Real-Estate Investors, House Prices and Rents: Evidence from Capital-Gains Tax Changes*

Itai Ater
Tel Aviv University & CEPR

Yael Elster
University of Haifa

Eran B. Hoffmann
The Hebrew University of Jerusalem

February 15, 2024

Abstract

We study the dual role of real-estate investors – households who own multiple housing units – in ownership and rental housing markets. Exploiting a series of capital-gains tax changes and rich administrative data from Israel, we first show that sales by real-estate investors increased by 50% after they were unexpectedly exempt from paying capital-gains taxes. Investors predominantly sold housing units that they had been renting out, and these units were sold to first-time homeowners. Next, we exploit spatial variations in the ownership distribution to examine how sales induced by the tax changes affected house prices and rents. We find that a 1 percentage point increase in investors' semi-annual sales rate decreases house prices by 14% and increases rents on new leases by 10%. Prices of smaller and older units, in which investors own a larger share of the stock, respond more to sales. Finally, we propose a dynamic heterogeneous-agent model of the housing market that features a lock-in effect due to taxes and downward-sloping demand curves for owned and rented housing, which are consistent with our main findings. The results suggest that policies that encourage investors to sell can achieve their stated objective of reducing house prices, but may also restrict the supply of rental housing and increase rents.

JEL: R31, R38, G12, H20

Keywords: real-estate investors, rental market, housing market, capital-gains tax, lock-in effect

*Ater: Tel-Aviv University Coller School of Management (email: ater@post.tau.ac.il); Elster: University of Haifa, Department of Economics (email: yelster@econ.haifa.ac.il); Hoffmann: The Hebrew University of Jerusalem, Department of Economics (email: eran.hoffmann@mail.huji.ac.il). We thank Alina Arefeva, Danny Ben-Shahar, David Genesove, Shao Lei (discussant), Wenli Li (discussant), Evan Mast, Tsur Sommerville (discussant), Yichen Su, Charles Nathanson (discussant) and seminar and conference participants at Tel-Aviv University, Hebrew University, Bar-Ilan University, University of Haifa, Stanford Institute for Theoretical Economics (SITE) Conference on Housing and Urban Economics 2021, AREUEA National Conference 2021, ESCP-TAU-UCLA Conference on low-income housing supply and housing affordability 2022, NBER summer institute 2022 (Real Estate), Syracuse - Chicago Webinar Series on Property Tax Administration and Design 2022, for helpful comments. Financial support from the Falk Institute and the Alrov Institute for Real Estate Research is gratefully acknowledged.

1 Introduction

A large share of the total housing wealth is held by households who own multiple housing units. These households, which we refer to as "real-estate investors", are also the primary suppliers of rental units in many countries. Therefore their decisions to buy or sell affect both the homeownership rate and the supply of rental housing. Investors' activity has increased over the past two decades in several countries and triggered a debate on their impact on housing markets.[1] A common view is that investor entry raises house prices, and their activity contributes to price fluctuations (Chinco and Mayer, 2016; Gao, Sockin, and Xiong, 2020; Bayer, Mangum, and Roberts, 2021; DeFusco, Nathanson, and Zwick, 2022) with potential adverse macroeconomic implications.[2] Another view emphasizes that housing units owned by investors are essential for the rental markets, and contribute to social mobility and housing affordability. Also, investors who trade frequently serve as middlemen and reduce frictions in housing markets (Bayer, Geissler, Mangum, and Roberts, 2020). Despite the growing interest and the ongoing policy debate, there exists little systematic evidence on how real-estate investors impact the housing ownership and rental markets simultaneously.

This paper studies the effects of a temporary capital-gains tax exemption, designed to induce investors to sell their housing units, on housing and rental markets. [3] Using rich ad-

[1]For instance, Australia, Canada, China, the Netherlands, New Zealand, Singapore, the UK, the US, and Israel witnessed an increase in investors' activity. In Canada about 40% of recently completed homes were purchased by investors. https://betterdwelling.com/canadian-cities-have-seen-up-to-90-of-new-real-estate-supply-scooped-by-investors/. In some US cities, investors were responsible for about a third of home sales in Q4/2021 (https://edition.cnn.com/2022/02/23/investing/premarket-stocks-trading/index.html). In Amsterdam, about a third of houses are held by private investors (https://www.amsterdam.nl/en/news/putting-stop-investors-buying-up/).

[2]The Bank of England stated that: "the scale and nature of Buy-To-Let activity makes it a significant potential amplifier of housing and credit cycles". https://www.bankofengland.co.uk/statement/fpc/2014/financial-policy-committee-statement-september-2014. Similar concerns were raised in New Zealand (Reserve Bank of New Zealand, 2016), Australia (Reserve Bank of Australia, 2017) and the Netherlands (De Nederlandsche Bank, 2018).

[3]Other countries use taxes that target real-estate investors. In the UK, starting in April 2016, second-home purchases are charged with an additional 3% tax (https://www.stampdutycalculator.org.uk/stamp-duty-buy-to-let.htm). The UK Treasury stated that the surcharge's goal is to support homeownership (https://www.gov.uk/government/consultations/consultation-on-higher-rates-of-stamp-duty-land-tax-sdlt-on-purchases-of-additional-residential-properties/higher-rates-of-stamp-duty-land-tax-sdlt-on-purchases-of-additional-residential-properties). Singapore introduced in 2022 a 17% tax on buyers of a second unit and 25% tax on buyers of a third. https://www.thestar.com/news/gta/2022/01/17/singapore-taxes-investment-homes-up-to-30-per-cent-could-a-similar-tax-cool-off-prices-here. South Korea also taxes investors since 2020. https://www.ibanet.org/article/BA51A729-4D36-4FD4-98D3-8F7CB0536ADB. In Israel, investors are subject to a special 8% purchase tax (as of November 2021).

ministrative data and a research discontinuity design, we show that the tax changes increased the sales rate by treated investors by 50%. Investors predominantly sold small housing units, that were typically offered for rent, and then purchased by first-time homeowners. We then exploit variation in investors' presence across local markets to examine how the induced supply shock from the policy change affected local house prices and rents. We find that sales by investors led to a decrease in the price of housing units and to an increase in rents. Our preferred estimate is that an increase in investors' sales by 1 percentage point (pp) of the housing stock, in a given half-year period, causes house prices to fall by 14% and rents on new leases to increase by 10%.

Theoretically, the effect of sales by investors to first-time homeowners on prices and rents is ambiguous. In many frictionless models, the ownership structure of the housing stock and the transaction volume do not affect prices. Each renting household that becomes a homeowner reduces the demand for rental housing and reduces the supply of rental housing by the exact amount that keeps prices unchanged. Likewise, the change in ownership increases demand for owned housing in the exact amount to offset the increased supply due to investors' sales. Some theories of the housing market capture the effects of transactions on prices. These theories emphasize search frictions and transaction costs (e.g., Han and Strange, 2015). Other theories emphasize household heterogeneity to link the ownership structure to prices. For instance, some households may prefer renting due to a greater need for mobility (Halket and Vasudev, 2014; Halket and di Custoza, 2015). Information asymmetry may also prevent high-quality tenants from signaling their type to landlords (Henderson and Ioannides, 1983). Under these conditions, owned and rented housing are imperfect substitutes, and both have demand that is downward sloping with respect to their share of the housing stock. As a result, when more investors sell rental units to first-time homeowners, house prices decline and rents increase. These channels may be amplified when housing markets are segmented by location and quality (Piazzesi, Schneider, and Stroebel, 2020).

We propose a model to capture and interpret two key elements from our empirical findings. First, history-dependent taxes capture realistic features of the tax system and generate significant lock-in effects in line with the data. Second, indivisibility of housing units and heterogeneity in households' ability to extract utility from owned and rented housing generate downward-sloping demand curves for owned and rented housing, which have different elasticities. The model combines the two elements into a single framework, and so is a first step toward analysis of the welfare implications of a policy that treats investors differently.

Our empirical analysis exploits a series of capital-gains tax changes that encouraged some investors, but not all, to sell their housing units at a significantly reduced capital-gains tax rate. Up until January 2011, investors who had sold a housing unit during the previous 4 years paid a capital-gains tax of 25%. Single-unit owners and other investors (i.e., those who have not

sold a housing unit in the previous 4 years) were exempt from capital-gains tax. Starting in January 2011, the government also exempted investors who had sold a housing unit in the past 1.5-4 years. Importantly, the eligibility for the tax exemption was determined based on whether the investor had sold another housing unit in the past 1.5-4 years anywhere in the country. The tax exemption ended in June 2013, and in January 2014, the government enacted a new capital-gains tax law that ties tax payments to the holding period of the unit sold. This sequence of tax changes allows to examine how investors responded to large decreases *and* increases in the tax rate over a short period. Since the tax changes applied only to a subset of investors, we can use other investors as a plausible control group. Moreover, as eligibility was neither based on the location of the unit nor its holding period, we can treat the decision to sell a unit independently from local demand factors and use spatial variation in the number of units held by real-estate investors in different local markets to estimate the impact of investors' sales on house prices and rents.

We base our analysis on rich administrative data on investors and housing markets from Israel. We use several new data sets. First, we use annual property tax records on the universe of housing units in the 76 largest municipalities across Israel. These municipalities account for nearly 90% of the total housing stock in Israel. Second, we use administrative records on the universe of housing transactions in Israel between 1990 and 2018. These data include information on the price, date, buyer and seller IDs, investor status of the buyer, indicators for sellers who pay capital gains tax, and unit characteristics. We match these records to the population registry, which includes the approximate location of the primary residence of individuals, and rich demographic information, such as age, marital status, and links to parents and spouses for all individuals who were either sellers or buyers of a housing unit in Israel between 2002 and 2018. We use this information to determine the identity of the household that owns each housing unit, track the number of units owned by this household on each date, and identify which of the units is the primary residence. Lastly, we use data on rents from a large longitudinal rent survey conducted by Israel's Central Bureau of Statistics for the official consumer price index. The survey includes information on the rental unit location and characteristics and distinguishes between new leases, which were signed in the last 12 months, and extended leases.

Our findings come from two main empirical exercises. In the first exercise, we measure the effects of the capital-gains tax changes on sales by investors that were subject to the capital-gains tax exemption, which we define as treated investors. We document a discontinuity in the probability of sale by treated investors when they were subject to the 25% capital-gains tax, i.e. just before and just after the 4-year mark. During the temporary exemption phase, this discontinuity around the 4-years cutoff disappears. We quantify the effect of the tax exemption on sales using a controlled difference-in-differences design in which we compare

the sales probability of housing units owned by treated investors to those owned by control investors (investors who sold another unit in the past 1.5, or more than 4 years). We find that the reduction of 25% in the capital-gains tax rate increased the sales probability of units owned by treated investors by 50%. In absolute terms, the semiannual sales rate by the treated investors went from 1.2 pp before the exemption to 1.8 pp after. In auxiliary analyses, we further show that most of these sales were of non-primary housing units (i.e., they are not the primary residence of the owner and thus more likely to be a rental unit) and were sold to non-investor buyers, such as first-time homeowners. When the temporary exemption expired, the sales probability reverted to values almost identical to its values in the pre-exemption period, strengthening the causal interpretation.

In the second exercise, we explore how the additional sales (or lack thereof) induced by the tax changes affect housing prices and rents. We document considerable spatial variation in the composition of investors across 360 local housing markets. We use this variation to construct a predictor of the housing units sold by investors, as a share of the local housing stock, for each half-year period and local market. Then, we regress transaction prices on the local sales predictor as a share of the local housing stock. We find that an increase in sales by investors by 1 pp reduces house prices by 14%. The coefficient is larger in magnitude when we restrict attention to sales of smaller units in which investors specialize. This points to imperfect substitution between small and large housing units. We also conduct the same analysis with rental data and find that a 1 pp increase in investor sales increases rent on new leases by 10%. Similar to transaction prices, the coefficient is larger for smaller units, which is the segment in which investors own a relatively larger share of the stock. We find an insignificant and negligible impact when we restrict attention to extended leases. This is consistent with rents on existing leases being driven by inertia and partly insulated from market conditions.

A back-of-the-envelope calculation based on these numbers and given the small share of units owned by treated investors (roughly 6%) suggests that the tax exemption reduced house prices by 0.4% and increased rents by 0.28%. We also consider a counterfactual scenario in which the capital gains-tax is increased for investors who did not sell another unit in the last 4 years from 0% to 25%. This is a relevant counterfactual since this is what the January 2014 tax change does, though gradually over a long time period. Under this scenario, which affects a much larger share of the housing stock, house prices would increase by 2.5% and rents decrease by 1.8%. Importantly, this calculation only accounts for the lock-in effect of capital-gains tax and does not adjust for changes in the demand for real-estate investment due to tax changes, and thus should only be used to gauge the estimated effects that we see in the data.

In the final step of the analysis, we use an instrumental variable approach, where the dependent variable in the first stage is the share of units held by investors in the local market, and the dependent variables in the second stage are either prices or rents. This specification

assumes that the change in ownership structure is driving the change in house prices and rents. Thus, each additional sale induced by the temporary tax exemption shifted one housing unit from investors to non-investors. We find that the estimates of the inverse-demand semi-elasticity of local house prices and rents with respect to investors' share are 2.7 and -1.0, respectively. These estimates point to a large potential effect of a small number of transactions by speculators on house prices (Piazzesi and Schneider, 2009; Bayer et al., 2020) and suggest that the local demand for housing is quite inelastic in the short term, in line with theories that emphasize segmentation of housing markets and household heterogeneity (Piazzesi, Schneider, and Stroebel, 2020).

The causal interpretation of the link between investors' activity and prices hinges on the validity of the sales predictor as an exogenous supply shifter. A key assumption is that the measure of predicted sales by investors in a local market which we derive from the tax changes and the pre-policy local composition of investors, is independent of local demand shocks. Similar to shift-share instrumental variable design, the predictor of sales can be considered as an exogenous supply shifter either because investor type shares are independent of local demand shocks, or because the aggregate changes in sales probability (the "aggregate shocks") are independent of them (Borusyak, Hull, and Jaravel, 2022). We believe that in our setting both conditions are plausible. First, the changes in the tax code were determined at the national level and were largely unanticipated. Second, the tax treatment of different investor types depended on whether they sold another unit in the 1.5-4 years time window. Thus, neither the holding period nor the location of a housing unit determined the eligibility of an investor for the tax exemption we consider. To further alleviate concerns that the composition of investors in the local market is nevertheless correlated with unobserved characteristics that predict house price appreciation, we show that the local investor composition in late 2010 (before the tax exemption came into effect) is uncorrelated with previous price appreciations at the local level. In contrast, when we repeat the same analysis for 2011, after the tax exemption was granted, we find a negative relationship between price increases and the share of housing units owned by investors in the local market.

To the best of our knowledge, we are the first to document a causal link between investors and the rental market. Our analysis shows how a policy that targeted investors, changed the ownership structure in the housing market, and affected housing prices and rents. This result reveals an important and underexplored policy trade-off: encouraging investor sales may reduce prices and benefit those who intend to buy, but at the same time may hurt renters who may not be able to buy due to credit constraints or search frictions. While our analysis considers the effects of a particular policy change, we believe that the implications extend to a wider set of policies, such as banning foreign real-estate investors or permanently raising transaction taxes on investors. In many countries, the recent rise in investors' activity is treated

with suspicion, especially due to a view that homeownership is a stepping stone for young households to build wealth. Yet, based on our results, policies that promote homeownership by reducing investors' ownership share may also unintentionally raise rents and hurt relatively poor households. Our paper also adds to the literature that examines the impact of tax policy on housing markets. In particular, very few studies (Gao, Sockin, and Xiong (2020); Shan (2011)) consider the impact of capital gains taxes on investors' and market outcomes.

