Supporting Young Homebuyers

Max Löffler Jonas Wogh

December 2025

Many countries support young homebuyers financially, but the effects of such targeted policies on homeownership, efficiency, and welfare are not well understood. We study a high-profile transfer tax exemption for young homebuyers introduced in the Netherlands in 2021. The reform offers quasi-experimental variation across time, buyers, and housing units, which we exploit using multiple identification strategies and high-quality administrative data. The policy substantially increased housing transactions, but we can exclude even modest effects on overall homeownership. Instead, we document large timing responses, with bunching around key policy dates and the age threshold. We also show that more than half of the exemption was capitalized into house prices, benefitting sellers rather than buyers. To assess the policy's welfare implications, we use our reduced-form estimates in a simple assignment model with indivisible housing. While the exemption increased economic surplus, the gains accrued primarily to upper-middle-class households, raising doubts about the policy's effectiveness as a tool for improving social welfare.

JEL Codes: R21, R28, H22, H31

Keywords: Housing Markets, Homeownership, Transfer Taxes, Welfare

* Löffler: Maastricht University, CESifo, IZA, and ZEW (email: m.loeffler@maastrichtuniversity.nl). Wogh: Maastricht University (email: j.wogh@maastrichtuniversity.nl). We thank seminar and conference participants at Cambridge, Helsinki GSE, FIT Tampere, Tinbergen, UEA, VfS-Regional, and ZEW for valuable comments. Löffler acknowledges funding from the Dutch Research Council (NWO, file number VI.Veni.221E.037).

1. Introduction

In most developed economies, governments support young or first-time buyers in the housing market. For instance, the US introduced the First-Time Homebuyer Credit in the aftermath of the Great Recession and today, more than 2,000 down payment assistance programs exist at the subnational level, typically targeting first-time buyers.[1] In the UK, first-time buyers are offered a substantial discount, which is repaid only when the property is later sold. In France, zero-interest loans cover up to half of a property's total value for first-time buyers. In Spain, young buyers are attracted to specific locations through outright cash grants. While these programs differ in design, their common aim is to promote homeownership by supporting young households who often face binding borrowing constraints.

Evaluating the effects of young homebuyer support schemes on homeownership, economic efficiency, and welfare is difficult, as it requires rich population-wide microdata and exogenous treatment variation. In this paper, we address these challenges and provide novel evidence on the impact of targeted young-buyer support programs using linked register data and quasi-experimental variation from a recent policy change in the Netherlands. The Dutch transfer tax reform in 2021 introduced an exemption for homebuyers under the age of 35. After three months, a price limit was added to the exemption. This reform offers an ideal laboratory, as it generates treatment variation over time, across buyers, and across housing units.

Using several complementary econometric approaches, we present four main findings. First, the starter tax exemption stimulated home purchases among young individuals. The effect was strongest during the initial quarter, but remained positive in the medium run after the introduction of the price limit. Second, we find no evidence for a higher homeownership rate. Our estimates exclude even modest increases. Many of the additional purchases were made by repeat buyers trading up, rather than by first-time buyers entering the market. We also find substantial timing responses. Third, we estimate that 65–90% of the tax relief is capitalized in higher prices and thus benefits sellers rather than buyers. Fourth, the policy mitigates existing distortions and raises economic surplus, with the most significant gains accruing to upper-middle-income households. This distributional pattern, however, raises doubts about the positive impact on social welfare.

We start the analysis by providing evidence on the causal effect of the starter tax exemption on home purchases. We do so in a difference-in-differences setup, comparing the purchase behavior of young, treated individuals with that of slightly older cohorts who were never eligible. Both groups move in parallel before the policy and diverge sharply just after the announcement. Transactions by young buyers doubled in the first quarter, when no price limit applied. In the medium run, the policy raised purchases

[1] Warden and Kadzielawski (2025) provide an overview of such programs at the state, regional, and local level in the US. Most of them offer second mortgages with preferential loan terms.

among younger cohorts by around 13% relative to pre-reform levels. Short-run responses are driven by individuals with higher liquid assets and incomes, whereas those with lower assets and incomes drive medium-run effects. Comparing the time trends across market segments offers little evidence for the crowding-out of older buyers.

However, an increase in purchases does not necessarily translate into higher homeownership rates. A first indication is that the majority of additional purchases generated by the policy were made by existing homeowners deciding to trade up. We move on by investigating various forms of timing responses to the policy. The specific design of the tax exemption creates strong incentives to optimize home purchases around cutoff dates—such as the policy introduction in January, the end of the exemption period in March 2021, and the 35th birthday. Dynamic treatment estimates suggest that approximately 60% of the spike in transactions in the initial three-month period is attributable to time bunching around policy cutoff dates. We complement the evaluation with a regression discontinuity analysis of home purchases around the 35th birthday. While behavior is smooth around the cutoff before the policy, a sharp discontinuity emerges after its implementation. Most of this discontinuity is driven by individuals just around their 35th birthday. Using a donut approach that excludes those months reduces the estimated discontinuity by about 80%, rendering it only marginally significant.

We then turn to the Dutch homeownership register to directly test differences in ownership rather than purchases. To do so, we estimate discontinuities in the homeownership probability using birth-year-month cohorts as the running variable. This analysis indicates a short-lived effect: Individuals born in January 1986 or later postponed home purchases in 2020 and see higher ownership rates after the first year of the policy. Nevertheless, we cannot rule out that the discontinuity equals its pre-reform level. In 2023, two years into the policy, the discontinuity returned to its pre-reform level and is no longer significantly different from zero. Overall, the timing analyses leave little room for extensive-margin ownership responses to the starter tax exemption, as the estimated confidence intervals exclude increases in the homeownership rate larger than 0.6 pp.

On the intensive margin, we find strong responses to the incentives created in the different phases of the policy. Eligible homebuyers used the absence of a price limit in the first three months to buy significantly more expensive units in higher-income neighborhoods. After the introduction of the price limit, this effect reverses, as treated buyers opt for significantly lower-quality units to remain eligible.

Finally, we investigate the policy's price effects. Theory predicts that the tax exemption will partly capitalize into higher prices unless housing supply is perfectly elastic. Using an auxiliary difference-in-difference analysis comparing units below versus above the price limit, we find that qualifying units experience stronger medium-run price increases. Estimates imply that 65–90 percent of the tax relief is capitalized into prices and benefits sellers rather than buyers.

Equipped with the reduced-form evidence, we set up a simple assignment model to

quantify the utility consequences of financial support for young homebuyers at the individual level (as in Määttänen and Terviö, 2022). In the model, individuals maximize utility by consuming a composite good and housing services from indivisible units. Real-estate transactions are subject to taxation, but the government offers an ad valorem subsidy to some homebuyers. We then take a sufficient-statistics approach (Chetty, 2009, Kleven, 2021): We assume the economy is in equilibrium and study the impact of small changes in the purchase subsidy on household utility. This perturbation approach shows that utility consequences are governed by three sufficient statistics, which we estimated before: (i) the effect of the subsidy on house prices, (ii) the policy's impact on home purchases, and (iii) its impact on tax revenues. Behavioral responses to the subsidy, in contrast, have no first-order welfare effect because of the envelope theorem. We argue that the underlying assumptions are likely to be met, given that we find no evidence of extensive margin effects on homeownership.

