpha Y}{P_t}\left(\frac{1}{(1+\alpha\theta_t)^2}\right).$$

At the equilibrium point:

$$f = \left(\frac{\partial D_t}{\partial \theta_t}\right)_0 = \frac{N\alpha Y}{P_0(1+\alpha\theta_0)^2} = \frac{D_0}{\theta_0(1+\alpha\theta_0)}. \tag{15}$$

Thus, coefficient f indicates that improving neighborhood quality increases demand, but the effect diminishes as the product $\alpha\theta_0$ increases. Since both numerator and denominator are positive, the partial derivative is also positive.

C.4 Linearized supply

We linearize S_t around the pre-ban equilibrium (P_0, θ_0) using a first-order Taylor expansion,

$$\Delta S \approx \left(\frac{\partial S_t}{\partial P_t}\right)_0 \Delta P + \left(\frac{\partial S_t}{\partial \theta_t}\right)_0 \Delta\theta,$$

where $\Delta S = S_t - S_0$, $\Delta P = P_t - P_0$, and $\Delta\theta = \theta_t - \theta_0$. The Taylor expansion can be interpreted in terms of a linear supply function $S_t = a - bP_t + c\theta_t$, which becomes $\Delta S = b\Delta P + c\Delta\theta$ in first-differences.

C.5 Sign of *b*

The aggregate supply function is:

$$S_t = N_I \mathcal{G}_{\mathcal{P}}(P_t - \delta\beta_V\theta_t), \tag{16}$$

where $\mathcal{G}_{\tilde{\mathcal{P}}}(\cdot)$ is the CDF of reservation prices. Compute the partial derivative with respect to P_t:

$$\frac{\partial S_t}{\partial P_t} = N_I \frac{\partial \mathcal{G}(P_t - \delta\beta_V\theta_t)}{\partial P_t} = N_I g_p(P_t - \delta\beta_V\theta_t), \tag{17}$$

where $g_p(\cdot)$ is the probability density function (PDF) of reservation prices. At the equilibrium point:

$$b = \left(\frac{\partial S_t}{\partial P_t}\right)_0 = N_I g_p(P_0 - \delta\beta_V\theta_0). \tag{18}$$

where $P_0 - \delta\beta_V\theta_0$ represents the adjusted equilibrium price that investors use as their reservation price after accounting for the influence of neighborhood quality. Since where $g_p(\cdot)$ is a PDF, it is always positive.

C.6 Sign of *c*

From the same aggregate supply function we differentiate with respect to θ_t:

$$\frac{\partial S_t}{\partial \theta_t} = N_I \frac{\partial \mathcal{G}(P_t - \delta\beta_V \theta_t)}{\partial \theta_t} = -N_I \delta \beta_V g_p(P_t - \delta\beta_V \theta_t). \tag{19}$$

At the equilibrium point:

$$c = \left(\frac{\partial S_t}{\partial \theta_t}\right)_0 = -N_I \delta \beta_V g_p(P_0 - \delta\beta_V \theta_0), \tag{20}$$

which directly contains the term from the previous subsection. Since this term, as well as δ and β_V, are always positive, the partial derivative is always negative.

D Replication of Rotterdam Results at National Level

Table A.1: 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
De Bilt	43,508	6-Oct-22	487,000	634,801	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
Diemen	31,822	01-Jul-22	512,000	465,573	Full
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
Nieuwegein	179,100	04-May-22	487,000	364,849	Full
Nijmegen	179,100	15-Nov-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
Sliedrecht	25,597	01-Jun-21	355,000	303,720	Full
Stichtse Vecht	65,240	01-Jul-22	487,000	495,126	Full
Zwijndrecht	65,240	01-Sep-22	355,000	306,639	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.

Table A.2: Impact of All 2022 Investor Bans on National Investor and First-time-Buyer Demand

	Investor buyer (0/1)		*1st-time owner-buyer*	
Estimator:	*TWFE*	*Deb et al. (2024)*	*TWFE*	*Deb et al. (2024)*
Treated × Post	-0.118***	-0.130***	0.086***	0.090***
	(0.015)	(0.015)	(0.016)	(0.010)
log(Tax appraisal Value)	No	Yes	No	Yes
Property controls	No	Yes	No	Yes
Neighborhood FE	Yes	Yes	Yes	Yes
Year-Month-of-sale FE	Yes	Yes	Yes	Yes
Observations	56,359	56,359	50,367	50,367
R^2	0.198	0.202	0.131	0.187