Related literature. Our paper is related to few strands of the real-estate literature. First, the literature examines the impact of various taxes, typically transaction taxes, on market outcomes such as volume, price and time on market. Best and Kleven (2018) use notches in the UK transaction-tax code and a temporary "tax holiday" in 2008-2009 to estimate the effect of taxes on transactions. Slemrod, Weber, and Shan (2017); Kopczuk and Munroe (2015) study house price responses using bunching at U.S. transaction tax notches. Few studies consider the impact of capital-gains taxes on the real-estate market. Gao, Sockin, and Xiong (2020) use cross-sectional state-level variation in capital-gains taxes to provide evidence that purchases by investors contributed to rising house prices and to the severity of the economic downturn which followed. Shan (2011) explores the impact of a change in the US capital-gains tax law in 1997. She uses a new rule which exempted from capital-gains tax sales of units with less than \$500K of capital gains to show that the exemption increased semiannual sales rate of units with sub-\$500K capital gains by 0.4-0.62 percentage point. For the most part, these papers did not consider the impact of such taxes on the rental market. Two recent papers that consider the linkage between tax policy and the rental market are: Han, Ngai, and Sheedy (2021) and Levy (2021). Han et al. (2021) estimate the effects of Toronto's imposition of property tax on both rental and ownership markets. Levy (2021) use French data and show that real-estate investors strongly prefer to own properties they rent close to their primary residence.[4]

Second, a nascent strand of the literature focuses on the impact of real-estate investors. For instance, Haughwout, Lee, Tracy, and Van der Klaauw (2011) and Albanesi, DeGiorgi, and Nosal (2022) document that investors were mostly active in US states that experienced the largest housing booms before the crisis. Bayer, Geissler, Mangum, and Roberts (2020) document a sharp rise in speculator investor activity in the Los Angeles metropolitan area in the years leading to the financial crisis. Focusing on China, Somerville, Wang, and Yang (2020) show that restrictions on the number of residential properties that each individual can buy had a significant cooling effect on newly constructed units, sales volume and prices. These studies typically do not distinguish between buy-to-let buyers and buyers who purchase housing

[4]Papers that study the impact of policies that target the rental market, without direct reference to investors include: Diamond, McQuade, and Qian (2019) and Gete and Reher (2018). Diamond et al. (2019) examine the impacts of rent control on tenants and landlords in San Francisco, and Gete and Reher (2018) show that an increase in the rate of mortgage denials in the US contributed to a rise in rents relative to the counterfactual.

units as a secondary residence. Studies that made this distinction, typically consider buyers that use these units as vacation homes or a second residence (Badarinza and Ramadorai, 2018; Cvijanović and Spaenjers, 2021; Favilukis and Van Nieuwerburgh, 2021; I. García, 2022). In contrast, in our study investors reside in their primary unit and rent rent out their non-primary unit. Finally, the real-estate literature examines the importance of homeownership (e.g., DiPasquale and Glaeser, 1999; Malmendier and Steiny) and the trade off that households face between renting vs. owning a house (e.g., Sinai and Souleles, 2005). Other papers investigate the impact of various taxes on the buy-vs-own decision. (Sommer and Sullivan, 2018; Floetotto, Kirker, and Stroebel, 2016) take a different approach and write a dynamic model of the US housing market while considering the impact of taxes on housing prices, homeownership and mortgage debt. These papers did not examine the role of households who own multiple housing units (i.e., investors).

The rest of the paper is organized as follows. In Section 2 we provide relevant background on the Israeli real-estate market and the tax policy. We also present the data and descriptive statistics. Section 3 examines the effect of the capital gains tax reform on investors' sales, and Section 4 explores the effect of investors' sales on house prices and rent. Section **??** builds a model that captures the key empirical findings. Section 6 concludes.

2 Background and Data

2.1 The Israeli housing markets

During the last decade and a half, house prices and rents in Israel steeply increased. Between 2009 and 2014, the period that we focus on, rents rose by 30% and house prices increased even more so that the price-to-rent ratio rose by an additional 33%. Figure 1 presents these patterns, based on a hedonic regression. Rising prices were accompanied by significant changes in the ownership structure of the housing stock. According to a household survey conducted by the Central Bureau of Statistics (CBS), the share of Israeli households who were real estate investors more than doubled–from 4.2% in 2009 to 9.1% in 2014. Figure A.3 in the appendix presents how the share of investors' share out of housing transactions in Israel changed between 2002 and 2018. The figure shows that investors were involved in a substantial share of real-estate transactions, both as sellers and as buyers. Until mid 2015, the net purchase share of investors (share of unit purchased by investors minus the share of units sold by investors) was positive, reaching more than 10 percent of transactions towards the end of 2010, shortly before the tax exemption that we study. The share of households living in a non-owned residence was 32.7% in 2014, which is similar to the OECD average.[5]

[5]See, Raz Dror and Shamir, 2017: `https://economy.pmo.gov.il/councilactivity/housing/documents/rent.pdf`. This share of renters is largely similar to the share of renters in the EU (31%), the US (36%),

The rise in house prices generated considerable public concerns about housing affordability, and the government enacted a series of initiatives aimed at slowing house prices, making housing more affordable to young households, and curtailing investor activities. In particular, in late 2010, the Israeli government decided to change the capital-gains tax levied on investors in an attempt to encourage investors to sell their existing units.[6] Up until 2011, investors who sold a housing unit paid a flat 25% capital-gains tax. But investors who did not sell any other unit in the preceding four years were exempt from this tax. Between January 2011 and June 2013, this exemption was extended to all investors who did not sell any other unit in the preceding 1.5 years. Between July 2013 and December 2013, the new exemption for investors who sold another unit in the past 1.5-4 years expired. Finally, in January 2014, a new tax code came into effect which determines the capital-gains tax rate for investors based solely on the holding period of the unit sold (i.e., no longer distinguished sales by the history of owners' transactions). The tax changes are illustrated in Figure 2.

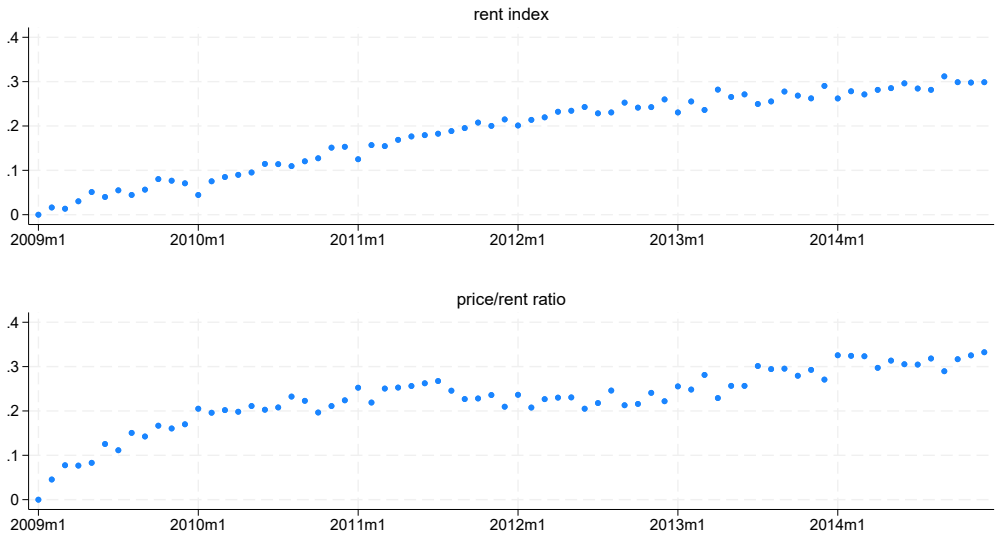

Figure 1: Rent index and Price-to-Rent ratio in Israel (2009-2015)

Notes. The figure displays rent index and price-to-rent ratio based on a hedonic regression of log prices on number of rooms, year-month and location (statistical area) fixed-effects as well as an interaction between transaction type (house purchase or rental lease) and year-month. The value of the index is normalized to 0 in January 2009. *Sources*. Data on house prices are from the Israel Tax Authority. Data on rent are from a rental survey conducted by the Central Bureau of Statistics.

Five features of the capital-gains tax changes make them particularly attractive for our

and the UK (37%). See Bracke 2015 for patterns in the UK.

[6]This was not the only measure adopted by the government. For instance, Tzur-Ilan (2019) and Laufer and Tzur-Ilan (2021) study the impact of loan-to-value limits on mortgages introduced in 2010 and readjusted in 2012. Our identification strategy controls for these and other changes that affect all investors.

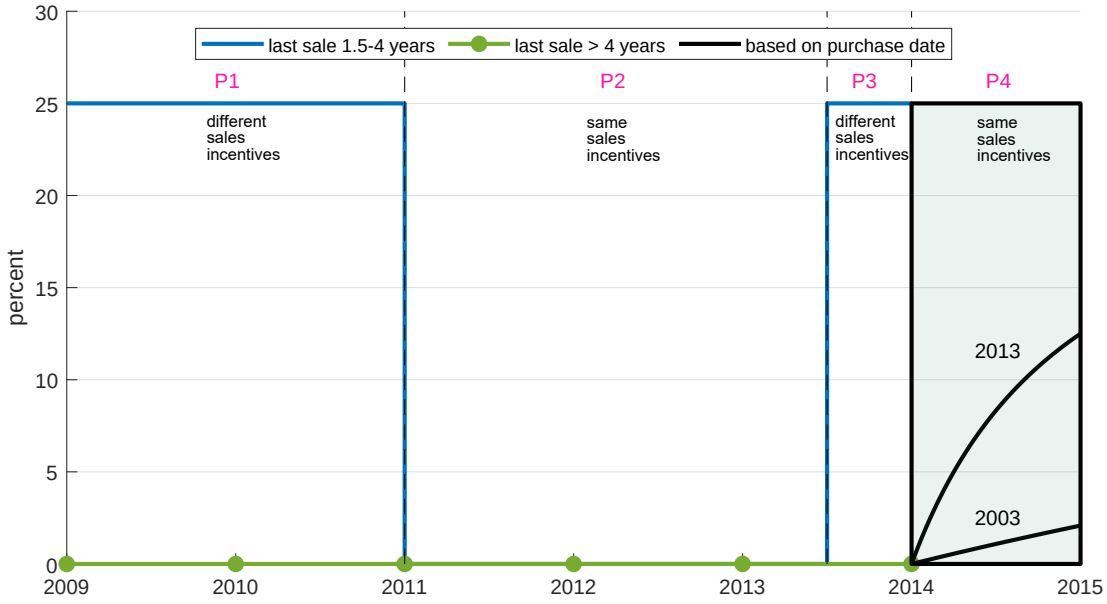

Figure 2: Changes in capital-gains taxes in Israel

Notes. The lines show the capital-gains tax rate for residential properties owned by investors between 2009 and 2014. Until 2011 (P1), a 25% tax rate applied only to investors who sold another property in the 4 years before the current sale (blue line). Investors who did not sell in the previous 4 years were exempt (green line). Between 2011 and 2013H1 (P2), investors who sold another property 1.5-4 years before the current sale were also exempted. In 2013H2 (P3) the temporary exemption expired. Starting from January 2014 (P4), the tax rate is determined by the holding period of the property. For instance, when selling a property purchased in 2003, capital-gains up to 2014 are exempted, but a 25% tax rate applies for capital gains from 2014 onward.

empirical analysis. First, the temporary exemption from January 2011 to June 2013 applied only to a subset of investors (investors who had sold another housing unit in the previous 1.5 to 4 years), whom we consider treated investors. Investors whose previous sale of a housing unit occurred 4 years or more before the current sale did not pay capital-gains tax both before and after the January 2011 tax changes and are included in the control group. This differential impact on a subset of investors allows us to compare the response of investors affected by the policy to other investors. Second, eligibility for the capital-gains tax exemption was based on the sale of another unit. Therefore, the allocation into treatment and control groups does not depend on the characteristics or location of the units themselves. For instance, an investor who considered selling a housing unit in Tel Aviv in 2012 was subject to the exemption if she sold another unit, say in 2009, anywhere in Israel. By contrast, an investor who considered selling a housing unit in Tel Aviv in 2012, and who had sold another housing unit in 2007 was not affected by the change in the tax law. This is because that investor was already exempt from the capital-gains tax before the law change. Third, the initial phase of the tax change was

unexpected. It was announced in late 2010 and came into effect in January 2011. Moreover, the change was scheduled to take effect for exactly two years (until December 2012).[7] Due to early elections that were announced in October 2012 and took place in January 2013, the exemption stayed in effect for additional six months. Since it was difficult to predict the tax changes, strategic behavior by investors in anticipation of tax changes is unlikely to have played a major role in their sale decisions. Fourth, the tax rates changed several times for the same treatment group. This provides more than one differential treatment points, and gives us the opportunity to consider what happens when the tax rate is both reduced and increased. Moreover, during the exemption period (P2) and the new tax regime (P4) the treatment group and the control group faced the same tax rate, whereas in the pre- and post-exemption periods (P1) and (P3) the tax rates were different. Lastly, the treatment was large: the capital-gains tax rate was reduced from 25% to 0% and then increased back to 25% before switching back to zero in early 2014. Since house prices were rising fast at the time, the exemption removed a large disincentive to sell. Therefore, we can expect considerable changes in the sales' behavior of the treated. To illustrate the effect, Figure 3 shows the number of sales by each of the three following categories of sellers: non-investors, treated investors and control investors. Investors own 2 housing units or more, and the treated investors are those who sold another housing unit 1.5-4 years prior to the current sale. The figure clearly shows an increase in sales volume by treated investors from 2011 to mid 2013 and from 2014 onwards.

Renting households are typically smaller, younger, and earn a lower income than homeowners. Investor households typically own 2 to 3 units each, including their primary residence, and supply 90% of rental housing units.[8]. Rental units are typically smaller and older than owner-occupied units. The typical rent contract is for 12 months (90% of contracts), and the mean lease duration is 4.5 years (Ater, Elster, Genesove, and Hoffmann, 2023). Finally, at the relevant time period, institutional investors had a negligible role in the Israeli residential real-estate market.

2.2 Data and descriptive statistics

Our analysis is based on several administrative data sources. The main data sources cover Israel's housing stock (2011-2018) and all housing unit transactions (1990-2018). For buyers and sellers of residential real estate units, we have information on demographic characteristics from the population registry, and for the rental market analysis, we use a large rental survey. Below, we describe each data source.

[7] See a news article from December 2010 (in Hebrew) reporting on the tax change: `https://www.themarker.com/realestate/1.562304`

[8] The rest are owned by firms and public sector entities. See for example Hausman, Ramot-Nyska, and Zussman (2022)

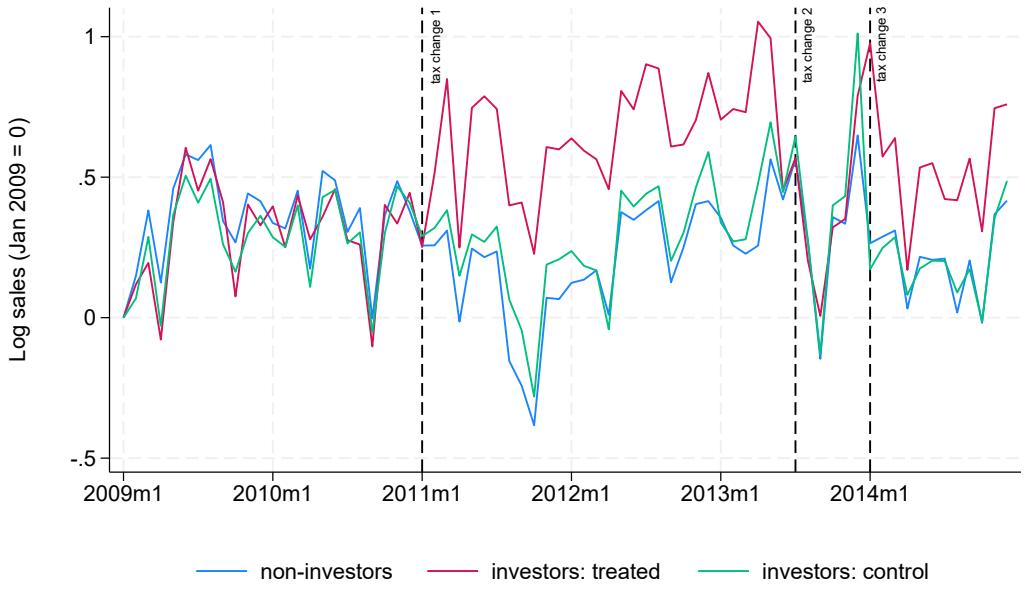

Figure 3: Changes in volume of sales by sellers' type in Israel, 2009-2014

Notes. The lines show the number of sales (in logs, normalized to 0 in January 2009) of three seller types: non-investors, treated investors and control investors. Investors own 2 housing units or more (before selling the current housing unit). Treated investors are investors who sold another housing unit 1.5-4 years prior to the current sale. *Sources*. Data on sales are from the Israel Tax Authority. Data on ownership are based on both Israel Tax Authority and municipalities' property tax files.