In the final part of the analysis, we quantify the utility consequences for the entire Dutch adult population. Based on their observed housing market activity, we classify individuals into different subpopulations—distinguishing between young and old, eligible and ineligible, marginal and inframarginal buyers and sellers, as well as never takers. Relying on the quasi-experimental estimates, we then calculate money-metric utility changes because of the starter tax exemption. We show that the policy's effect on economic surplus is positive, though small in magnitude. This positive impact on efficiency reflects that the policy is effective at mitigating existing distortions due to the real-estate transfer tax. However, the policy redistributes resources toward the upper middle class, at the expense of those at the bottom and the top of the income distribution. This is because many eligible homebuyers (i.e., those under 35 buying units below the price limit) and many sellers are from this group. In contrast, there are few homebuyers in the lowest income quartile, even with the tax exemption. Households in the highest income quartile typically buy units above the price limit and therefore rarely qualify for the tax exemption. Overall, this distributional pattern makes it unlikely that the starter tax exemption improved social welfare.

Related Literature. Even though support schemes for young homebuyers have received broad attention from policymakers, there is little evidence on their overall economic effects. Most existing studies focus on the impact of individual schemes on their respective housing markets. Berger, Turner, and Zwick (2020) study the First-Time Homebuyer Credit program, which was implemented as a temporary stimulus for the US housing market after the Global Financial Crisis, and document its success in promoting purchase activity, ownership transitions, and aggregate prices. Hembre (2018) confirms the success of the policy in stimulating home purchases but contrasts it with a high fiscal cost, leaving the net welfare effects unresolved. Goodwin and Zumpano (2011) focus on the distributional aspect of the policy and highlight that it was particularly successful

in promoting homeownership among low-income and minority groups. Gobillon and Le Blanc (2008) study the effect of zero-interest loans for first-time homebuyers in France, based on a tenure choice model. They document that while the subsidy is successful in promoting home purchases, it involves a large windfall effect, as most beneficiaries only bring forward their already planned purchase. Carozzi, Hilber, and Yu (2024) study the UK Help to Buy scheme, which extended government loans to first-time and repeat buyers, with a focus on prices and supply. In line with theory, they find that the policy promoted construction only in areas with elastic supply, whereas it further drove up prices in supply-constrained urban areas. We add to this existing work by bringing into focus the welfare effects of supporting young homebuyers. Similar to previous studies, we find that such policies do promote housing market activity. However, based on detailed administrative data, we establish that the increase in activity does not translate into higher homeownership rates and that the policy mainly redistributes resources along the income distribution.

Studies examining outright home purchase support schemes are related to a broader literature on fiscal subsidies for homeownership. Two aspects of this are particularly relevant. On the one hand, tax exemptions for imputed rental income, which Poterba and Sinai (2008) identify as a substantial subsidy for homeowners. On the other hand, tax deductions for mortgage interest payments have received more scholarly attention. Glaeser and Shapiro (2003) argue that the huge cost of this subsidy cannot be justified, given that it merely incentivizes owners to consume more housing, which entails no positive externality. Hilber and Turner (2014) show that the mortgage interest deduction can even have an adverse effect on homeownership to the extent that it is capitalized into higher aggregate prices. Using administrative data from Denmark, Gruber, Jensen, and Kleven (2021) confirm these concerns and show that the mortgage interest deduction operates only at the intensive margin, by inducing buyers to purchase more expensive homes. In contrast, it has no effect on aggregate ownership rates. One study that attempts to incorporate different types of housing market subsidies in a general equilibrium framework is Floetotto, Kirker, and Stroebel (2016). They argue that home purchase subsidies have a negative welfare effect for the majority of the population, who finance the program without benefiting directly. Moreover, in agreement with the previous literature, their findings suggest that removing tax benefits of home ownership, including the mortgage interest deduction, would be welfare-enhancing in the long run.

Another stream of the literature examines the effects of real-estate transaction taxes. Most studies in this field rely on the analysis of notches, i.e., discontinuous increases in the tax rate above a given sales price. Kopczuk and Munroe (2015) study the effect of the "Mansion Tax" in New York and document a substantial bunching response that leads to the local unraveling of the market: around the notch, transactions that would yield positive surpluses for buyers and sellers cease to occur. The distortionary effect of transfer taxes is emphasized by Best and Kleven (2018), who study their temporary

removal in the UK housing market. They find that the elimination of transfer taxes has a profound effect on transaction activity in the short run and, by extension, on consumer spending due to housing-related expenditures. Moreover, they show that this effect is only partially reversed after the tax is reintroduced, suggesting that transfer taxes may reduce ownership transitions at the extensive margin. In contrast, Slemrod, Weber, and Shan (2017) find that an increase in the transfer tax did not reduce transaction volumes for home values just above the notch. They therefore argue that transfer taxes have little long-run effect on housing market activity and welfare. Finally, Dolls et al. (2025) study the effect of transfer taxes on aggregate prices in a difference-in-differences framework, exploiting variation in tax changes across German states over time. They estimate that a one-percent increase in transfer taxes lowers house prices by 3% on average. This effect is more pronounced for apartments, which have short periods and hence incur tax payments more frequently. We contribute to this literature by shedding light on the distributional effects of transfer tax exemptions. In addition, we complement the analysis of purchasing activity by investigating various intensive-margin responses.

The remainder of this paper is organized as follows. Section 2 details the institutional background. Section 3 introduces the data used. Section 4 presents the empirical results. Section 5 introduces the theoretical framework to evaluate the utility consequences of subsidies for young homebuyers. Section 6 simulates welfare consequences, based on the reduced-form empirical parameters. Section 7 concludes.

2. Institutional Setting

2.1. Background

Homeownership rates in the Netherlands, at around 60%, are close to the European average. One defining institutional feature is the well-developed mortgage market, which supports broad availability. Dutch mortgage debt levels are the highest in Europe (DNB, 2021), in part due to historically generous loan-to-value (LTV) limits. While lending standards were somewhat tightened in the aftermath of the global financial crisis, the statutory LTV limit has remained at 100%. In theory, this makes homeownership more accessible for young, liquidity-constrained households than in other countries with binding down payment requirements. However, tight limits on loan-to-income (LTI) ratios constrain the borrowing capacity of many prospective homebuyers. Since these limits are based on current, rather than permanent income, they are particularly binding for young households with short employment histories. As in many other advanced economies, the sharp increase in house prices over the past decade has therefore put homeownership out of reach for many young households. Over the decade from 2010 to 2020, the average age of first-time homebuyers has increased by more than a year, from 28.6 to 29.8 years (Tweede Kamer, 2020).