Notes:
This table presents DiD estimates of the demand effect of the 2022 investor ban for all property transactions in the Netherlands, using transactions data from the Dutch Land Registry from 2021Q2-2022Q4. In Columns 1 and 2, the dependent variable is a dummy variable that equals 1 if the buyer is an investor. In Columns 3 and 4, the dependent variable is a dummy variable that equals 1 if the buyer is a first-time owner-buyer. Columns 1 and 3 show two-way fixed effects regressions (TWFE) with neighborhood and time fixed effects; The coefficients in Columns 2 and 4 are based on a treatment-weighted average of city-level treatment effects following Deb et al. (2024) and add appraisal value and property controls. Property controls include square meters, property type, building year, and energy label. All variables are detailed in Appendix Table A.1. Standard errors are clustered at the municipality level. Significance levels: * p <0.10, ** p <0.05, *** p <0.01.

Table A.3: Impact of All 2022 Investor Bans on National Investor and Owner-Occupier Supply

	Dependent variable: Investor dummy				
Estimator:	*TWFE*		*Deb, Norton, Wooldridge & Zabel (2024)*		
Heterogeneity:	None	None	Group	Group-Time	Cohort-Time
Treated × Post	-0.125***	-0.118***	-0.118***	-0.119***	
	(0.016)	(0.015)	(0.007)	(0.010)	
Treated × Q1 2022 × Cohort Q1					-0.165***
					(0.024)
Treated × Q2 2022 × Cohort Q1					-0.153***
					(0.021)
Treated × Q3 2022 × Cohort Q1					-0.174***
					(0.026)
Treated × Q4 2022 × Cohort Q1					-0.151***
					(0.030)
Treated × Q2 2022 × Cohort Q2					-0.081***
					(0.013)
Treated × Q3 2022 × Cohort Q2					-0.077***
					(0.018)
Treated × Q4 2022 × Cohort Q2					-0.065***
					(0.022)
Treated × Q3 2022 × Cohort Q3					-0.084***
					(0.026)
Treated × Q4 2022 × Cohort Q3					-0.065***
					(0.023)
Treated × Q4 2022 × Cohort Q4					-0.071***
					(0.030)
log(tax appraisal Value)	No	Yes	Yes	Yes	Yes
Property controls	No	Yes	Yes	Yes	Yes
Neighborhood FE	Yes	Yes	Yes	Yes	Yes
Year-Month-of-sale FE	Yes	Yes	Yes	Yes	Yes
Observations	56,359	56,359	56,359	56,359	56,359
R^2	0.138	0.198	0.203	0.203	0.200

Notes:
This table presents DiD estimates of the supply effect of the 2022 investor ban for all property transactions in the Netherlands, using the entire stock of owner-occupied housing from administrative data from 2021Q2-2022Q. In Columns 1 and 2, the dependent variable is a dummy variable that equals 1 if the buyer is an investor. In Columns 3 and 4, the dependent variable is a dummy variable that equals 1 if the buyer is a first-time owner-buyer. Column 1 uses standard two-way fixed-effects without control variables, Column 2 is similar to Column 1 but adds control variables. Column 3 presents coefficients based on a treatment-weighted average of treatment effects that are heterogeneous across municipalities following Deb et al. (2024). Column 4 is similar to Column 3 but allows for effects heterogeneous across municipalities and over time. Column 5 is similar to Column 4 but allows for effects heterogeneous across municipalities, over time, and cohorts based on each municipality's introduction date of the ban. Property controls include square meters, property type, building year, and energy label. All variables are detailed in Appendix Table A.1. Standard errors are clustered at the municipality level. Significance levels: * p <0.10, ** p <0.05, *** p <0.01.