2.2.1 Housing Stock

We use data on the housing stock from the property tax records covering the largest 76 municipalities in Israel. Starting in 2011, municipalities are required to report to the Central Bureau of Statistics annual information on all housing units in their jurisdiction. The records include identifiers of the owner of each unit and the unit's location at the statistical-area level (equivalent to a census tract). In a statistical area (4-digit code) there are between 3,000 and 5,000 residents, and an average of 1,000 housing units. Overall, the data covers 1.87 million housing units as of 2011, and it grows by about 30,000 every year to reach 2.13 million housing units by mid-2019.[9]

2.2.2 Housing purchase transactions

The transaction data include 2.8 million residential real-estate ownership transactions carried out in Israel between 1990 and 2018, based on a registry by the Israel Tax authority.[10] Each transaction includes identifiers of the buyer and seller, transaction price, date, location (at

[9]Growth in the number of housing units is likely due to improved coverage rather than construction of new units.

[10]We later only keep the 2.5 million transactions where the part of property that is being transacted is reported

the statistical-area level), investor status of the buyer, indicator for sellers who pay capital gains tax, and property characteristics: number of rooms, area, and building age. For the 2.5 million buyers and sellers that appear in the data, we also have annual demographic information based on Israel's population registry for the years between 2002 and 2018. This information includes the location of the primary residence, identifier of spouse, and other individual characteristics, including marital status, gender and year of birth.

2.2.3 Building households' transaction history

We use the demographic information on buyers and sellers to construct households, as the Israeli tax code on house transactions is based on households' overall ownership of housing units. Starting in 2002, we use the buyer or seller ID, their gender and their spouse ID to comprise each household ID (HHID) of male and female IDs for married couples, or a single ID in case the seller or buyer is not married that year. For years earlier than 2002, we do not have demographic information, so we assume for simplicity that the buyer or seller had the same HHID between 1990-2002. We added to each household and year information about last year's HHID. We kept unique HHIDs and year observations- which led us to 1.7 million observations (note that we had 2.5 million observations of buyers and sellers, but this number includes the same household appearing more than once in the same year).

We then merge this information into the housing stock data, adding to each owner their HHID. This allows us to calculate for each year, the number of housing units owned by each household. We only use mid-2015 data, and find that 85% of households hold 1 housing unit, 11% have 2 housing units, and only 4% hold more than 2 housing units.

Our next aim is to create a measure of the housing stock held by each household based on their transaction history. We thus look at each transaction as two entries: one that subtracts a unit from the seller's household, and one that adds a unit to the buyer's household. This should lead us to 5 million observations, but we only have demographic data on 1.67 million sellers and 2.45 million buyers, which leaves us with 4.12 million observations (this is mainly due to non-private sellers, i.e. real estate contractors who sell new housing units).

Between 1990 and 2002, since households do not vary, we simply add up the number of housing units bought and subtract the number of units sold. In this manner we know for each transaction the number of housing units held by the household historically, and the date of their previous transactions. Later, we use these dates to calculate the time span between current and previous sales, which is crucial for identifying treated investors in our setting. We also use the investor status of buyers as a proxy for seller investor; a seller in the current sale is considered an investor if he is labeled as an investor buyer in his previous transaction.

Starting in 2003, HHID may vary over the years. If it is the same HHID as last year, we use the procedure described in above. If the household splits (divorces for example), we

divide the number of housing units of that household by 2, and keep the same last transaction dates. If the household unites (marriage) we add their housing units and keep the most recent transaction date.

Finally, we fix the number of housing units held by each household using two steps. First, we take a snapshot of each household stock of housing units in mid-2015, and compare it to the data we have from municipalities' property tax records. We add the corresponding housing units to the household in case the latter is larger than the former.[11] Second, if we end up with a negative number of housing units held by a household, we add housing units to that household so it will not be smaller than zero.[12]

2.2.4 Sellers' investor status

A seller household is considered an investor in a specific sale if he either has housing units after he sells or if he was labeled an investor buyer in his previous transaction. Focusing on the 2.4 million observations of purchases between 1990-2018, 51% of sellers are not investors, 36% are investors, and 13% are real estate contractors who sell new housing units. When we narrow down the analysis to 2009-2014 (see table 1), we are left with 374,603 transactions. Here 46% of sellers are not investors, 35% are investors, and 19% are real estate contractors who sell new housing units.

Table 1 provides summary statistics of housing units sold by different types of sellers over the sample period. Overall, we observe 374 thousand transactions (Column 1). 81% are sales of existing housing units (Column 2) and the rest are new units. The convention in Israel is to measure the number of rooms including a living room (a 3-rooms unit is equivalent to a 2 bedroom unit in the US). The mean number of rooms (including a living room) is 3.7 and the floor area is 85.7 squared meters (920 squared feet). The mean age of the structure is 25.6 years and the mean price is just over 1 million Shekels (approximately 300 thousand US dollars). About a third of the sales are by investors. We further split sales by investors into sales of their primary unit of residence (Column 3) and any other units they own (Column 4). Non-primary units (which are likely rental units) are smaller, older, and cheaper than the primary units. Columns 5 and 6 present characteristics of transacted housing units separately for sales by treated investors (investors who sold another property 1.5-4 years earlier) and by control investors (all other investors). While the differences in the number of rooms and prices are insignificant, the differences in size and age between the control and treatment sales are statistically significant but economically small.

[11] The idea here is to account for heritage or other non-transactional manners of accumulation of housing units. We do not subtract housing units, because the data on municipalities is limited in scope, while the transactions data covers all of Israel

[12] For instance, if a household sold 3 units but has no record of buying housing units, we add these units to their total units owned. This may happen due to housing purchases before 1990

The table also reports the share of buyers that own a single housing unit after the transaction, the share of upgrades who own another unit but declared to the tax authority that they will sell it within two years, and the share of buyers who are investors (i.e. own at least one other housing unit). While investors buy 25% of housing units transacted overall (Column 1), their share of non-primary (rental) units is 32% (Column 4) vs only 20% of primary units (Column 3). The last row of table 1 reports the average number of units owned by sellers before they sell. Sellers on average own 1.9 units before they sell. Investors who sell non-primary units own 2.59 housing units vs only 1.9 housing units owned by investors who sell their primary unit. Treated investors own 2.99 housing units vs. control investors who own 2.34 units on average before they sell.

Table 1: Mean characteristics of housing units that were sold, 2009-2014

	all sellers		investors			
	all	resale only	primary	non-primary	treated investors	control investors
	(1)	(2)	(3)	(4)	(5)	(6)
Number of rooms	3.7	3.6	3.9	3.3	3.4	3.5
Area (m^2)	85.7	79.7	88.6	73.4	77.1	78.2
Building age (years)	25.6	31.7	29.5	35.6	35.6	33.6
Price (mil. ILS)	1.11	1.01	1.22	0.96	1.03	1.04
buyer statuss:						
single homeowner	0.67	0.68	0.71	0.63	0.46	0.66
upgrader	0.08	0.07	0.09	0.06	0.07	0.07
investor	0.24	0.25	0.20	0.32	0.47	0.27
units owned by seller	–	1.93	1.90	2.59	2.99	2.34
N	374,603	301,263	41,172	89,732	8,181	122,729

Notes. The table reports mean characteristics of housing units sold between 2009 and 2014, by seller and unit types. Column (1) includes the full sample, and column (2) focuses on housing units sold by a previous owner (as opposed to new units sold by contractors). Columns (3) and (4) compare primary and non-primary housing units sold by investors. Investors reside in larger, newer and more expensive housing units compared to housing units that investors rent out. Columns (5) and (6) compare the units sold by treated and control investors. These units have similar characteristics. The average ILS-USD exchange rate during this period was 3.7. See text for more details. *Sources*. Data on housing purchase transactions are from the Israeli Tax Authority.

Figure 4 illustrates the impact of the tax exemption that we study on investor sales. The blue round dots represent the sales rates of investor-owned units in 2009 by the time since last sale. In 2009, the mean semiannual sales rate by investors who sold another unit 1.5-4 years before the current sale was 1.2 pp. The mean semiannual sales rate by investors who

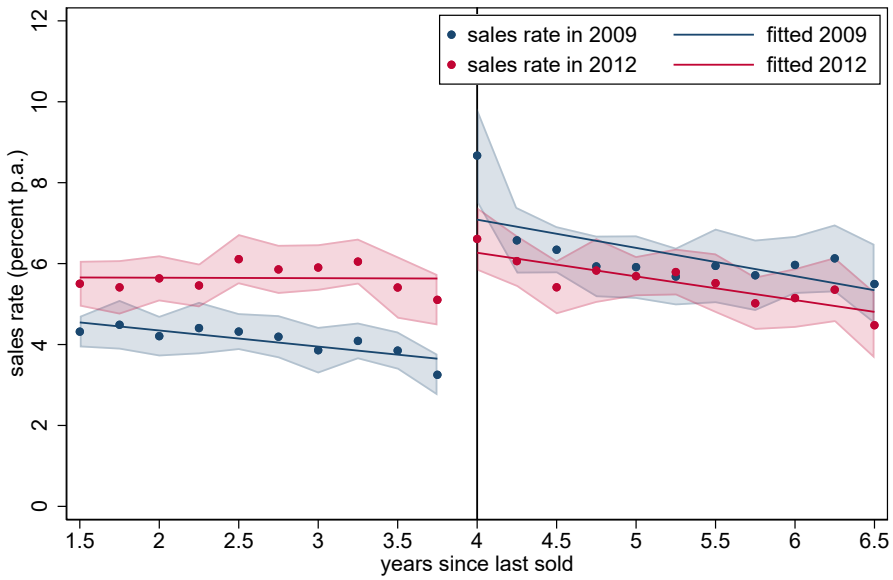

Figure 4: Investors' sales rate by time since the last sale

Notes. The figure displays the sales rate of investors relative to the time since they sold another unit, focusing on the 4-year cutoff. Dots represent the semiannual sales rate of housing units owned by investors defined by the time since their last sale. Values for 2009 are marked with blue dots and lines and values for 2012 are marked with red dots and lines. Lines are linearly fitted separately to investors who sold another unit in the last 4 years and those who have not. In 2009 there is a discontinuity in sales rate at the 4 years since the last sale. In 2012, when the tax exemption also applies to investors who sold another unit during the previous 4 years, this discontinuity disappears. Moreover, sales of units by investors who have not sold another unit in the last 4 years are virtually unchanged. See text for more details. *Sources*. Data on the housing stock are from 76 municipalities' property tax files. *Source*: Data on sales are from the Israeli Tax Authority.

sold another unit between 4 and 5 years earlier was over 2 pp. The blue lines show the linearly fitted values, which reveal a discontinuity in sales rates around the 4-year cutoff. This discontinuity reflects a lock-in effect for sellers who are required to pay the capital-gains tax based on the history of their sales. The red triangular dots in Figure 4 represent the same statistics in 2012, the year in which investors that were above and below the 4-year cutoff faced the same tax treatment. In contrast to what we see in 2009, the discontinuity around the 4 years cutoff disappears during the exemption period. Furthermore, it is investors who had sold another unit in previous 1.5-4 years period (to the left of the cutoff) who substantially changed their sales' rate.[13] A surprising feature of Figure 4 is that there doesn't appear to be a large spike in sales right after the tax cutoff. The absence of such a spike may suggest that the additional sales by the investors on the left of the cutoff during the exemption period did

[13] Appendix Figure A.5 displays the same calculation conducted in every half-year between 2009 and 2014. The discontinuous jump at 4 years is there in 2009H1-2010H2 and 2013H2, and is absent in all other periods confirming the result. Rows of figure A.5 represent different tax periods, P1-P4, as defined above.

not come at the expense of future sales.

To better understand the tax liability of treated investors, we examine whether they actually paid capital-gains taxes or not. Figure 5 shows the share of treated investors who paid capital-gains tax in each period between 2009 and 2014. Before the exemption period (during 2009 and 2010) about 30% of "treated" sellers paid the tax. This share drops abruptly in 2011 and continues to be very low until the second half of 2013.[14] The share of treated sellers who pay the tax returns to around 30% in the second half of 2013, when the temporary exemption expires. During 2014, the share who pay capital-gains tax doubles (reaching 60% of sales by treated investors) as the new tax code determines the tax rate based on the holding period of the specific unit and applies to all sales by investors, regardless of their sales' history.[15]

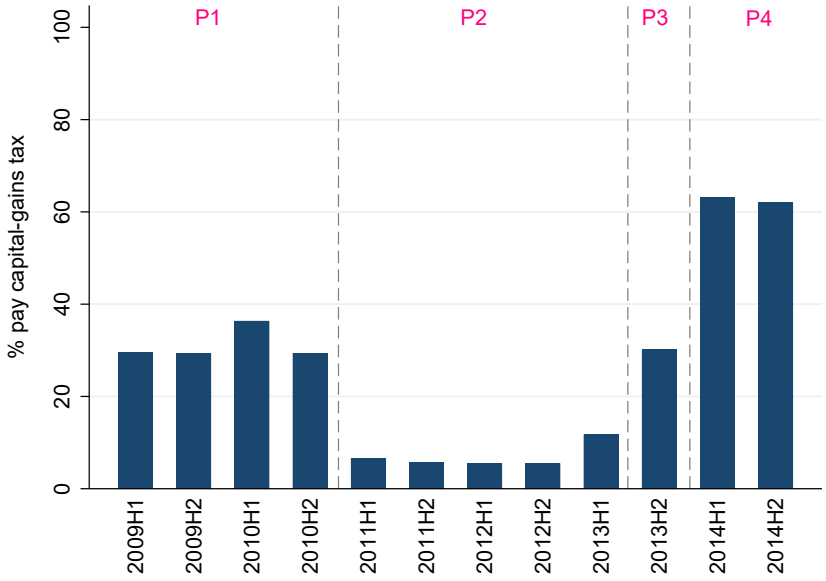

Figure 5: The effect of the tax reform on capital-gains tax payments by treated investors

Notes. The figure displays the share of sales by treated investors who paid a capital-gains tax in every half-year period between 2009 and 2014. The figure shows that when the tax exemption was granted (P2), very few treated investors paid capital-gains taxes. In P1 and P3, treated investors were potentially subject to the tax and many of them pay. In P4, the capital-gains tax was no longer based on elapsed time from the previous sale. See text for details. *Sources*. Data on sales are from the Israeli Tax Authority.

[14]The share of treated investors who pay the tax is not zero during this period, since additional requirements need to hold for the exemption. For instance, the sale price needs to be lower than 2.2 millions Israeli Shekels, and holding period of at least 1.5 years. Both are relatively rare in our data.

[15]Not all treated investors pay the tax since not every seller has a positive capital gain, for example when they can provide evidence for significant improvements in the unit.

2.2.5 Rent survey

Data on rents are based on a monthly rent survey conducted by the Central Bureau of Statistics. This survey covers a representative sample of rental units and is used to construct the consumer price index in Israel. Importantly, truthful response to this survey is mandated by law. Each observation includes information on the rent, month, location and number of rooms in the property. The survey is conducted on a rotating panel, in which the first observation is always a household who recently moved in, and is sampled as long as it did not move out. We therefore consider the first observations of each household to be "new leases" and subsequent observations "extended leases". The survey includes 65 thousand total observations, including 19 thousand new leases between 2009 and 2014. 61% of renters live in a unit with 3 rooms or less (including a living room). Average rent is 3.2 thousand Shekels per month (approximately 920 dollars).