In response to these trends, the Dutch government has introduced a range of housing

market interventions aimed—directly or indirectly—at supporting young and first-time buyers. These include tax exemptions for inter-generational wealth transfers related to housing (Bos, Kok, and Wogh, 2024), mortgage assistance schemes, and restrictions on buy-to-let transactions in certain municipalities (Francke et al., 2023).

2.2. The Starter Tax Exemption of 2021

In September 2020, the Dutch parliament passed a real estate transfer tax exemption for homebuyers under the age of 35 (in Dutch: *startersvrijstelling*). Despite its name, the starter tax exemption applied to both first-time and repeat buyers. Older buyers continued to pay the statutory two percent tax rate. The policy was first mentioned in public in July 2020 and finally implemented as of January 2021, shortly before the Dutch housing market reached its peak. The tax exemption was granted only if the acquired unit is used for owner-occupied living, whereas the applicable tax rate for buy-to-let transactions increased concurrently from six to eight percent, irrespective of the buyer's age. If individuals made a joint purchase, the exemption applied in proportion to ownership shares. For instance, a co-purchasing couple with one partner under 35 and one partner above 35 paid 1% of the price in tax if ownership was split equally.

In November 2020, prior to the policy's introduction, legislators announced a price limit for qualifying transactions that would later lead to a notch in the tax code. As of April 2021, three months after the policy took effect, only units trading below €400,000 qualified for the tax exemption, while units priced above this limit were again subject to the regular two percent tax rate, regardless of the buyer's age. In January 2023, the price limit was raised to €440,000, to reflect higher property prices.

The Dutch starter tax exemption policy shares many features with other support schemes for young buyers worldwide, such as outright cash grants, tax credits, or down payment assistance. While these policies differ in their specific designs, they all subsidize young buyers by reducing upfront fixed costs, intending to facilitate earlier access to homeownership. These targeted subsidies usually mitigate existing market distortions, like those caused by real-estate transaction taxes, and may thus increase surplus.

3. Data

Our main data source is anonymized administrative records on the Dutch population provided by Statistics Netherlands (CBS). These data include demographic characteristics, residential location, and homeownership status for every person living in the Netherlands. One limitation of the data is that birth dates are often missing for people born outside the country. As we heavily rely on birth-date information, we restrict the sample to individuals born in the Netherlands. The CBS data also provide various household-level income and balance sheet variables, which we utilize in heterogeneity analyses to group households into subsamples of the population. In addition to the

individual-level data, the CBS records track various characteristics of the residential housing stock, including first and second owners, as well as tax-assessed values (*WOZ* values).[2]

We match the CBS records with two housing transaction datasets. First, we use data from the Dutch Land Registry (*Kadaster*), which covers the universe of existing home sales in the Netherlands but includes only a limited set of variables. Because the identities of buyers and sellers are not directly recorded in the transaction data, we impute home purchases from changes in property ownership records. Since these records capture both the first and second owners of a property, we identify all purchases made individually or with a partner. Second, we augment this data with information from the country's largest association of real estate brokers, NVM. This auxiliary data covers only around 60% of all transactions, but provides more detailed information on each transaction, including property characteristics and listing prices.

Main Sample. Combining these sources, we construct a person-by-month panel dataset spanning the period from January 2019 to December 2023. We track several key variables for Dutch-born adults, most importantly, their home purchases and ownership, as well as transaction prices, tax-assessed values, and characteristics of purchased units. Throughout the analyses, we restrict the sample to individuals near the policy cutoff at age 35, dropping individuals under 30 and those 43 or older. As a result, we work with a semi-balanced panel in which individuals enter and exit the sample over time as they approach the lower or pass the upper age thresholds. Each month, the sample covers around 1.7 million individuals.

4. Reduced-Form Evidence

In this section, we provide empirical evidence on the effects of young homebuyer support policies, using quasi-exogenous variation from the Dutch starter tax exemption. In Section 4.1, we first analyze the policy's impact on home purchase behavior. Section 4.2 presents regression discontinuity analyses of timing responses. We evaluate intensive-margin responses and bidding behavior in Section 4.3. Finally, we study the policy's effect on aggregate house prices in Section 4.4.

4.1. Home Purchase Behavior

We begin the empirical investigation by analyzing home purchases of young individuals around the introduction of the starter tax exemption. Each month before the reform, we observe around 4,500 home purchases in our baseline sample. This corresponds to an individual monthly purchase probability of 0.28%. To distinguish the causal effect

[2] The Dutch ownership register reported only the first owner until 2018. We therefore perform our analyses only from 2019 onwards to ensure a consistent ownership register over time.

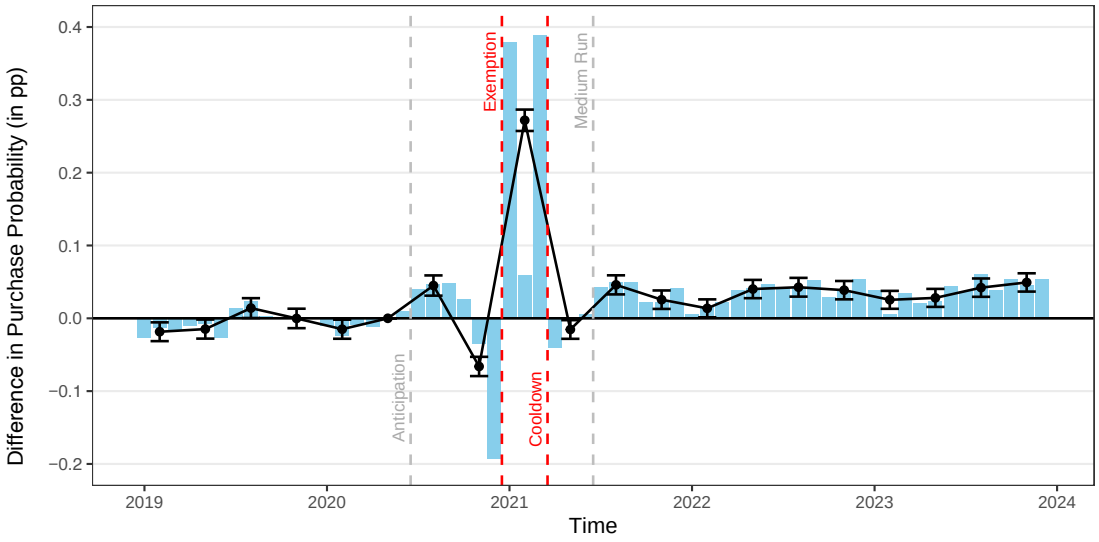


Figure 1: The Effect of the *Starter Tax Exemption* on Home Purchases

Notes: This figure shows the difference in home purchase probabilities (in percentage points) between individuals aged 30–34 and those aged 38–42 between January 2019 and December 2023. Blue bars show raw, monthly differences in the share of homebuyers among both groups. Black dots depict quarterly differences in the share of homebuyers estimated based on Equation (1). Differences are normalized to their 2020Q2 level, before the announcement of the policy. Black vertical bars indicate 95% confidence bounds. Raw time series of the monthly home purchasing behavior in both groups and depicted regression results for each quarter can be found in Appendix A.1.

of the policy from overall housing market trends, we compare eligible and ineligible individuals in a difference-in-differences framework. Concretely, individuals aged 30–34 form the treatment group and those aged 38–42 form the control group.[3] The identifying assumption is that the home purchase behavior of both groups would move in parallel, absent the introduction of the tax exemption for young buyers.