Figure A.2: Heterogeneity Across Cities

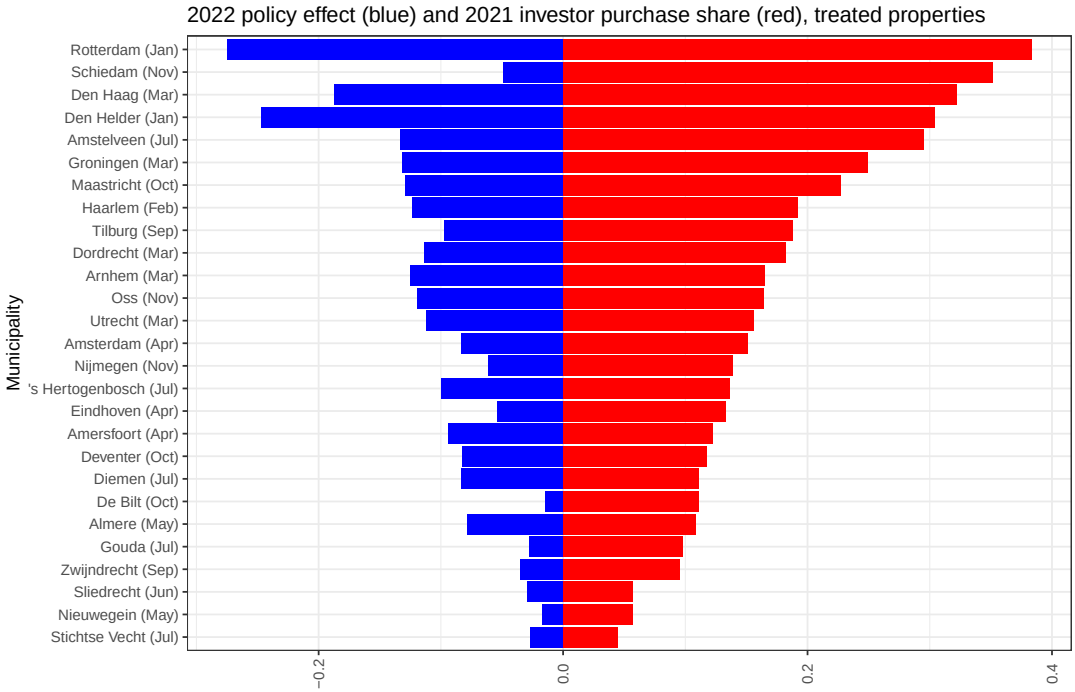

Notes: This Figure shows the treatment coefficients on the individual cities in the national specification in the left side of the panel in blue, and the average share of investors purchases in 2021 in red on the right side. All these estimates are based on properties treated by the law alone. We report also the month of introduction in brackets behind the name of each municipality. The figure shows the reduction due the policy is proportional to the amount of ex-ante investor purchases on treated properties in each municipality.

Table A.4: Impact of All 2022 Investor Bans and Residential Sorting Across Dutch Property Markets

	Dependent variable:													
	Log(non-RE wealth)		Income pctile		Dutch-born (0/1)		Log(distance)		#Residents		Life event (>0)		Move <2y	
Treated × Post	0.133***	0.050	2.418***	0.218	0.033***	0.005	-0.092***	-0.041	-0.033**	-0.001	-0.029***	-0.003	-0.080***	-0.033***
	(0.035)	(0.032)	(0.484)	(0.438)	(0.010)	(0.009)	(0.033)	(0.034)	(0.014)	(0.014)	(0.007)	(0.007)	(0.007)	(0.006)
Investor-owned		-0.829***		-17.530***		-0.158***		0.346***		0.209***		0.152***		0.319***
		(0.007)		(0.122)		(0.002)		(0.004)		(0.012)		(0.002)		(0.002)
Log(tax appraisal value)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Neighborhood FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month-of-sale FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	40,786	40,786	93,594	93,594	119,027	119,027	107,730	107,730	48,287	48,287	107,819	107,819	93,772	93,772
R^2	0.043	0.077	0.117	0.166	0.213	0.235	0.073	0.081	0.166	0.180	0.084	0.104	0.142	0.206

Notes:
This table presents DiD estimates of wealth and income effects, neighborhood homogeneity effects, and household mobility of all 2022 investor bans for property transactions participating municipalities, using administrative data from 2021Q2-2022Q4. In Columns 1-4, the dependent variables correspond to Table 6. In Columns 5-10, the dependent variables correspond to Table 8. In Columns 11-14, the dependent variables correspond to Table 8. All Columns are based on a treatment-weighted average of city-level treatment effects following Deb et al. (2024) and add appraisal value and property controls. Odd-numbered Columns show baseline regressions without a dummy variable that equals one if the property is investor-owned, i.e., owned by non-residents. Even-numbered Columns add in the dummy variable. Property controls include square meters, property type, building year, and energy label. All variables are detailed in Appendix Table A.1. Standard errors are clustered at the municipality level. Significance levels: * p <0.10, ** p <0.05, *** p <0.01.