In sum, we construct a panel of the housing units' stock, sales and a representative sample of rental units in 360 Israeli local markets for each half-year between 2009 and 2014. Four patterns in the data are important for our analysis. First, investors in Israel hold a substantial share of the housing stock (36%). Second, 13% out of investors' housing units are held by treated investors, i.e. investors who were affected by the temporary capital gains tax exemption we analyze in this paper. Third, investors more often sell their non-primary housing units, which are smaller than their primary units. This distinction will be useful for our analysis of the effect of investors on housing and rental markets. Fourth, treated investors increased their sales during the tax exemption period. In the next section, we estimate the magnitude of the effect.

3 The Effect of Capital-Gains Tax on Investors' Sales

During the tax exemption period investors who sold another housing unit between 1.5 and 4 years ("treated investors") were exempt from the tax. Figure 4 above illustrates the impact of the tax changes on investor sales around the 4 years cutoff. We now apply a difference-in-difference framework to quantify the effect of the tax rate on sales of units owned by treated investors. We use units owned by non-treated investors as the control group. Arguably, these investors face similar time-varying factors, such as shifts in beliefs and financial conditions, that affect decisions by treated investors to sell. As shown in Table 1, the characteristics of housing units sold by control and treated investors are quite similar, suggesting that the treated and control groups face similar underlying trends.

Figure 6 presents the semiannual sales rate of the treatment and control groups. The sales rate of the treated investors was lower than that of the control group before 2011 (P1), in line with the evidence presented earlier. The mean semiannual sales rate by treated investors

is 1.2 percentage points in this period. Then, shortly after the adoption of the tax reform, the sales rate of treatment units rises all the way to the level of sales by the control group. During 2011-2013H1 (P2) the sales rates of the treatment and control units follow a similar pattern and are almost identical. The mean semiannual sales rate by treated investors is 1.8 percentage points in this period. Following the expiration of the temporary exemption for the treatment group in 2013H2 (P3) the sales rate of treatment units plummets, while the sales rate of control units remains high. Finally, starting from January 2014, the capital gains tax rate is no longer determined by the timing of the previous sale, eliminating all differences in taxes between treatment and control, and the sales rates of the two groups converge at a lower level.

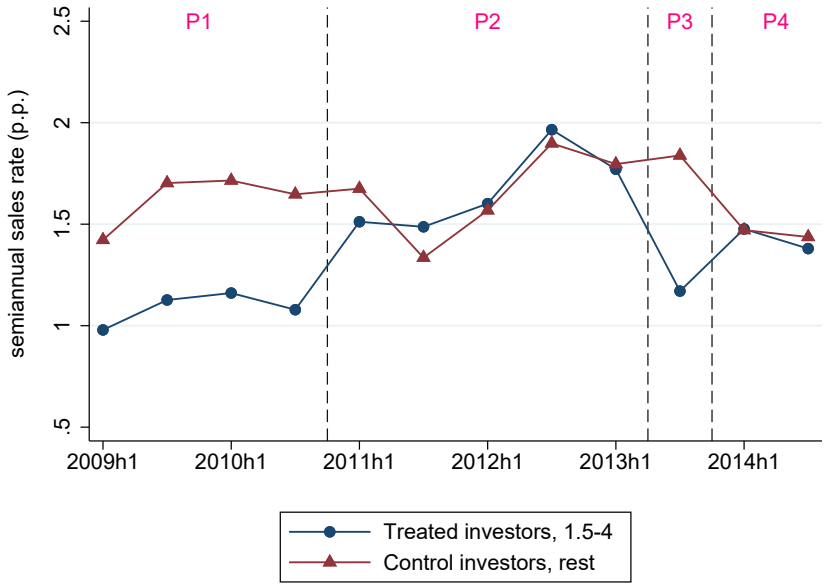

Figure 6: The effect of the tax reform on sales by investors

Notes. The figure displays semiannual sales rates separately for treated investors and for all other investors. The sales rate by treated investors increased to the level of sales of control investors when the capital-gains tax was removed (P2), diverged when the differential tax rate was in place again (P3), and similar when the differential tax rate was eliminated (P4). See text for more details. *Sources*. Data on the housing stock are from 76 municipalities' property tax files. Data on sales are from the Israel Tax Authority.

The statistics presented in Figure 6 do not take into account the composition of treatment and control groups over local markets. For example, demand for housing units in areas in which treated investors own many housing units may have increased during the exemption period (P2). If this increase is unrelated to the tax exemption, we may erroneously attribute it to the tax exemption. To address such concerns we conduct a two-way fixed-effect regression

using the following difference-in-differences specification:

$$\text{sales rate}_{ijt} = \beta_{it}\text{treated investors}_i + \delta_j + \theta_t + \varepsilon_{ijt}, \quad (1)$$

where "sales rate" is the share of sales out of the stock of housing units owned by each investor type (i) in every local market (j) and half year (t), "treated investors" is an indicator which is equal to 1 for investors who sold another property in the last 1.5-4 years, and 0 for all other investors, θ is a half-year fixed-effect, δ is a local market fixed effect and ε is an error term. The equation is estimated at the local area level using the number of units owned by each type as frequency weights, essentially capturing the sales rates at the unit level. Standard errors are clustered at the local market level. Figure 7 plots the value of the β coefficients from Equation (1). Similar to the aggregate trends, before the 2011 tax change (P1) and during the second half of 2013 (P3) treated investors sell less frequently compared to investors in the control group. In contrast, during the exemption period (P2) and the equal treatment period (P4), the two groups have a similar tendency to sell.[16] As in Figure 4, we do not observe apparent spikes in sales by the control group due to delays in sale of units.

We now examine the type of units sold by treated investors and the type of buyers to whom they sell these units. Panel A of Figure 8 focuses on the sales rate of units owned by treated investors before and after the tax reforms, calculated separately for primary and non-primary units. In general, investors more often sell housing units they do not reside in, which are likely to be rental units. Both types of housing units are sold more often during the exemption period between 2011-2013H1 (P2) and from 2014 onward (P4), but most of the added sales consist of non-primary (rental) housing units. Panel B of Figure 8 further considers the identity of buyers of housing units to whom treated investors sell, distinguishing between investors and non-investors. The figure illustrates that housing units sold during the exemption periods (P2 & P4) were mostly purchased by non-investors.[17]

Overall, this analysis shows that the existence or absence of capital-gains tax at a given point in time has a pronounced impact on the tendency of investor to sell. Treated investors, who sold on average 1.2 pp of their housing units each half-year period when they were subject to the tax, increased their sales to 1.8 during the exemption period–a 50% increase– which is also equal to the sales rate of the control investors. Moreover, when the exemption temporarily expired the sales rate dropped back to the pre-exemption values. We also provide evidence that these additional sales affected the ownership structure of housing units as the

[16] Appendix Figure A.4 uses a similar specification but with a narrower definition of the treatment and control groups. In panel A, we focus on investors who sold another housing unit in the past 3-4 years and control investors who sold another housing unit in the past 4-5 years. Panel B further restricts the estimation to 3.5-4 and 4-4.5 years, respectively. The results are qualitatively similar.

[17] Sales are divided by the overall stock of housing units owned by treated investors. Similar results are attained when dividing the sales by the number of transactions in the half-year period.

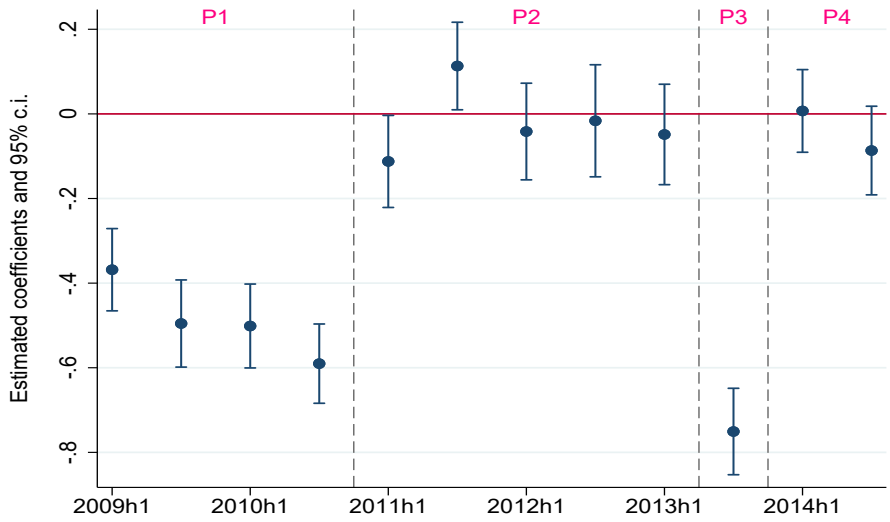

Figure 7: Differential tax incentives lead to different sale volume

Notes. The figure displays β coefficients (and 95% confidence intervals) from equation (1). The regression results suggest that during the the tax exemption period (P2) treated investors (those who sold another housing unit in the previous 1.5-4 years period) sold as many units as other investors. Before and after the exemption period (P1 and P3) the share of sales by treated investors is lower. Finally, during P4, when the elapsed time since the previous sale is no longer relevant for calculating the capital-gains tax, we also do not observe differences between types of investors. See text for details.

majority of the sales were of non-primary housing, which are likely to be rental housing, and to non-investors. We are now ready to explore the impact of this policies on prices.

4 The Effects of Investors' Sales on House Prices and Rents

4.1 Measuring the effects of sales on prices

How do sales by investors affect local house prices and rents? Our aim is to identify the impact of investors' sales on market prices, rather than their impact on the prices of investors' units. We build on our findings that the capital-gains tax exemption increases sales by the treated investors to estimate causal links between investor sales and prices. The relationship between sales and prices is typically difficult to measure as prices and quantities are determined simultaneously based on a multitude of supply and demand factors. In particular, prices of assets such as housing units, their demand, and their supply are all affected by beliefs which cannot be directly controlled for. To overcome these endogeneity issues, we construct a supply shifter based on the local-market composition of investor ownership, and use it to measure the response of house prices and rents along the local demand curve. The idea is that local markets in which there were more treated investors would see a larger increase

Panel A: Type of units sold by investors

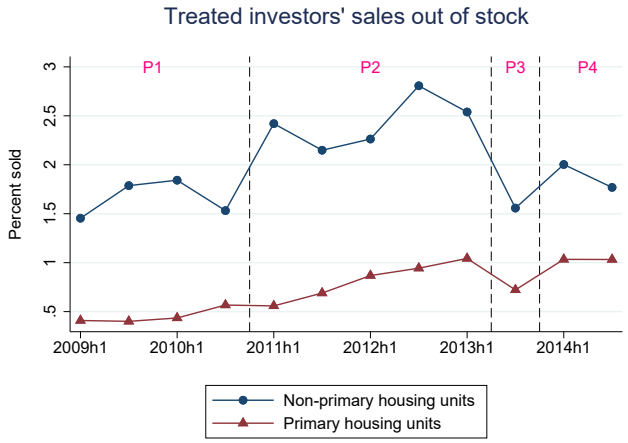

Panel B: Type of buyers to whom investors sell

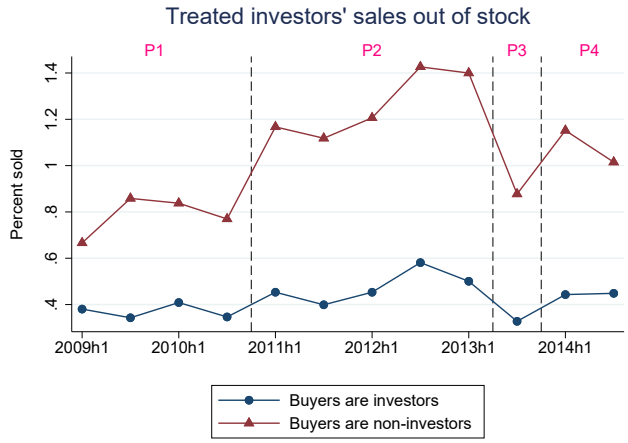

Figure 8: The effect of tax changes on treated investors

Notes. The figure displays the share of sales out of housing stock by "treated investors" between 2009 and 2014. Panel A plots these shares separately for sellers who sell their primary housing units and sellers who sell their non-primary units. Panel B plots the shares of sales separately for housing units sold to investors vs. non-investors buyers. Panel A show that investors are more likely to sell non-primary units, and Panel B shows that investors sell predominantly to non-investors, particularly during the tax exemption period (P2). *Sources*. Data on sales are from the Israel Tax Authority. Data on seller's residence are from Israel's population registry.

in supply which is uncorrelated with other local-market demand shocks. This approach relies on having considerable variation in investors' ownership share across local markets.

There is a significant variation in the composition of ownership across the 360 local housing markets in our data.[18] Figure 9 shows the ownership composition in 2010 across 171 local markets in Israel's central district (which is the most populated and includes Tel-Aviv). The left panel shows the share of units owned by investors. For instance, the local markets in Tel-Aviv (center-left) show a much larger share of investor-owned housing. The right panel shows the share of treated investors out of all investors. It demonstrates that a large variation in the composition of investors exists among adjacent local markets, even among those with a similar overall share of investors. Our approach exploits this variation to construct a supply shifter for each of market and half-year period.

We construct the local-market supply shifter in the following way. First, we use the diff-in-diff regression presented in Figure 6 to construct the expected sales rate, $\widehat{\text{sales rate}}_{sjt}$, for investors of type s (*treatment* or *control*), in local market j, and at time t (every half-year from

[18]Our definition of a local-housing market is based on "sub-quarters". These geographic units are constructed by the Central Bureau of Statistics (CBS) and are supposed to cover between 3,000 and 5,000 housing units. The delineation process aims at capturing meaningful local social and demographic characteristics. In addition, each sub-quarter is under the jurisdiction of one municipality.

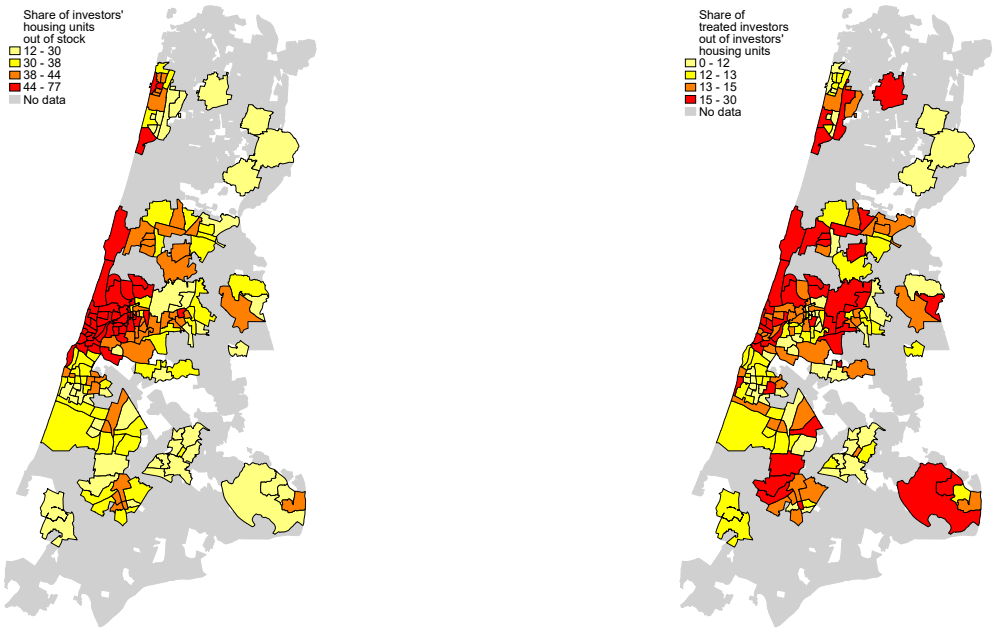

Figure 9: Ownership composition of housing units, Israel's central district, 2010

Notes. The map highlights 171 local markets (sub-quarters) in 30 municipalities at the center of Israel. For each market, investors' share of housing stock (on the left) and the share of treated investors' stock out of housing units owned by all investors (on the right) are presented. The map displays considerable heterogeneity in the share of housing units owned by investors, and in housing units owned by treated investors. *Sources*. Data on the housing stock are from 76 municipalities' property tax files. Data on sales are from the Israel Tax Authority.