In Figure 1, we plot the difference in monthly purchase probabilities between the two groups (blue bars), normalized to zero in the second quarter of 2020.[4] Despite some differences between the two groups, we find no evidence of systematic deviations from parallel trends prior to the policy change being first discussed. In 2020Q3, just around the official announcement of the policy change, there was increased activity by younger buyers. We attribute this temporary divergence to the market turmoil around the COVID-19 pandemic, rather than to a causal anticipation effect.[5] In 2020Q4, immediately before the policy was implemented, individuals under 35 signed substantially

[3] We exclude individuals aged 35–37 from the analysis as individuals in that age range switch from the treatment to the control group during our study period. Consider, e.g., an individual turning 35 in January 2023. This person was treated for two years, but is no longer eligible from then onward, and would erroneously become part of the control group if we included individuals aged 35–37.

[4] Before the reform, the younger group is around 0.15 pp more likely to purchase a home in any given month. Appendix Figure A.1 depicts the raw time series of monthly home purchases for both groups.

[5] If buyers were anticipating the tax exemption, we would expect a decline in purchases, as observed in 2020Q4. Moreover, the transaction date in our data reflects the signing of the notarial deed, which typically occurs weeks after negotiations. It is therefore unlikely that the announcement in Q3 affects purchases within the same quarter. The uptick in young buyers' purchases during the early months of the pandemic is also consistent with patterns observed in other countries (see, e.g., Appendix Figure A.2 for corresponding evidence based on US data).

fewer contracts. In turn, we observe a strong spike in purchases in 2021Q1, when the starter tax exemption was enacted but did not yet include the €400,000 price limit. After a cooldown period in 2021Q2, the purchase volume of younger individuals remained systematically elevated until the end of our study period in December 2023.

We formalize this graphical comparison by estimating a dynamic difference-in-differences model:

$$Y_{it} = \sum_{k \neq 20Q2} \beta_k \cdot Treat_i \cdot \mathbb{1}_{(t=k)} + \gamma_t + \delta \cdot Treat_i + \varepsilon_{it}, \tag{1}$$

where the dependent variable is an indicator equal to one if an individual purchased a home in a given period, and zero otherwise. $Treat_i$ is the binary indicator for those aged 30–34 in a given month and γ_t is a period fixed effect. The treatment estimates $\widehat{\beta}_k$ capture the average difference in monthly home purchase probabilities between younger and older individuals, relative to 2020Q2, the omitted reference period. Standard errors are clustered at the individual level. The resulting estimates are shown in Figure 1 (in black). To summarize the policy's impact, we further aggregate the estimates for the five sub-periods indicated in the figure. Column (1) of Table 1 presents the resulting coefficients. As the figure suggests, pre-trends are largely flat before the policy announcement. After the announcement, but prior to implementation, the average monthly purchase probability of eligible buyers declines by 0.011 pp. In the exemption period—during which no price limit is in place—purchases by young buyers rise sharply by 0.272 pp. The purchase probability thus doubles during this period. After a short cooldown period in 2021Q2, we estimate a medium-run effect of the policy of around 0.035 pp. This translates into a 12.5% increase, relative to the reference period.

Heterogeneous Effects. Financial relief is likely most beneficial to those young homebuyers facing a binding borrowing constraint. Some countries explicitly target this group by offering support only to low-income homebuyers. In the Netherlands, loan-to-income limits are strict, such that individuals with low incomes are most likely to be constrained. Similarly, although Dutch loan-to-value limits are relatively generous, we expect those with few liquid assets to be constrained, as transfer taxes and other transaction costs must be paid out of pocket. To investigate heterogeneity in treatment effects across these dimensions, we identify each individual's disposable household income and bank balance, respectively, and classify individuals into quartiles based on the 2019 distribution of both measures. We then re-estimate Equation (1) via sample splits for each population quartile.

The resulting estimates are reported in columns (2)–(5) of Table 1. We first note that pre-trends become somewhat unstable when cutting along these margins.[6] After

[6] Note that, despite these precisely estimated small deviations before the reform, the quarter-by-quarter estimates exhibit clear trendbreaks and sharp changes in purchase behavior exactly around the policy thresholds. We relegate event study plots and further regression results, as well as details on the

Table 1: The Effect of the *Starter Tax Exemption* on Home Purchases

		By Liquid Assets		By Income		By Ownership	
	Base-line (1)	First Quartile (2)	Fourth Quartile (3)	First Quartile (4)	Fourth Quartile (5)	Non-owner (6)	Exist. Owner (7)
Treated ×							
Pre-Policy Period (19Q1–20Q1)	−0.007 (0.005)	−0.017** (0.008)	0.027** (0.012)	−0.031*** (0.008)	−0.002 (0.012)	−0.006 (0.007)	−0.006 (0.008)
Anticipation Period (20Q3–20Q4)	−0.011* (0.006)	−0.006 (0.010)	−0.036*** (0.013)	−0.001 (0.009)	−0.020 (0.013)	−0.018** (0.009)	−0.001 (0.008)
Exemption Period (21Q1)	0.272*** (0.007)	0.133*** (0.012)	0.350*** (0.016)	0.175*** (0.011)	0.401*** (0.017)	0.294*** (0.011)	0.250*** (0.010)
Cooldown Period (21Q2)	−0.015** (0.007)	0.019* (0.011)	−0.069*** (0.014)	0.038*** (0.010)	−0.105*** (0.014)	−0.007 (0.010)	−0.048*** (0.008)
Medium Run (21Q3–23Q4)	0.035*** (0.005)	0.043*** (0.008)	0.007 (0.011)	0.064*** (0.007)	0.010 (0.011)	0.016** (0.008)	0.038*** (0.007)
Person-by-Quarter Observations	33.3mn	8.3mn	8.3mn	8.3mn	8.3mn	13.9mn	17.8mn

Notes: This table summarizes the estimated effect of the starter tax exemption on the monthly purchase probability among individuals aged 30–34. Coefficients are to be interpreted as percentage point changes relative to the omitted quarter 2020Q2. Estimated treatment effects $\widehat{\beta}_k$ are based on the difference-in-differences model specified in Equation (1). The estimation is based on a person-by-month panel dataset from January 2019 until December 2023. Standard errors, clustered at the individual level, are reported in parentheses (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

the introduction of the policy, two patterns stand out. On the one hand, the spike in purchases during the exemption period (2021Q1) is most pronounced for individuals with higher assets and income. On the other hand, we observe a lasting medium-run increase in purchases only among those with fewer financial resources. Taken together, these findings suggest that price limits—a common feature of many young-buyer support policies around the world—can be effective in targeting financial support to the most-constrained individuals.