2009H1 to 2014H2). We then construct a local predictor of sales by investors using the investor composition at the beginning of the period, share of units$_{sjt}$,

$$\widehat{\text{investors' sales}}_{jt} = 100 \times \sum_{s} \widehat{\text{sales rate}}_{sjt} \times \text{share of units}_{sjt},$$

The variable "investors' sales" is measured in units of percentage points as a share of the local housing stock. Our main specification is a linear regression of the outcome variable y_{ijt} at the transaction or the lease level i on the local-market supply shifter $\widehat{\text{investors' sales}}_{jt}$, controling for housing unit characteristics X_i, and time and local-market fixed effects,

$$\log y_{ijt} = \beta \widehat{\text{investors' sales}}_{jt} + \gamma X_i + \delta_j + \theta_t + \varepsilon_{ijt} \tag{2}$$

The outcome variables are the price of a transacted housing unit and the rent on a lease. The coefficient of interest is β. Housing unit controls include the number of rooms, the logarithm of the floor area, and the building age.[19]

[19] In case of rents, the floor area and building age are not available, so we only control for the number of rooms.

4.2 Main results: house prices and rents

Table 2 reports the estimates for specification (2) with transaction house prices as the outcome variable. Column 1 includes all transactions. We find that an increase in sales by investors of 1 percentage point of the housing stock reduces local house prices by 14%. Columns 2-5 of Table 2 report estimates of the same specification on restricted samples based on unit characteristic. Column 2 restricts the sample to units with 3 rooms or less (equivalent to a two-bedroom unit in the US). The effect is larger and more significant for these smaller units. Column 3 reports the estimate for the complementary group (large units) and finds a similar point estimate as the full sample, albeit less accurately measured and as a result statistically insignificant. Columns 4 and 5 split the sample into new units, which are bought by households from a constructor, and resale units. The point estimate in the resale-units sample is similar to the effect in the full sample, but the estimate in the new-units sample is closer to zero and insignificant.

Table 2: The effect of investors' predicted sales on house prices

	(1) all units	(2) $\leq$ 3 rooms	(3) $>$ 3 rooms	(4) new units	(5) resale units
Predicted sales	-0.14*	-0.24***	-0.14	-0.05	-0.13*
by investors (/stock)	(0.08)	(0.08)	(0.09)	(0.09)	(0.08)
Unit characteristics	✓	✓	✓	✓	✓
Half year FE	✓	✓	✓	✓	✓
Local market FE	✓	✓	✓	✓	✓
R^2	0.75	0.72	0.69	0.75	0.74
Observations	399,081	179,974	219,107	80,258	318,823

Notes. The table presents OLS regression results of Equation (2), where transacted price is the dependent variable. Column (1) includes all housing units transacted between 2009 and 2014. Columns (2)-(5) use different subsets of transacted housing units; "$\leq$ 3 rooms" are units with 1.5-3 rooms; "$>$ 3 rooms" are units with 3.5-5 rooms; "new units" are recently built units sold by companies, and "resale units" are non-new units. "Predicted sales by investors" is the number of housing units predicted to be sold by investors (divided by the total number of housing units) in each local market and half year. All regressions include unit characteristics (rooms, area and building age), half-year and local market fixed-effects. Standard errors are clustered by local market and shown in parentheses. The results suggest that additional sales of 100 units out of stock decrease prices by 14%. The decrease is larger for smaller housing units (Column 2) and is statistically significant for resale units (Column 5). For new units, we do not find evidence for a change in house prices due to the increase in sales by investors.

The differences across sub-samples can be explained by imperfect substitution between housing quality segments. If, say, small and large units are not close substitutes, and the supply of small units shifts out, the prices of small units should typically move more than

that of larger units. In the data, investor owned non-primary housing units tend to be smaller and older than owner occupied housing (see Table 1). Thus, the additional units sold by investors are also more likely to have 3 rooms or less. Also, according to our definition (and the tax law) a new housing unit sold by a constructor is not considered as a sale by investors. Thus, the price movement is more pronounced in the small-unit and resale-unit segments than for new and large units.[20]

Table 3: The effect of investors' predicted sales on rent

	(1) All units	(2) new leases	(3) new leases $\leq$ 3 rooms	(4) new leases >3 rooms	(5) extended leases
Predicted sales	0.01	0.10*	0.14**	-0.03	-0.03
by investors (/stock)	(0.04)	(0.06)	(0.07)	(0.12)	(0.05)
Unit characteristics	✓	✓	✓	✓	✓
Half year FE	✓	✓	✓	✓	✓
Local market FE	✓	✓	✓	✓	✓
R^2	0.75	0.77	0.74	0.75	0.74
Observations	64,817	19,436	11,872	7,564	45,381

Notes. The table presents OLS regression results of Equation (2) using rent as the dependent variable. Column (1) includes all rental units in the survey. Columns (2)-(5) utilize the panel structure of the survey as tenants are repeatedly surveyed periodically until they leave the rental unit. New leases (Column 2) are leases that appear for the first time and extended leases (Column 5) are subsequent appearances of the same rental unit and tenants. Columns (3)-(4) are based on a subset of new leases; "$\leq$ 3 rooms" are units with 1.5-3 rooms, "$>$ 3 rooms" are units with 3.5-5 rooms. "Predicted sales by investors" is the number of housing units expected to be sold by investors (divided by the total number of housing units) in each local market and half year. All regressions include the number of rooms, half year fixed-effects and local market fixed-effects. Standard errors, clustered by local market, are in parentheses. The results suggest that rent at new leases increased in local markets where treated investors are predicted to sell more units. This increase is significantly larger for small units.

Table 3 reports the estimates with survey rents as the outcome variable. Column 1 uses all the available observations, including new leases and extensions. The estimated effect of investor sales on rent is close to zero and statistically insignificant. Column 2 repeats the estimation for new leases only. These are rents of leases signed with new tenants in that period. While the sample size greatly diminishes (19 thousand compared to 65 thousand), the estimated effect is much larger and significant: 1 pp increase in sales by investors increases

[20] Appendix Table A1 repeats this analysis but replaces the predicted sales by investors as a percentage of the local housing stock with the predicted number of units sold by investors divided by 100. We find that a 100 additional sales by investors (2.5 pp of stock on average) reduce prices by 34%, which closely matches our main specification with similar significance levels, which points to the robustness of these estimates.

rent by 10%. This is our preferred estimate of the effect of investors' sales on rents. Column 5 repeats the exercise for the extended leases only, in which no effect of predicted sales on rents is observed. A likely explanation is that extended leases exhibit inertia, and are more insulated from changing market conditions. Columns 3 and 4 further split the sub-sample of new leases into smaller and larger units based on the number of rooms. The estimated effect increases to 0.14 and is significant for small-unit leases (3 rooms or less, equivalent to two-bedroom units in the US). We do not find an effect on rent for larger units.[21]

Our findings are consistent with our estimates for the impact of investors sales on prices shown in Columns 2 and 3 of Table 2. There we also find that the price of small units decreased more than the price of larger units. A likely explanation is that investors own a relatively larger share of small units, and therefore the impact of their sales on both prices and rents is larger for this segment of the market.

4.3 Validity of the supply shifter

Similar to standard shift-share research designs, the interpretation of regression (2) as a causal demand relationship hinges on the exogeneity of the constructed supply shifter (see Borusyak et al., 2022). This is plausible under the condition that the aggregate shocks (here, driven by tax changes) are independent of local demand shocks, or the exposures (here, the investor composition) are independent of them. In our setting, we believe both assumptions are plausible. First, the tax changes, which drive the changes in expected sales rate of different investor types, were decided by the government and aimed at reducing house prices at the aggregate level. Second, the exposure of investors to the tax changes is based on their sale of another unit between 1.5 and 4 years earlier. Since the tax exemption was largely unanticipated and lasted only two and a half years, it would have been difficult to manipulate the exposure with respect to local demand shocks that were not realized yet.

One remaining concern is that the investor composition is correlated with unobserved characteristics that determine house prices trends, and thus predict their appreciation. In this case, the investor composition would also be correlated with house price appreciation *before* the tax changes. Figure 10 plots the growth of a hedonic price index from 2009H2 to 2010H2 (before the first tax change) and from 2010H2 to 2011H2 (right after) for each local market against the share of treated investors. In the pre-treatment period there is no visible or statistical correlation between price appreciation and the composition of investors (the coefficient of correlation is -0.06 and insignificant). But after the tax changes there is a clear negative relationship. This suggest that there was no pre-trend associated with the

[21] Appendix Table A2 repeats this analysis with the predicted the number of housing units sold by investors divided by 100 instead of their predict investor sales out of stock. Similar to the estimates of the effects on prices the magnitudes remain similar, though less significant.

composition of investors.

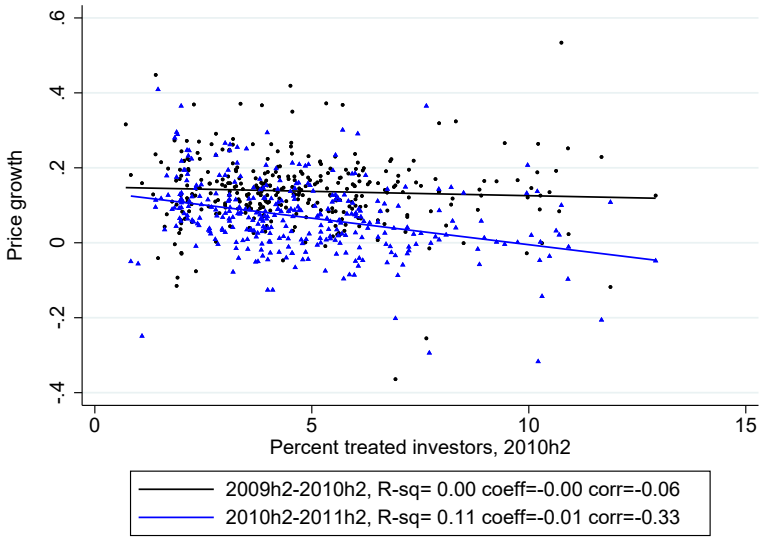

Figure 10: House price growth and percent of treated investors at the end of 2010

Notes. The figure displays house price growth across local markets separately for 2010 (i.e from 2009H2 to 2010H2) and for 2011 (2010H2 to 2011H2), against the share of stock held by treated investors in the end of 2010. According to the figure, the presence of treated investors in a local market had no association with the local price growth when investors were subject to the tax in 2010. In contrast, when treated investors were exempted from the capital-gain tax in 2011, a negative relationship arises between their presence in a local market and house prices. See text for details. *Sources*. Data on the housing stock are from 76 municipalities' property tax files. Data on sales are from the Israel Tax Authority.

4.4 Discussion of magnitudes

We conduct two back-of-the-envelope calculations to gauge the magnitude of the estimated coefficients of sales by investors on house prices and rents. We consider two scenarios. The first calculation aims at estimating the effect of the actual exemption of treated investors on prices. Since investors own 36% of all housing units on average, 13% of them are in the treatment group, and those treated by an exemption increased their expected sales rate by 0.6 pp, the mean increase in sales by investors during the exemption period is then

$$36\% \times 13\% \times 0.6\text{pp} = 0.028\text{pp}.$$

This is the share of the housing stock sold each half-year due to the exemption. We then multiply this number by our preferred point estimates for the coefficient on predicted sales by investors. We find that the exemption reduced house prices by 0.4% and increased rents on new leases by 0.28%. These are modest, though not negligible impacts on prices. The calculation suggests though that the impact of this particular policy on prices could not have been large since the treated group is quite small.

The second calculation considers a counterfactual policy in which all exemptions from capital-gains tax are removed, and the tax rate is set at 25%. Up to January 2014 investors who did not sell another unit in the previous 4 years, which account for roughly 82% of all investors, were exempt from capital-gains tax (those who sold another unit in the last 1.5 years, account for the remaining 5%). This is a relevant counterfactural as the changes in the tax code in January 2014 gradually increase the capital-gains taxes to that value. Under the assumption that the impact of an increase in capital-gains tax is the same for this group, their sales rate should decrease by 0.6 pp in the counterfactual scenario. Thus, the share of the housing stock that is *not* sold due to the counterfactual is

$$36\% \times 82\% \times 0.6\text{pp} = 0.18\text{pp}.$$

Therefore, removing all exemption from capital-gains tax on investors would result in an increase of 2.5% in house prices, and a decline of 1.8% in rents on new leases. These are larger values than the impact of the actual policy due to the larger group of affected investors.

These simple calculations highlight an important channel through which policymakers can affect housing affordability and homeownership. That said, these calculations should be read carefully, as they only account for the effect of the tax policy through their lock-in effect on existing investors and do not take into account other equilibrium effects. For instance, increasing the capital-gains tax rate also reduces the attractiveness of residential real-estate investment, and so investors may choose to avoid this type of investment.

4.5 The impact of the ownership structure

In the final analysis we consider a more speculative exercise that explores a particular channel through which sales by investors can affect prices – the ownership structure channel. The idea of this channel is that sales by investors change the share of housing units owned by investors ("the investor share"). When the investor share increases, the supply of housing to single-unit owners decreases, and the supply of rental housing units increases which affect prices in their respective markets. If we ignore other potential effects of investor sales, say, through search frictions, and attribute all the effect to the ownership structure channel, then the constructed variable of sales by investors can serve as an instrument for changes in the ownership share of investors.

We thus consider an instrumental variable estimates of the impact of changes in the ownership share of investors, ΔS_{jt} on house prices and rents,

$$y_{ijt} = \beta \Delta S_{jt} + \gamma X_i + \delta_j + \theta_t + \varepsilon_{ijt},$$

using the predicted sales by investors $\widehat{\text{investors}'}\ \text{sales}_{jt}$ as an instrument.

Table 4 reports the results of this exercise. Column 1 shows a regression of changes in the ownership share of investors on the predicted sales by investors, which are negative and

Table 4: Change in investors' share of stock as an instrument

	Δ investor share	house prices			rents		
		all	≤ 3 rooms	resale only	all	≤ 3 rooms	new leases
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Predicted sales by investors (/stock)	-0.05*** (0.01)						
Δ investor share		2.66** (1.30)	4.08*** (1.37)	2.68*** (1.43)	-1.01 (0.87)	-2.75*** (1.02)	-1.53 (1.11)
Unit characteristics		✓	✓	✓	✓	✓	✓
Half year FE	✓	✓	✓	✓	✓	✓	✓
Local market FE	✓	✓	✓	✓	✓	✓	✓
R^2	0.30	0.10	0.26	0.20	0.06	0.24	0.46
Observations	3,960	365,141	165,933	293,565	56,013	34,586	21,439

Notes. Columns (2)-(7) in the table presents IV regression results where house prices (columns 2-4) and rents (columns 5-7) are the dependent variable, and the change in investors share of stock is the instrument. In column (1) we estimate equation (2) where the outcome variable is the change in investors' share of stock. "Δ investors share" is the share of investors out of stock at the end of the period (half-year) minus their share at the beginning of the period. "Predicted sales by investors" is the number of housing units expected to be sold by investors (divided by the total number of housing units) in each local market and half year. All regressions include half year and local market fixed-effects. IV regressions include unit characteristics. Standard errors, clustered by local market, in parentheses. The results suggest that the share of investors fell in markets where investors' sales were predicted to be higher (column 1). Moreover, an increase in investors' local activity likely result in higher prices, particularly for non-new, smaller units, and also for lower rents at small housing units.

highly significant. This regression is an unweighted version of the first stage in all the other IV regressions. Columns 2-4 show IV regression of house prices on changes in investor shares for three samples: the full sample, units with 3 rooms or less, and resale only sample. All three regressions yield significant values. Based on Column 2, a 1pp increase in the investor share increases local-market prices by 2.66%. Under the assumptions outlined above, this is an estimate of the inverse semi-elasticity of demand concerning the investor's share of ownership. The point estimate for the smaller unit is 4.08, which suggests that local demand in this segment is quite inelastic. Columns 5-7 report estimates for IV regressions with rent as the outcome variable. Column 5 reports the estimate on the full sample, which is -1.01 and is insignificant. Column 6 reports the estimate for units with 3 rooms or less, which are larger in magnitude and more significant (-2.75, significant at the 0.01 level). The lower elasticity of demand in the small unit segments in both the renal and ownership markets supports the

hypothesis that the households that consume housing services in these markets have a more constrained choice set.

5 Model

We propose a model of a local housing economy with heterogeneous households. There are two types of households in the economy: "residents" and "investors". The only real asset is houses, which are homogeneous, indivisible, and in fixed supply. Trade is conducted in two linked housing markets: a market for houses and a market for rental housing services.

The model combines features from two frameworks. First, the demand for owned and rental housing is determined by residents who choose to own or rent based on their heterogeneous ability to extract utility from each housing option, as was proposed by Weiss (1978). This heterogeneity links the ownership structure of the housing stock to the rent-to-price ratio. We extend this framework with housing indivisibility, which forces some residents to consume no housing and so decouples the elasticity of demand for owned houses and rental housing services.

Second, the supply of rental housing is determined by investors who enter and exit the market infrequently, similar to the way firms do in the framework proposed by Hopenhayn (1992). Investors buy and sell houses and rent out the houses they own. They may own one or two houses at any moment. Their individual trade decisions are driven by differences in taxes and idiosyncratic maintenance costs, which are subject to persistent shocks. Trend growth in the demand for housing services by residents generates price appreciation. Capital gains accumulate, and capital-gains taxes on investors which are history-dependent generate a lock-in effect. House prices, rents, and the ownership composition of the housing stock are determined in equilibrium.

The model replicates two key features of the data. First, investors' sales rate jumps discontinuously when they become exempt from capital gains tax. Importantly, in this model, the timing of the tax exemption is endogenous, as investors who own multiple houses form a sophisticated contingency plan for selling their own houses in anticipation of the future tax rate they will face. Second, a higher homeownership rate by residents–for instance, due to increased taxes on rent–is associated with lower house prices and higher rents. The primary mechanism is as follows. When investors sell their owned houses to residents, some resident-renters stop consuming housing services completely and new resident-owners buy the houses. Differences between the former renters and the new owners allow house prices and rents to move in opposite directions by eliminating the user cost condition that appears in standard models.

5.1 Technology and preferences

Time is continuous and $t \geq 0$ denotes the date. The economy includes a unit mass of indivisible houses in fixed supply. Each house, at any date, is either owned by a resident or owned by an investor and rented to a resident (no vacancies). Houses and rental housing services are traded in spot markets at prices p_t and r_t, respectively. In addition, there is a market for risk-free bonds which households can use for saving and borrowing, up to a natural borrowing constraint.

There is a mass $R > 1$ of resident households who are risk-neutral, derive utility from housing h and a numeraire consumption good q, and discount time at rate ρ. Let $h_t(j)$ be the housing choice of household j at date t, which can be either "own", "rent", or "none",

$$h_t : [0, R] \to \{\text{own}, \text{rent}, \text{none}\}.$$

Since houses are indivisible and homogeneous, when households own or rent they consume the housing services provided by exactly one house. The flow utility from housing services is captured by a two-dimensional preference shock

$$\theta_t(j) = \left(\theta_{\text{own},t}(j), \theta_{\text{rent},t}(j)\right),$$

and is given by

$$u(h, q; \theta, t) = \begin{cases} e^{gt}\theta_{\text{own}} + q & \text{if } h = \text{own}, \\ e^{gt}\theta_{\text{rent}} + q & \text{if } h = \text{rent}, \\ q & \text{if } h = \text{none}, \end{cases}$$

where $0 < g < \rho$ is the rate of growth of housing demand. The preference shock $\theta_t(j)$ is updated at a Poisson arrival rate μ, in which new i.i.d. values are drawn from a distribution $F(\theta)$ with support $\mathbb{R}^2_+$. Residents' preferences are therefore quasi-linear, which implies that the demand for housing is not affected by aggregate wealth or the wealth distribution. The growth rate g stands in for the effect of growing aggregate income and wealth on the demand for housing services.[22]

New investors are born at rate ν with zero wealth and no housing units. At birth, each investor i is assigned a maintenance cost $M_t(i) > 0$. The logarithm of the maintenance cost, $m_t = \log M_t(i)$, follows an Ornstein-Uhlenbeck process,

$$dm_t(i) = -\kappa(m_t(i) - \bar{m})dt + \sigma dB_t(i),$$

[22]This specification can be seen as a reduced form for the case in which resident households perfectly share risks within larger "families". The family assigns the housing choice and consumption choice of each household to maximize a welfare function. Based on such interpretation, g can be seen as the result of growth in the family income, and the distribution of θ as a reduced form reflecting differences in wealth across families.

where $\kappa, \sigma, \bar{m} > 0$ are parameters and $B_t(i)$ is an independent standard Brownian motion process. The initial logarithm of maintenance cost is drawn from the stationary distribution of this process, which is normal with mean $\bar{m}$ and standard deviations $\sigma/\sqrt{2\kappa}$. Upon realizing their maintenance cost, investors decide whether to buy a house.

Investors can own 1 or 2 houses. Let $n_t(i) \in \{1, 2\}$ be the number of houses owned by household i at date t. The flow income from rental housing services is

$$y_t = n_t(i) r_t - e^{gt} e^{m_t(i)} \phi(n_t(i)),$$

where $\phi(n)$ is a maintenance cost multiplier that satisfies $\phi(1) = 1$ and $\phi(2) > 2$, reflecting convexity in maintenance costs. The exponential growth component in maintenance costs reflects the assumption that input prices of maintenance, e.g., wages, growing at the same rate as housing demand.

Investors can buy or sell houses only infrequently. Buying and selling opportunities arrive randomly at a Poisson arrival rate λ_B for buying opportunities and λ_S per house for selling opportunities. When *selling*, investors pay capital gains tax and transaction costs. Both the tax rate and the tax base depend on their transaction history. The tax rate depends on the time since the last sale, which we denote $s_t(i)$ ($s_t(i)$ increases with time until reaches a maximum value $\bar{s}$). The dependence is captured by a function $\tau(s_t(i))$, which can be discontinuous, and we assume is weakly decreasing. We also assume that the tax rate has an upper bound $\bar{\tau} < 1$. The tax base is the accumulated capital gains which depend on the time of purchasing. We denote the time since buying house number $n \leq n_t(i)$ by $b_{n,t}(i)$ (we assume $b_{1,t}(i) > b_{2,t}(i)$ if investor i owns two houses). The tax base for selling the nth house at date t is $\max\{p_t - p_{t-b_{n,t}(i)}, 0\}$. In addition, selling investors pay transaction costs at the rate Δ of the sale price. The selling opportunities arrive independently, and so the decision to sell is made separately for each house. The payout from the sale of the nth house is the sale price net of transaction cost minus the capital gains tax,

$$(1 - \Delta) p_t - \tau(s_t(i)) \max\{p_t - p_{t-b_{n,t}(i)}, 0\}.$$

Immediately upon sale, the time since the last sale $s_t(i)$ is set to zero, and, if an investor has owned two units before the sale and house number 1 was sold, the number of house number 2 is updated to 1.

When *buying* a house investors pay a purchase tax at a rate of T on the market price, so that the after-tax cost is $(1 + T) p_t$. Upon purchase, the number of owned houses $n_t(i)$ increases by 1, and the time since buying the house is set to zero, $b_{n_t(i),t}(i) = 0$.

Investors *exit* when they own one unit and sell it, or when they receive an exogenous "death" shock at rate δ. We assume that upon death, all houses are immediately sold. Thus, the mass of active investors and the joint distribution of their maintenance costs and ownership

of houses, including their history of transactions are determined endogenously. We denote the set of active investors by I_t.

A *feasible allocation* includes the joint distribution of preference shocks θ and tenure choices h across residents and the joint distribution of investors across maintenance cost m, the number of houses owned n, and the transaction history s and b at all dates, such that each house has exactly one owner and that each investor-owned house is rented to exactly one resident. The conditions that the services generated by all houses are consumed can be written as

$$\int_0^R I(h_t(j) = \text{own})dj + \int_0^R I(h_t(j) = \text{rent})dj = 1,$$

and market clearing in the rental housing services market equates the number of investor-owned houses and residents who rent,

$$\int_{I_t} n_t(i)di = \int_0^R I(h_t(j) = \text{rent})dj.$$

Implicitly, the two conditions above also imply that the spot market for houses clears and that all houses have owners.

An important note is that this model abstracts from wealth effects by assuming all households have quasi-linear preferences. This means that the model remains silent on the distribution of income, numeraire consumption, and wealth, which are left outside the scope of the model.

5.2 Stationary equilibrium

We define a stationary equilibrium as prices $\{p_t, r_t\}$ for all dates and a feasible allocation in which house prices and rents grow at a constant rate g, residents and investors follow optimal paths, and their distributions remain constant over time. That is, for residents, the joint distribution of housing choice and preference shocks is time-invariant, and for investors, the joint distribution of the number of houses owned, maintenance costs, time since the last sale, and time since the houses were bought is time-invariant.[23] We now describe the equilibrium conditions derived from the optimal choices of residents and investors.

Residents. The constant growth rate in prices implies that they can be written as $p_t = e^{gt}p_0$ and $r_t = e^{gt}r_0$. Since residents can trade continuously in houses and risk-free bonds, the resident problem is simpler. Let $c_t = p_t(\rho - g)$ be the user cost of a house, and c_0 be its initial

[23] We resolve indifference between housing options in the following way. Residents who are indifferent between owning and renting will choose to own, and if indifferent between consuming and not consuming housing services choose to consume. Investors who are indifferent between selling or buying or staying put, choose to stay put.

value. Then, the resident choice can be written as a time-invariant rule given by

$$h_t(j) = \begin{cases} \text{own} & \text{if } \theta_{\text{own},t} \geq c_0 \text{ and } \theta_{\text{rent},t}/\theta_{\text{own},t} \leq r_0/c_0, \\ \text{rent} & \text{if } \theta_{\text{rent},t} \geq r_0 \text{ and } \theta_{\text{rent},t}/\theta_{\text{own},t} > r_0/c_0, \\ \text{none} & \text{otherwise.} \end{cases}$$

The optimal housing choice by residents is illustrated in Figure 11. The two dashed lines represent the conditions that owning or renting is preferred to staying put. The diagonal line represents the indifference between choosing to own or rent. The blue region is residents who choose to own and the green region to buy. Since the distribution of $\theta_t(i)$ is stationary, and the decision rule (given stationary equilibrium prices) is time-invariant, residents stay at the same housing choice until the arrival of a new preference shock.

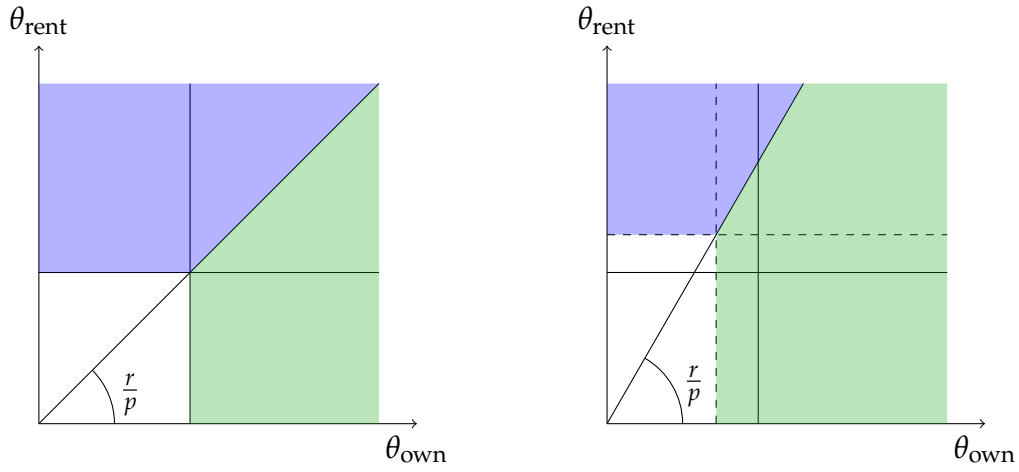

(a) low r/c (b) high r/c

Figure 11: Determination of the demand for housing services

Investors. Optimal decisions by investors are captured by a value function and Hamilton-Jacobi-Bellman (HJB) equations. The state of an investor is captured by her log maintenance cost $m_t(j)$, the number of units she owns $n_t(i)$, and items from her history of transactions: the time since the last sale $s_t(i)$ and the time since buying each unit $b_{n,t}(i)$ for $1 \leq n \leq n_t(i)$. Since we focus on a stationary equilibrium in which maintenance costs, rents, and prices grow at a constant rate we can safely conjecture that the expected present discounted value (PDV) of the net income by an investor also grows at rate g. The PDV can be expressed by $e^{gt}V^1(m, s, b_1)$ and $e^{gt}V^2(m, s, b_1, b_2)$ where V^1 and V^2 are the normalized value functions for investors who

own 1, and 2 houses, respectively. They must satisfy the following HJB equations:

$$
\begin{aligned}
(\rho+\delta-g)V^1(m,s,b_1) &= r_0 - e^m\phi(1) - \qquad (3)\\
&\quad -\kappa(m-\bar{m})\frac{\partial V^1}{\partial m}+\frac{1}{2}\sigma^2\frac{\partial^2 V^1}{\partial m^2}+\frac{\partial V^1}{\partial s}+\frac{\partial V^1}{\partial b_1}+\\
&\quad +\lambda_B \max\{V^2(m,s,b_1,0)-(1+T)p_0-V^1(m,s,b_1),0\}+\\
&\quad +\lambda_S \max\{p_0(1-\Delta-\tau(s)(1-e^{-gb_1}))-V^1(m,s,b_1),0\}+\\
&\quad +\delta p_0(1-\Delta-\tau(s)(1-e^{-gb_1})),\\
(\rho+\delta-g)V^2(m,s,b_1,b_2) &= 2r_0 - e^m\phi(2) - \qquad (4)\\
&\quad -\kappa(m-\bar{m})\frac{\partial V^2}{\partial m}+\frac{1}{2}\sigma^2\frac{\partial^2 V^2}{\partial m^2}+\frac{\partial V^2}{\partial s}+\frac{\partial V^2}{\partial b_1}+\frac{\partial V^2}{\partial b_2}+\\
&\quad +\lambda_S \max\{V^1(m,0,b_2)+p_0(1-\Delta-\tau(s)(1-e^{-gb_1}))-V^2(m,s,b_1,b_2),0\}+\\
&\quad +\lambda_S \max\{V^1(m,0,b_1)+p_0(1-\Delta-\tau(s)(1-e^{-gb_2}))-V^2(m,s,b_1,b_2),0\}+\\
&\quad +\delta p_0(2-2\Delta-\tau(s)(1-e^{-gb_1})-\tau(s)(1-e^{-gb_2})).
\end{aligned}
$$

The entry decision depends on the initial draw of the maintenance cost. Investors enter if their expected present discounted value of their cashflow is greater than the after-tax cost of a house,

$$V^1(m,\bar{s},0)\geq p_0.$$

Therefore, there is a threshold maintenance cost, m_e, such that investors enter if $m_t(i)<m_e$. The entry rate is given by $\nu\Phi\left(\sqrt{2\kappa}\frac{m_e-\bar{m}}{\sigma}\right)$, where $\Phi(\cdot)$ is the standard normal CDF.

The HJB equations illustrate how elements of dynamic asset pricing affect the private value of houses to investors. The effective discount rate applied to the normalized value function is $\rho-g$, which reflects the expected appreciation of houses. This implies that if log maintenance costs were constant and equal to the mean $\bar{m}$, and taxes were set to zero, the present discounted value at date 0 of the cash flows from a single house held in perpetuity would be $(r_0-e^{\bar{m}})/(\rho-g)$.