Finally, despite its name, the Dutch starter tax exemption was not restricted to first-time homebuyers. Therefore, next to promoting access for aspiring homeowners, it offered existing owners the chance to trade up. We examine the policy's impact on these two groups separately, by splitting the sample into individuals who did already own a property in the previous quarter (existing owners) and those who did not (non-owners). The resulting coefficients are reported in columns (6) and (7) of Table 1. Strikingly, we find that the majority of the policy's medium-run effect is driven by existing owners making another purchase. Whereas the monthly volume of first-time purchases increased by only 1.6%, repeat purchases rose by 3.8%. This suggests that the policy gave a strong incentive for existing owners to trade up. In Appendix A.2, we analyze the mechanism behind this more closely.

Crowding-Out of Older Homebuyers. A potential concern with our identification strategy is that individuals from both the treatment and the control group may compete

temporal stability of these population quartiles relative to their 2019 distributions, to Appendix A.2.

for similar homes. Subsidizing the younger group could then lead to a crowding out of older homebuyers from the market if supply is sufficiently inelastic. In that case, the estimated treatment effect ($\widehat{\beta}_k$) would be biased upward. We address this concern by exploiting cross-sectional policy variation stemming from the €400,000 price limit that was introduced in 2021Q2, regardless of the buyer's age. As the tax exemption applied only to units in the lower segments of the market, potential crowding out would be expected to be most pronounced in those segments as well.

To investigate this prediction, we rerun the difference-in-differences analysis from Equation (1) but replace the dependent variable Y_{it} for any purchase with indicators Y_{it}^s that are equal to one only if an individual purchased a home in a given price segment s. We assign segments by splitting all transacted units into two groups based on their 2019 (i.e., pre-policy) tax-assessed value. We define units with tax values below €250,000 as low-value and those with higher tax values as high-value. This value corresponds roughly to the median tax value in 2019, and therefore provides a useful ex ante proxy for the policy's €400,000 price limit. Whereas 87% of low value units sold below the price limit—and therefore qualified for the tax exemption—only 19% of high value units did. The distinction between low and high value units will be revisited in Section 4.4.

In Figure 2, we compare the policy's effect on purchases between low- and high-value units (see Appendix A.3 for detailed results). In contrast to the baseline results, we visualize not only the estimated treatment effect (in black), but also quarterly fixed effects (in grey, normalized to 2020Q2), which capture average purchase probabilities among older individuals. This is central to our test of crowding out: if the subsidy for young homebuyers has a direct negative effect on older ones, then their purchase activity should decline most in the lower-price segment, where the policy had the most bite. However, the results do not support this. Whereas purchases by older buyers decrease sharply among higher-value units (right panel), the decrease is much more muted among lower-value units (left panel). These patterns suggest that the primary effect of the exemption was to increase purchases among eligible buyers, rather than to displace purchases by ineligible buyers.

4.2. Extensive Margin vs. Timing Responses

The difference-in-differences analysis reveals that the starter tax exemption causally increased home purchases among eligible individuals. A central question that follows from this is whether the additional transactions correspond to an extensive margin response—and therefore a lasting increase in homeownership rates—or merely reflect timing responses. Making this distinction is crucial for policy and welfare implications. For individuals who would have only bought slightly later (or earlier) without the tax exemption, the main benefit is a financial windfall. In contrast, those who would never have purchased absent the policy—or only much later—may gain additional utility from higher housing consumption and homeownership.

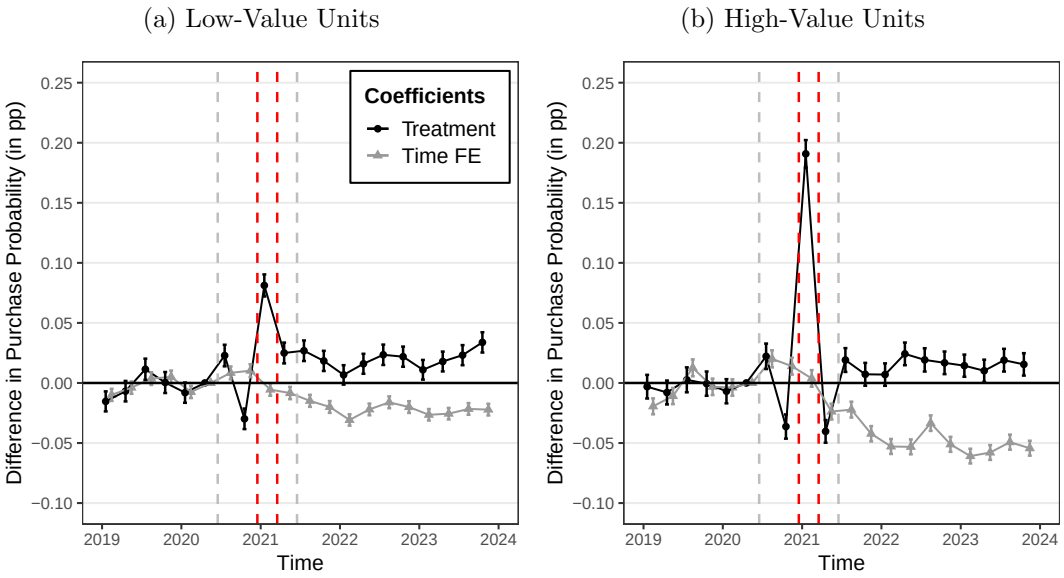


Figure 2: Testing Crowding-Out by Comparing Market Segments

Notes: This figure shows differences in home purchase probabilities separately for low vs. high value units, as defined by their 2019 tax assessed value. The coefficients are based on equation (1). However, in contrast to Figure 1, we here also visualize time fixed effects (in grey), normalized to their 2020Q2 level. These capture the purchase behavior of individuals in the control group, i.e., those aged between 38-42. Vertical bars indicate 95% confidence bounds.

The first piece of evidence on extensive margin responses comes from the split between first-time and repeat purchases. As documented in Table 1, the majority of additional transactions due to the policy involve existing homeowners, which by definition cannot contribute to higher ownership rates.

In addition, we investigate the reoptimization of already planned home purchases around policy cutoffs. As discussed in Section 4.1, purchases by young buyers declined just before the introduction of the tax exemption and rose sharply immediately afterwards, suggesting substantial bunching in the initial three months when no price limit applied. A back-of-the-envelope calculation of the monthly differences in purchases between November 2020 and June 2021 suggests that at least 57 percent of the strong policy response in the initial period is attributable to timing effects within this narrow eight-month window.[7] It is likely that some individuals brought forward their home purchases by even longer periods, given that the end of the exemption period and the introduction of the price limit in April 2021 were known well in advance.