The value of new investors is derived from the cash flow coming from their house plus the option to buy more houses: They have the right, but not the obligation, to buy another house. The effect of uncertainty in maintenance cost, as captured by the diffusion parameter σ, depends on the second derivative of the value function with respect to m. Typically, options generate convexity in value, in which case the value of the investor increases with uncertainty. Due to the option value and the variability of maintenance costs, the value of a house to an investor may be higher than the market price. However, due to infrequent sales opportunities and capital gains, the private value of owning a house may also be lower than the market price and potentially even negative. Therefore, when maintenance costs are high enough, the investor chooses to sell when possible.

Value is decreasing in the time since buying a house b_n. Under the assumption that $g > 0$, a larger share of the market price of a house is considered capital gains, and is therefore taxable. The direction of the effect of the time since the last sale, s, crucially depends on the tax policy $\tau(s)$. Since we assumed that the tax policy rate is non-increasing in s, the value of an investor increases with s.

While the trading technology allows investors who own one house to be sellers (i.e., willing to sell the house given the opportunity) and buyers (i.e., willing to sell the house given the opportunity) simultaneously, this would not be the case under broad conditions. Investors who own one house and face the maximal capital gains tax will only exercise the option to buy or the option to sell, not both, even when other taxes are set to zero. The reason why is captured by the following series of lemmas.

Lemma 1. The value of an investor who just bought a second house $V^2(m, s, b_1, 0)$ is less or equal to the combined value of two investors owning one house each which were bought on the same dates, $V^1(m, s, b_1) + V^1(m, s, 0)$ for all possible m, s, and b_1.

Proof. Suppose by way of contradiction that $V^2(m, s, b_1, 0) > V^1(m, s, b_1) + V^1(m, s, 0)$. Then,

$$\begin{aligned} V^1(m,0,0) + \hat{p}(1-\Delta-\tau(s)(1-e^{-gb_1})) - V^2(m,s,b_1,0) \leq \\ \leq V^1(m,s,0) + \hat{p}(1-\Delta-\tau(s)(1-e^{-gb_1})) - V^2(m,s,b_1,0) < \\ < \hat{p}(1-\Delta-\tau(s)(1-e^{-gb_1})) - V^1(m,s,b_1), \end{aligned}$$

and

$$\begin{aligned} V^1(m,0,b_1) + \hat{p}(1-\Delta) - V^2(m,s,b_1,0) \leq \\ \leq V^1(m,s,b_1) + \hat{p}(1-\Delta) - V^2(m,s,b_1,0) < \\ < \hat{p}(1-\Delta)) - V^1(m,s,0). \end{aligned}$$

Therefore the combined flow value from the option to sell each one of the houses by the two-house investor is lower than the flow value from the combined option to sell of the two one-house investors. In addition, the cash flow from owning two houses is less than twice the cash flow of owning one house due to the maintenance cost factor $\phi(2) > 2\phi(1)$. Finally, two one-house investors have buying options, which the one investor who owns two houses does not. Thus, the two one-house investors generate a combined flow value that is higher at any possible m, s, and b_1, and so their combined value must be greater than that of the single two-house investor. *QED*

Lemma 2. The difference between the value of an investor who just bought the first house, $V^1(m, s, 0)$, and the value of an investor who bought a house b_1 time before, $V^1(m, s, b_1)$, holding m and s fixed, is at most $p_0\bar{\tau}(1 - e^{-gb_1})$.

Proof. Let $\tilde{V}^1(m,s,b_1)$ be the value of an investor whose time since she last bought a house is b_1, but instead of acting optimally, makes buying and selling decisions as if she has just bought the house. This is not optimal, and so $\tilde{V}^1(m,s,b_1) \leq V^1(m,s,b_1)$. But compared with the optimal plan of the investor who just bought a house, the only difference in cash flow is one future extra tax payment when selling house number 1. Fix a date t and let $\tilde{t} \geq t$ be the time in which the investor sells the house. The present value of the tax payment at date t is

$$e^{-(\tilde{t}-t)r}\tau(s_{\tilde{t}}(i))(p_{\tilde{t}} - p_{t-b_{1,t}(i)}).$$

Since $r > g$ the contingency that maximizes this present value is if $\tilde{t} = t$ and $\tau(s_{\tilde{t}(i)}) = \bar{\tau}$. Since value functions are normalized, this implies that

$$V^1(m,s,0) - p_0\bar{\tau}(1-e^{-gb_1}) \leq \tilde{V}^1(m,s,b_1) \leq V^1(m,s,b_1),$$

which completes the proof.

Lemma 3. An investor subject to the maximal capital gains tax, $\tau(s_t(i)) = \bar{\tau}$, is either a seller or a buyer buy, not both.

Proof. If an investor with one house and m, s, and b, such that $\tau(s) = \bar{\tau}$ is a seller, then

$$V^1(m,s,b_1) - p_0(1-\Delta-\tau(s)(1-e^{-gb_1})) \leq 0.$$

Lemmas 1 and 2 imply that

$$V^2(m,s,b_1,0) - (1+T)p_0 - V^1(m,s,b_1) \leq V^1(m,s,0) - (1+T)p_0 \leq$$

$$\leq V^1(m,s,b_1) + p_0\bar{\tau}(1-e^{-gb_1}) - (1+T)p_0 < V^1(m,s,b_1) - p_0(1-\Delta-\tau(s)(1-e^{-gb_1})).$$

The last expression is less or equal to zero, and therefore a seller investor is not a buyer. *QED*

The main conclusion is that investors trade infrequently due to taxes, transaction costs, and infrequent trading opportunities. They accumulate houses when they have a low enough maintenance cost. They gradually sell their housing stock when they have very high maintenance costs, and take into account the tax implications of the sales. Investors also have wide regions of their state variables in which they just stay put.

5.3 Quantitative implementation

We select parameters of the model to capture realistic features of the Israeli tax code. The tax policy is set so that capital gains tax rate $\tau(s)$ is 0.25 when $s < 4$ and 0 when $s \geq 4$. Transaction costs are set to 6%, equivalent to typical real-estate agent fees. The purchase tax is set to 8%, which was its average rate in the sample period. The discount rate ρ is set to 3%, equal to the mean mortgage rate in the period 2009-2010. The expected growth rate of prices is 2%.

There are 8 remaining parameters on the investors' side. Parameters of the maintenance process, κ, σ, and m; the maintenance cost factor for owning two houses $\phi(2)$; the rate of arrival of trade opportunities λ_S, and λ_B; the flow of new investors ν; and the rate of arrival of the death shock δ. In addition, we parameterize the distribution of resident preference shocks so that their logs are drawn from a multivariate normal distribution with parameters μ_{own} and μ_{rent} and variance-covariance matrix Σ. The following illustration of the model is done with preliminary values, which capture the sales rate of investors around the cutoff.

5.4 Investors' sales rate in equilibrium

Figure 12 presents the discontinuity in investors' sales in the model simulation.

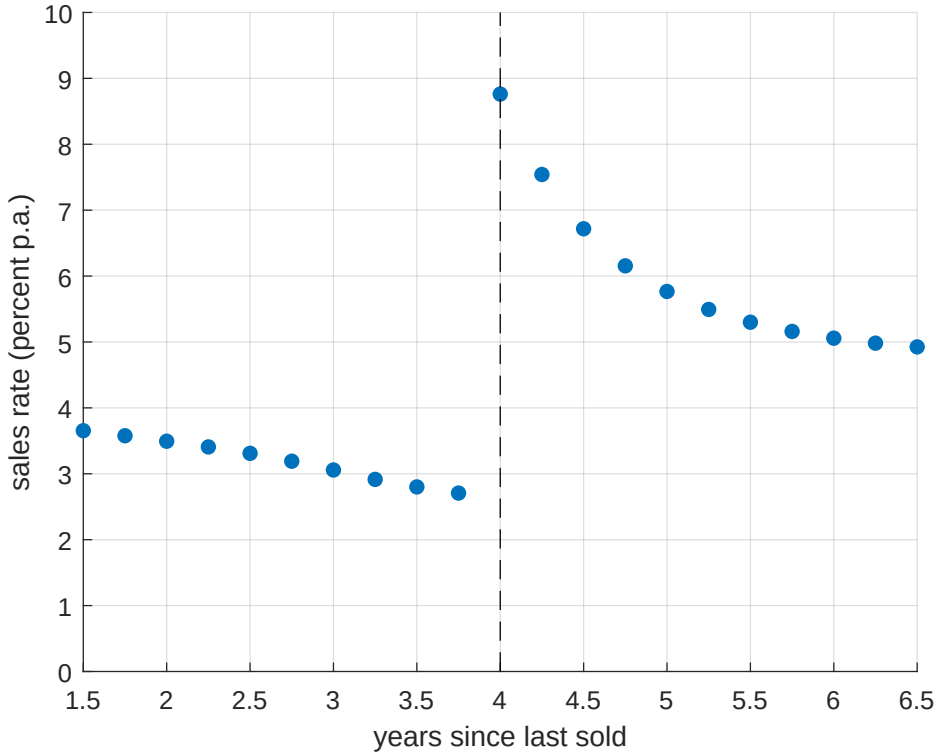

Figure 12: Discontinuity in investors' sales rate in the model

6 Conclusion

This paper achieves two goals. First, it sheds light on the efficacy of tax policy on investors in changing the ownership structure of the housing stock. We study a temporary tax exemption offered to some real-estate investors in Israel to encourage them to sell their housing units. We begin with estimating the effect of the tax exemption on sales by investors using a difference-in-differences approach. While the direction of the estimated effect is not surprising in itself, that is, the exemption eliminated a lock-in effect and led to a rise in sales, the magnitude is remarkable: A temporary exemption from a 25% capital-gains tax led the semiannual sales

rate of investor owned unit to increase by 50%, from 1.2 percentage points to 1.8. Importantly, we find that the additional sales are mostly of non-primary units, which are likely to be rentals, and to non-investor buyers, which implies a change in the ownership type of those units.

Second, we measure the implications of the ownership changes on local housing market conditions. We use variation in investor composition across 360 local markets to construct a local predictor of sales by investors. In our main regression specification, we estimate the local demand response of prices to the additional predicted sales by investors. Our preferred estimate is that an increase of sales by investors of 1 percentage point of the total housing stock is associated with a 14% drop in house prices. The same increase of sales by investors is also associated with lower supply of rental units, which leads to a 4% increase in rents of new leases. Both of these impacts of the investor sales on house prices and rents are stronger as we restrict attention to smaller units, in which real-estate investors specialize.

These results suggest that tax policies aimed at investors have the potential to significantly change the investors' behavior and the ownership structure of the housing stock. Moreover, a small local increase in the supply of housing units by investors significantly reduces local house prices and increases local rents. A high house price sensitivity suggests that the demand for housing units in a local market is quite inelastic in the short-term, and so is the demand for rental housing. From a theoretical perspective, this is at odds with theories that view local housing markets as close substitutes, and supports theories of spatially segmented markets (Piazzesi, Schneider, and Stroebel, 2020).

Our results also highlight the unintended distributional effects of housing policies. The specific policy under consideration–a temporary exemption from capital-gains taxes on investors–has clear winners and losers. Investors who sell their rental units with reduced taxes and first time home buyers who buy at a reduced price benefit, but poorer and more credit constrained renting households who face higher rents on new leases and other taxpayers who need to make up for lost tax revenue lose. We believe this trade-off should be an integral part of housing policy debates.

References

Stefania Albanesi, Giacomo DeGiorgi, and Jaromir Nosal. Credit growth and the financial crisis: A new narrative. *Journal of Monetary Economics*, 132:118–139, 2022.

Itai Ater, Yael Elster, David Genesove, and Eran B Hoffmann. Must agreements be kept? residential leases during covid-19. *The Economic Journal*, 133(649):477–492, 2023.

Cristian Badarinza and Tarun Ramadorai. Home away from home? Foreign demand and London house prices. *Journal of Financial Economics*, 130(3):532–555, 2018.

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

Patrick Bayer, Kyle Mangum, and James W Roberts. Speculative fever: Investor contagion in the housing bubble. *American Economic Review*, 111(2):609–51, 2021.

Michael Carlos Best and Henrik Jacobsen Kleven. Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK. *The Review of Economic Studies*, 85(1): 157–193, 2018.

Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1):181–213, 2022.

Philippe Bracke. How much do investors pay for houses? *Real Estate Economics*, 2015.

Alex Chinco and Christopher Mayer. Misinformed speculators and mispricing in the housing market. *The Review of Financial Studies*, 29(2):486–522, 2016.

Dragana Cvijanović and Christophe Spaenjers. "We'll always have Paris": Out-of-country buyers in the housing market. *Management Science*, 67(7):4120–4138, 2021.

Anthony A DeFusco, Charles G Nathanson, and Eric Zwick. Speculative dynamics of prices and volume. *Journal of Financial Economics*, 146(1):205–229, 2022.

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

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

Jack Favilukis and Stijn Van Nieuwerburgh. Out-of-town home buyers and city welfare. *The Journal of Finance*, 76(5):2577–2638, 2021.

Max Floetotto, Michael Kirker, and Johannes Stroebel. Government intervention in the housing market: Who wins, who loses? *Journal of Monetary Economics*, 80:106–123, 2016.

Zhenyu Gao, Michael Sockin, and Wei Xiong. Economic consequences of housing speculation. *The Review of Financial Studies*, 33(11):5248–5287, 2020.

Pedro Gete and Michael Reher. Mortgage supply and housing rents. *The Review of Financial Studies*, 31(12):4884–4911, 2018.

Jonathan Halket and Matteo Pignatti Morano di Custoza. Homeownership and the scarcity of rentals. *Journal of Monetary Economics*, 76:107–123, 2015.

Jonathan Halket and Santhanagopalan Vasudev. Saving up or settling down: Home ownership over the life cycle. *Review of Economic Dynamics*, 17(2):345–366, 2014.

Lu Han and William C Strange. The microstructure of housing markets: Search, bargaining, and brokerage. *Handbook of regional and urban economics*, 5:813–886, 2015.

Lu Han, L Rachel Ngai, and Kevin D Sheedy. To own or to rent? The effects of transaction taxes on housing markets. 2021.

Andrew Haughwout, Donghoon Lee, Joseph S Tracy, and Wilbert Van der Klaauw. Real estate investors, the leverage cycle, and the housing market crisis. *FRB of New York Staff Report*, (514), 2011.

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

J Vernon Henderson and Yannis M Ioannides. A model of housing tenure choice. *The American Economic Review*, 73(1):98–113, 1983.

Hugo A Hopenhayn. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1127–1150, 1992.

Daniel I. García. Second-home buying and the housing boom and bust. *Real Estate Economics*, 50(1):33–58, 2022.

Wojciech Kopczuk and David Munroe. Mansion tax: The effect of transfer taxes on the residential real estate market. *American Economic Journal: Economic Policy*, 7(2):214–57, 2015.

Steven Laufer and Nitzan Tzur-Ilan. The effect of LTV-based risk weights on house prices: Evidence from an israeli macroprudential policy. *Journal of Urban Economics*, 124:103349, 2021.

Antoine Levy. Housing policy with home-biased landlords: Evidence from French rental markets. 2021.

Ulrike Malmendier and Alexandra Steiny. Rent or buy? the role of lifetime experiences of macroeconomic shocks within and across countries. *Working paper*.

Monika Piazzesi and Martin Schneider. Momentum traders in the housing market: Survey evidence and a search model. *American Economic Review*, 99(2):406–11, 2009.

Monika Piazzesi, Martin Schneider, and Johannes Stroebel. Segmented housing search. *American Economic Review*, 110(3):720–59, 2020.

Hui Shan. The effect of capital gains taxation on home sales: Evidence from the taxpayer relief act of 1997. *Journal of Public Economics*, 95(1-2):177–188, 2011.

Todd Sinai and Nicholas S Souleles. Owner-occupied housing as a hedge against rent risk. *The Quarterly Journal of Economics*, 120(2):763–789, 2005.

Joel Slemrod, Caroline Weber, and Hui Shan. The behavioral response to housing transfer taxes: Evidence from a notched change in DC policy. *Journal of Urban Economics*, 100:137–153, 2017.

Tsur Somerville, Long Wang, and Yang Yang. Using purchase restrictions to cool housing markets: A within-market analysis. *Journal of Urban Economics*, 115:103189, 2020.

Kamila Sommer and Paul Sullivan. Implications of US tax policy for house prices, rents, and homeownership. *American Economic Review*, 108(2):241–74, 2018.

Nitzan Tzur-Ilan. The real consequences of LTV limits on housing choices. *Available at SSRN 2927783*, 2019.

Yoram Weiss. Capital gains, discriminatory taxes, and the choice between renting and owning a house. *Journal of Public Economics*, 10(1):45–55, 1978. ISSN 0047-2727. doi: https://doi.org/10.1016/0047-2727(78)90004-X. URL https://www.sciencedirect.com/science/article/pii/004727277890004X.

A Appendix

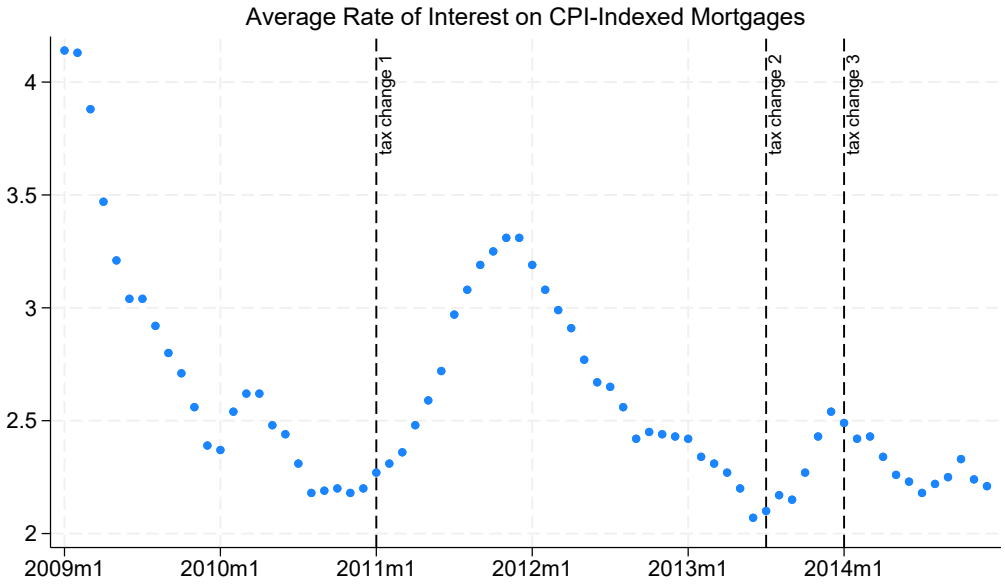

Figure A.1: Average Rate of Interest on CPI-Indexed Mortgages in Israel (2009-2014)

Notes. The figure displays Average Rate of Interest on CPI-Indexed Mortgages. *Sources*. Data on interest rates are from the bank of Israel (available online: `https://www.boi.org.il/information/interestrates/mortgage/`).

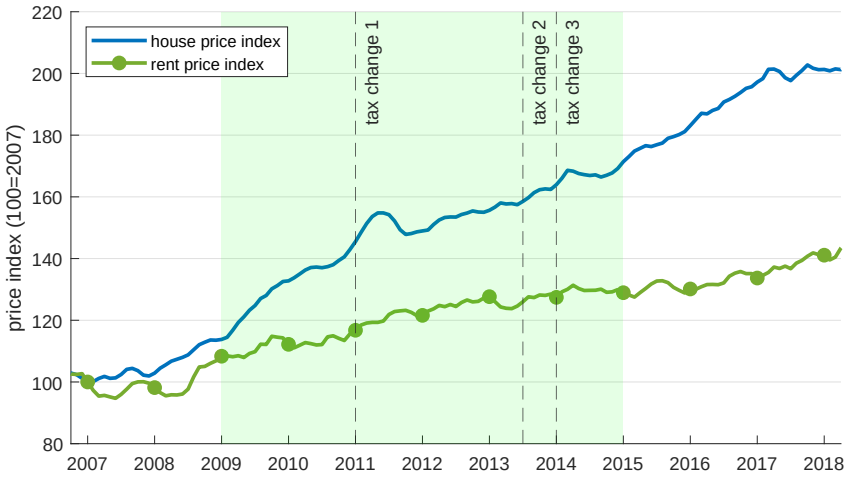

Figure A.2: House prices and rents in Israel (2007 - 2018)

Notes. The figure displays house price and rent indices based on hedonic regression of log house price (or rent) on unit characteristics (number of rooms for rents, number of rooms, log area, and age of structure for house prices), year-month and location (statistical area) fixed-effects. The value of each index is normalized to 100 in January 2007. *Sources*. Data on house prices are from the Israel Tax Authority. Data on rent are from a rental survey conducted by the Central Bureau of Statistics.

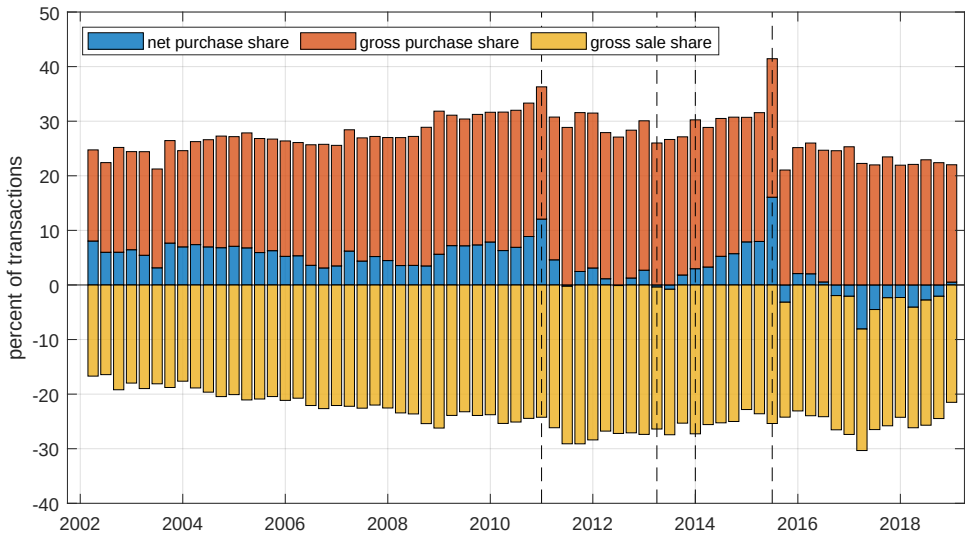

Figure A.3: Investors' share of housing transactions in Israel (2002 - 2018)

Notes. The figure displays the share of purchases and sales by investors out of all transactions in each quarter in Israel. Investors' share of purchases grew from about 25% in early 2000s to about 35% by the end of 2010. Even when deducting the housing units sold by investors in each quarter, the stock of housing units owned by investors grew steadily. In 2011, this trend slowed significantly after the government introduced higher purchase tax and an exemption from capital-gains tax for a subset of investors. The dates of the main tax reforms are marked by dashed vertical lines. *Sources*. Data on house transactions are from the Israel Tax Authority.

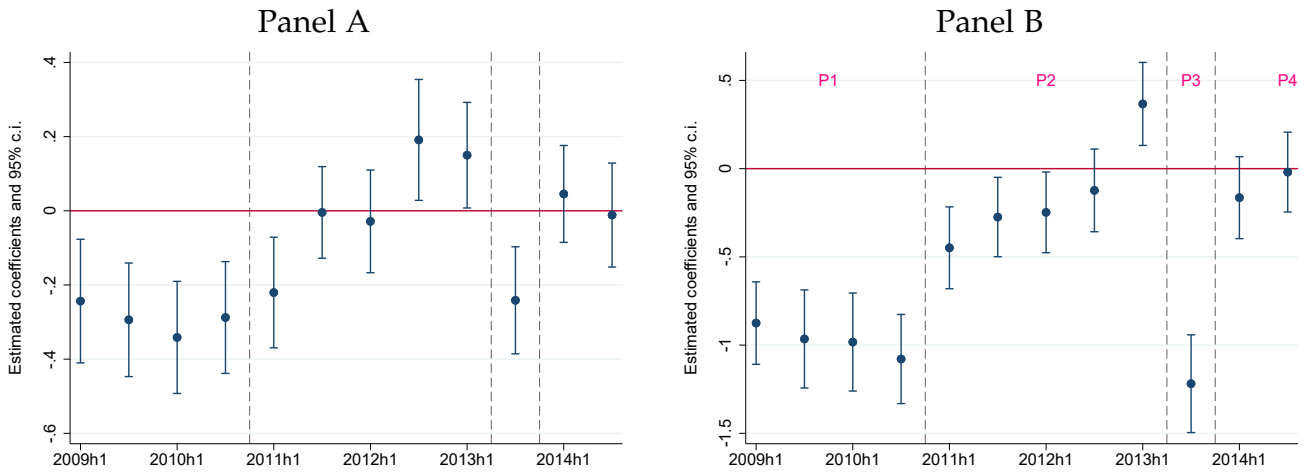

Figure A.4: The effect of the tax change on sales by investors

Sources. Data on the housing stock are from 76 municipalities' property tax files. Data on sales are from the Israel Tax Authority: Karmen Database.

Notes. The figure displays β coefficients (and 95% confidence intervals) from equation (1). In panel A, the estimation is limited to treated investors who sold another housing unit in the past 3-4 years and control investors who sold another housing unit in the past 4-5 years. Panel B further restricts the estimation to 3.5-4 and 4-4.5 years, respectively.

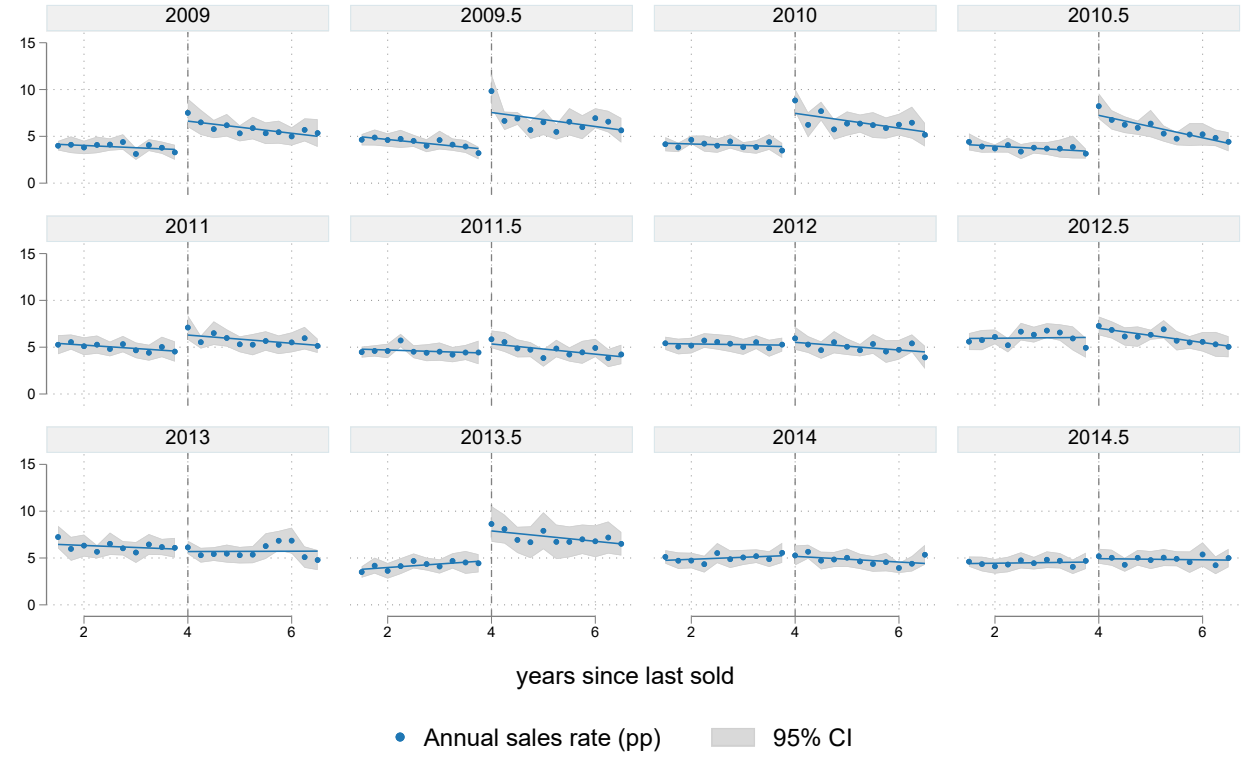

Figure A.5: Discontinuity by half year

Notes. Blue dots represent the semiannual sales rate of housing units owned by investors in each quarter-year bin defined by the time since the last sales by the owner. Lines represent second-degree polynomials fitted separately below (red) and above (green) 4 years. Rates are calculated at a quarterly frequency then averaged over the half-year period, 2009H1-2014H2. Each row represents a treatment period (P1)-(P4). See text for details. *Sources*. Data on the housing stock are from 76 municipalities' property tax files. Data on sales are from the Israel Tax Authority.

Table A1: The effect of investors' predicted sales on house prices

	all units	$\leq$ 3 rooms	$>$ 3 rooms	new units	resale units
	(1)	(2)	(3)	(4)	(5)
Predicted sales	-0.34***	-0.41***	-0.20*	0.18	-0.33***
by investors (/100)	(0.07)	(0.08)	(0.08)	(0.12)	(0.07)
Unit char.	✓	✓	✓	✓	✓
Half year FE	✓	✓	✓	✓	✓
Local market FE	✓	✓	✓	✓	✓
R^2	0.75	0.72	0.69	0.75	0.74
Observations	399,081	179,974	219,107	80,258	318,823

Notes. The table presents OLS regression results of estimating Equation (2), where transacted price is the dependent variable. Column (1) includes all housing units transacted between 2009 and 2014. Columns (2)-(5) use different subsets of transacted housing units; "$\leq$ 3 rooms" are units with 1.5-3 rooms; "$>$ 3 rooms" are units with 3.5-5 rooms; "new units" are recently built units sold by companies, and "resale units" are non-new units. "Predicted sales by investors" is the number of housing units predicted to be sold by investors (divided by 100) in each local market and half year. All regressions include unit characteristics (rooms, area and building age), half-year and local market fixed-effects. Standard errors are clustered by local market and shown in parentheses. The results suggest that additional sales of 100 units decrease prices by 34%. The decrease is larger for smaller housing units (column 2) and for resale units (column 5). For new units, we do not find evidence for a change in house prices due to increase in sales by investors.

Table A2: The effect of investors' predicted sales on rent

	(1) All units	(2) new leases	(3) new leases <3 rooms	(4) new leases >3 rooms	(5) extended leases
Predicted sales	0.02	0.03	0.15*	-0.19*	0.00
by investors (/100)	(0.04)	(0.06)	(0.08)	(0.12)	(0.04)
Unit characteristics	✓	✓	✓	✓	✓
Half year FE	✓	✓	✓	✓	✓
Local market FE	✓	✓	✓	✓	✓
R^2	0.75	0.77	0.74	0.75	0.74
Observations	64817	19436	11872	7564	45381

Notes. The table presents OLS regression results estimating equation (2) using rent as the dependent variable. Column (1) includes all rental units in the survey. Columns (2)-(5) utilize the panel structure of the survey as tenants are repeatedly surveyed periodically until they leave the rental unit. New leases (column 2) are leases that appear for the first time and extended leases (column 5) are subsequent appearances of the same rental units and tenants. Columns (3)-(4) are based on a subset of new leases; "$\leq$ 3 rooms" are units with 1.5-3 rooms, "$>$ 3 rooms" are units with 3.5-5 rooms. "Predicted sales by investors" is the number of housing units expected to be sold by investors (divided by 100) in each local market and half year. All regressions include the number of rooms, half year fixed-effects and local market fixed-effects. Standard errors clustered by local market, are in parentheses. The results suggest that rent at new leases of small housing units increased significantly in markets where treated investors are predicted to sell more units.