Bunching Around Birthdays. Another cutoff date that young buyers may use to optimize the timing of their purchases is their 35th birthday, which marks the eligibility threshold. We investigate this optimization strategy from a regression-discontinuity perspective on the policy. Figure 3 plots the monthly purchase probabilities for people aged between 32.5 and 37.5, i.e., thirty months on either side of the age threshold. After the

[7] The details of this calculation can be found in Appendix A.4, Table A.2.

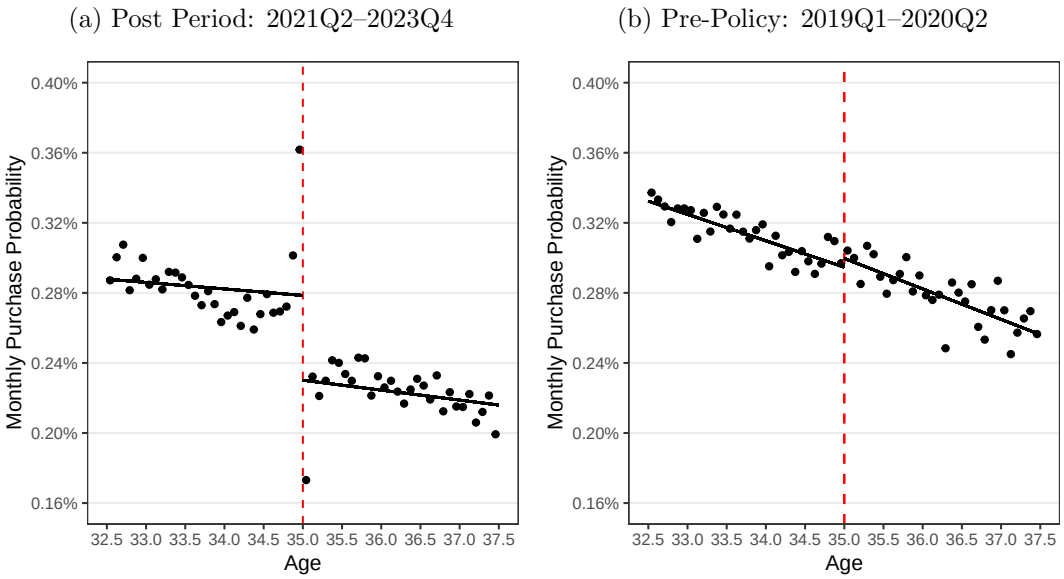


Figure 3: Regression Discontinuity Evidence on the *Starter Tax Exemption*

Notes: This figure shows the average monthly home purchase probability by age in months during our sample period. Panel (a) is based on the period from 2021Q2 until 2023Q4. Panel (b) shows pre-policy evidence for the period from 2010Q1 until 2020Q2. The solid black lines represent estimates from a linear RD model as specified in Equation (2). Corresponding estimation results are provided in Table 2, panel A. Complementary visualizations for regression-discontinuity models that exclude specific age-in-month groups around the threshold are presented in Appendix Figure A.8 ("donut RDs").

subsidy is introduced and the price limit is in place, we find a sharp drop in purchases when turning 35 (panel A), whereas purchases were smooth before the policy (panel B). We complement the graphical evidence with a linear regression discontinuity approach, estimating the following model on the pooled individual-by-month panel, separately for the periods before and after the policy introduction:

$$Y_{it} = \alpha + \beta \cdot \mathbb{1}(Age_{it} < 35) + \gamma \cdot Age_{it}^{(35)} + \delta \cdot \mathbb{1}(Age_{it} < 35) \cdot Age_{it}^{(35)} + \varepsilon_{it}, \qquad (2)$$

where Y_{it} is a purchase indicator and $Age_{it}^{(35)}$ denotes an individual's current age, recentered around their 35th birthday. The coefficient of interest, β, captures the discontinuity in purchasing behavior at the threshold. Estimation results are reported in panel A of Table 2. The baseline model, reported in column (1), suggests that the purchase probability drops significantly by 0.048 pp when turning 35, an effect that is slightly larger than the medium-run difference-in-differences estimate of 0.035 pp.

However, as is evident from Figure 3, the discontinuity is largely driven by the purchasing behavior a few months around the 35th birthday.[8] We quantify the impact of retimed purchases in this narrow age range by implementing a set of donut RD models. Estimates decrease when excluding two or four months on either side of the cutoff. Excluding half a year on the left and on the right of the cutoff lowers the linear RD estimate to 0.01 pp with a t-statistic of only 1.76 (see Table 2, columns (2)–(4); corresponding

[8] More flexible functional forms in individuals' age would thus yield even larger RD estimates.

Table 2: Extensive Margin vs. Timing Responses to the *Starter Tax Exemption*

Panel A – Discontinuity in Purchases Around 35th Birthday

	2021Q2–2023Q4					
	Base-line	±2 Months Omitted	±4 Months Omitted	±6 Months Omitted	Exemption Period	Pre Policy
	(1)	(2)	(3)	(4)	(5)	(6)
Treated ($Age < 35$)	0.048***	0.021***	0.014***	0.010*	0.213***	–0.005
	(0.004)	(0.004)	(0.005)	(0.006)	(0.017)	(0.006)
Person-by-Month Obs.	28.9mn	26.2mn	24.3mn	22.5mn	2.5mn	14.6mn

Panel B – Discontinuity in Homeownership by Birthyear

	01.01. 2019	01.01. 2020	01.01. 2021	01.01. 2022	01.01. 2023	01.01. 2024
	(7)	(8)	(9)	(10)	(11)	(12)
Treated ($Birthyear \geq 1986$)	0.271	0.252	0.074	0.438**	0.234	0.195
	(0.215)	(0.214)	(0.212)	(0.210)	(0.208)	(0.207)
Person-by-Month Obs.	842k	842k	842k	842k	842k	842k

Notes: This table summarizes the estimated effect of the starter tax exemption on the purchase probabilities and homeownership rates of eligible individuals. Coefficients are to be interpreted as percentage point differences. Pre-policy is 2019Q2–2020Q2. Post-policy is 2021Q2–2023Q4. Exemption period is 2021Q1. In Panel A, estimated treatment effects are based on the regression discontinuity model specified in equation (2), and the running variable is an individual's current age, measured at monthly level of detail and recentered around the 35th birthday. In Panel B, the running variable is an individual's birth-year-month (MoB_i), re-centered around January 1986. The following model is estimated: $Owner_i^t = \alpha^t + \beta^t MoB_i^{(\text{Jan'86})} + \gamma^t T_i + \delta^t \left(T_i \times MoB_i^{(\text{Jan'86})}\right) + \varepsilon_i^t$. The "ever-treated" indicator T_i equals 1 for all individuals born after January 1, 1986 (i.e., those with $MoB_i^{(\text{Jan'86})} > 0$). The coefficients of interest, δ^t, capture the discontinuous change in homeownership at the eligibility cutoff, with t ranging from January 2019 up to January 2024. In both models, 30 month-cohorts are included on either side of the discontinuity. Standard errors, clustered at the individual level, are reported in parentheses (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

donut RD plots are relegated to Appendix Figure A.8). We conclude that around 80 percent of the policy's effect on transaction volumes is driven by short-lived retiming of purchases around the 35th birthday, rather than by additional extensive-margin transactions.

Moreover, these estimates likely still overestimate the additional transactions in response to the policy. An RD analysis of the tax exemption faces the inherent challenge that individuals move along the running variable over time. Each period, some individuals switch from the treatment to the control group as they turn 35 and cross the age threshold. The longer the policy has been active, the greater the number of individuals aged 35 or older who had been eligible before. Their home purchases may thus understate the counterfactual homebuying activities of older cohorts if they had never qualified for the policy. Although we find significant discontinuities in each quarter since policy implementation in 2021Q1, these estimates likely overstate the true discontinuity, especially in later quarters (see Appendix Figure A.9 for quarter-by-quarter RD graphs).

Differences in Homeownership. The preceding analyses suggest that the policy did little to promote longer-term homeownership. To examine this directly, we turn to the Dutch homeownership register, which records the owner(s) of each unit in the Netherlands as of January 1 of each year. If the tax exemption increased homeownership among treated individuals, we should observe lasting differences across birth-year-month cohorts. Individuals born in December 1985 or earlier were never eligible for the policy since they turned 35 before it was introduced. Everyone born afterwards was at least briefly eligible and may have benefited from the tax exemption.

We exploit this variation by estimating linear RD models with birth-year-month as the running variable in the annual snapshots of the homeownership register. Panel B of Table 2 shows the resulting estimates of the discontinuity at January 1986, i.e., when turning eligible. Before the policy, estimates are positive but insignificant. In January 2021, the discontinuity appears to decrease, which mirrors the observation that eligible individuals delayed their purchases to benefit from the subsidy. By January 2022, one year into the policy, we observe an increase in ownership among treated individuals, which—while significantly different from zero—is not significantly different from the pre-policy discontinuity. Moreover, the gap disappears by 2023, when the estimated effect returns to and stays at its pre-policy magnitude. Taken at face value, these estimates suggest a relatively precisely estimated null effect of the tax exemption on homeownership. The 95% confidence intervals around the estimates presented in columns (11) and (12) exclude discontinuities exceeding 0.6 pp relative to a baseline ownership rate of around 50%.

A very similar picture emerges when taking a longer-term perspective using administrative income records. Each year, the Statistical Office publishes household income records, including imputed rent from owner-occupied properties. We use this data from 2012 onward to calculate annual ownership rates for each birthyear cohort from 1981 to 1990. In line with the previous analysis, we find a small, temporary dip in 2021 and a corresponding uptick in 2022, but no long-term effect on ownership rates (results are presented in Appendix Figure A.10).

Definitely evaluating the policy's long-term effect on homeownership would require data spanning a longer period after the policy's introduction, enabling a comparison between treated and untreated individuals later in life—once most home purchase decisions have been made. However, we argue that it is unlikely that the effect on homeownership only emerges years or decades after the policy introduction, especially given that most activity is observed during the brief exemption period. We therefore conclude that the primary impact of the tax exemption for young buyers was to accelerate the timing of home purchases for individuals who would have bought only slightly earlier or later in the absence of the policy.

Table 3: Intensive-Margin Responses to the *Starter Tax Exemption*

	Purchase Amount (1)	Tax Value (2)	Floor Area (3)	Postcode Income (4)	Over-bidding (5)
Treated ×					
Pre-Policy Period (19Q1–20Q1)	−40.5**	0.000	−0.001	−0.002	−0.002
	(21.5)	(0.009)	(0.009)	(0.003)	(0.002)
Anticipation Period (20Q3–20Q4)	−44.0**	0.006	0.016	0.000	−0.002
	(25.0)	(0.010)	(0.010)	(0.003)	(0.002)
Exemption Period (21Q1)	1,270.7***	0.055***	−0.004	0.008***	−0.002
	(33.8)	(0.012)	(0.011)	(0.003)	(0.002)
Cooldown Period (21Q2)	−151.6***	−0.121***	−0.040***	−0.029***	0.009***
	(28.1)	(0.012)	(0.012)	(0.004)	(0.003)
Medium Run (21Q3–23Q4)	240.1***	−0.021***	0.014	−0.007***	−0.003***
	(21.0)	(0.009)	(0.009)	(0.003)	(0.002)
Number of Observations	33mn	238k	116k	244k	116k

Notes: This table summarizes the estimated effect of the *starter tax exemption* on home purchase behavior among individuals aged 30–34 along various intensive margin dimensions. Corresponding event-study graphics are shown in Appendix Figure A.11. Estimated treatment effects $\widehat{\beta}_k$ are based on the difference-in-differences model specified in Equation (1). Column (1) is based on the baseline person-by-quarter panel dataset. Estimates in columns (2) and (4) come from a transaction-level dataset including all purchases involving buyers aged 30–34 and those aged 38–42. Estimates in columns (3) and (5) are based on a subset of transactions for which we have additional broker data from NVM. Treatment is defined either at the individual level (column (1)) or at the buyer level (columns (2)–(5)). Dependent variables are: (1) an individual's monthly home purchase volume in euros, (2) the natural logarithm of a unit's tax-assessed value as of January 2019, (3) the natural logarithm of a unit's floor area, (4) the natural logarithm of median postcode income in 2018, and (5) the log ratio between a units final sale price and the initial list price. Appendix Table A.3 provides additional evidence from splitting the sample between first-time and repeat buyers and estimating heterogeneous effects. Standard errors, clustered at the individual/transaction level, are reported in parentheses (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

4.3. Intensive Margin Responses

We next turn to changes in the type of homes bought. To analyze such intensive-margin responses, we reestimate the difference-in-differences model from Equation (1) for various transaction characteristics. Table 3 reports the estimation results aggregated to the five subperiods that were introduced above.[9]

We begin on the baseline individual-level panel and examine how the policy affected the monetary home purchase volume of younger, eligible individuals relative to older ones. In line with the estimates for the binary purchase indicator (Section 4.1), we find that average purchase amounts spike significantly by 1,271 euros in the initial three-month period without a price limit, and stabilize around an additional 240 euros in the medium run (see column (1) of the table). In addition to the increased likelihood of making a purchase at all, this effect reflects changes in the selection of homes, as well as price changes.

To investigate these additional margins of adjustment, we shift attention away from

[9] Appendix A.5 provides the corresponding event-study graphs and examines heterogeneous effects among first-time and repeat buyers.

purchase probabilities and turn to the transaction-level data, comparing purchases by individuals 30–34 with those of individuals 38–42 for the remainder of this section. The resulting sample includes around 240,000 transactions during our study period. In column (2), we present estimates for differences in the average (2019-level) tax-assessed value as a proxy for the quality of a unit. We find a highly significant six percent increase in the average tax value in the exemption period, suggesting that eligible homebuyers systematically purchased higher-value units in the initial three months while those units were also exempt from transfer taxation. In contrast, we find a significant two-percent drop in average tax values among young buyers in the medium run. This finding most likely reflects the incentive for eligible homebuyers to stay below the price limit after April 2021.

In column (3), we investigate a unit's floor area as an additional quality measure (based on a subsample with broker data). We find little systematic evidence for any policy responses outside the brief cooldown period from April to June 2021. In column (4), we shed light on the type of neighborhood in which treated homebuyers purchase their home. Using postcode-level average income as a proxy for neighborhood quality, we find a qualitatively similar pattern to that observed for unit-level tax values. Young homebuyers move their attention to higher-income neighborhoods during the exemption period. In contrast, they are significantly more likely to buy in lower-income areas after the introduction of the price limit.

Finally, we examine bidding behavior, again relying on the broker data, which includes both sale prices and initial list prices for each available transaction. To capture overbidding, we calculate the log-difference between these two prices. We find little evidence for changes in bidding behavior by eligible homebuyers in response to the tax exemption.

4.4. Aggregate Price Responses

Finally, we investigate to what extent the tax exemption for young homebuyers has been capitalized into property prices. Theory suggests that stimulating housing demand raises equilibrium prices unless housing supply is perfectly elastic. Part of the economic incidence of the tax exemption may thus fall on sellers rather than buyers.

Price effects cannot be identified with the preceding research designs, given the policy's intensive margin effects, infrequent sales of a given home, and secular market trends. To address these challenges and estimate the impact of the starter tax exemption on prices, we propose an auxiliary difference-in-differences analysis exploiting variation in treatment exposure at the dwelling level. More specifically, we group units below and above the policy's price limit as treatment and control units, respectively. As in Section 4.1, we use the units' tax-assessed values from January 2019 for this classification, coding all units with tax assessments below €250,000 as treated. Working with ex ante values rather than realized prices avoids conditioning on an outcome variable, which is par-

ticularly problematic in the presence of bargaining around the cutoff (Slemrod, Weber, and Shan, 2017). This strategy, however, comes at the cost of misclassification: Some units with lower tax values will be classified as treated but may ultimately transact above the price limit, and conversely. The estimates thus have to be interpreted as intention-to-treat (ITT) effects.

We estimate the following difference-in-differences model on the transaction dataset:

$$\ln P_{it} = \sum_{k \neq 20Q2} \beta_k \cdot Treat_i \cdot \mathbb{1}_{(t=k)} + \gamma_{c(i),t} + \delta \cdot Treat_i + \varepsilon_{it}, \tag{3}$$

where $\ln P_{it}$ is the natural logarithm of property i's transaction price at time t. We include flexible postcode-by-period dummies, $\gamma_{c(i),t}$, to compare only units from the same local market c. $Treat_i$ equals one for units with low tax-assessed values and zero for those with higher ones. In contrast to preceding transaction-level analyses, we now include transactions irrespective of the buyer's age, but restrict the sample to units with tax assessments of €200,000–€230,000 (treatment) and €270,000–€300,000 (control group) in the baseline specification. Selecting the range for treated and control units involves a trade-off between the comparability of units on the one hand and differences in exposure on the other. We return to this issue below and present results for various alternative ranges. Standard errors are clustered at the postcode level.

Estimation results are reported in Table 4. From column (1) to (4), we tighten the comparisons by including more fine-grained regional time trends. Starting with uniform national trends, we then control for MSA-by-quarter and city-by-quarter dummies, respectively, and add postcode-by-quarter fixed effects in the baseline specification presented in column (4). In all specifications, pre-trends are small and statistically insignificant. After the price limit was introduced in 2021Q2, the price of lower-valued units increased sharply, suggesting that the subsidy is immediately capitalized into prices. Effect sizes decrease in the medium run, but prices in the treated group remain significantly above their counterfactual trend. In the first specification, including only quarterly fixed effects, we estimate a 1.9% price increase in the medium run. However, highly localized trends lower the effect size. In our preferred specification with postcode-by-quarter fixed effects, we find a medium-run price increase of 1.3% relative to 2020Q2.

Given the ITT nature of this analysis, we relate the estimate to the difference in policy exposure. 95% of the units in the lower price range eventually transact below the price limit and qualify for the tax exemption. In the control group, 23% of units with relatively high tax assessments remain below the price limit. Extrapolating the baseline estimate from column (4), we obtain a price response of $0.013/(0.95 - 0.23) \approx 1.8\%$.

Assigning treatment based on ex ante tax-assessed values involves a critical trade-off. On the one hand, choosing treatment and control units from a narrow band just around a cutoff value ensures high comparability between both groups. On the other hand, a narrow band around a cutoff also increases the risk of misclassification, as it raises

Table 4: The Effect of the *Starter Tax Exemption* on House Prices

	Different Time Trend Specifications				Varying Bandwidth	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated ×						
Pre-Policy (19Q1–20Q1)	−0.003	−0.004	−0.003	−0.005	−0.006*	−0.013***
	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.004)
Anticipation (20Q3–20Q4)	0.006*	0.005	0.004	0.003	0.001*	−0.002
	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)
Exemption (21Q1)	−0.007*	−0.006	−0.006	−0.002	−0.003	−0.011**
	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)	(0.005)
Cooldown (21Q2)	0.020***	0.019***	0.019***	0.020***	0.008*	0.020***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.006)
Medium Run (21Q3–23Q4)	0.019***	0.015***	0.015***	0.013***	0.009**	0.012**
	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.004)
Adjusted-R^2	0.580	0.598	0.621	0.728	0.687	0.761
Number of Transactions	214,652	213,102	214,652	214,652	217,587	211,016
Region × Quarter Fixed Effects	National	MSA	City	Postcode	Postcode	Postcode
Treated Price Range	200–230k				210–240k	190–220k
Control Price Range	270–300k				260–290k	280–310k

Notes: This table summarizes the estimated effect of the starter tax exemption on aggregate house prices. Coefficients are to be interpreted as price growth relative to the omitted quarter 2020Q2. Appendix Figure A.12 visualizes the corresponding quarterly estimates. Estimated treatment effects $\widehat{\beta}_k$ are based on the difference-in-differences model specified in Equation (3), which compared lower-value units (which likely qualify for the exemption) with higher-value units (which likely do not qualify). The estimation is based on the repeated cross-section of all Dutch property transactions from 2019Q1 until 2023Q4 in the respective price ranges. Price ranges refer to tax-assessed values as of January 2019. Column notes indicate the spatial level of the included time fixed effects. MSAs refer to the 40 Dutch metropolitan statistical areas (known as COROP or NUTS 3 regions), cities to the 360 municipalities (*gemeenten*), postcodes to the 3,788 four-digit postcodes. Standard errors, clustered at the postcode level, are reported in parentheses (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

the likelihood that treated (control) units will eventually transact above (below) the price limit. In columns (5) and (6) of Table 4, we demonstrate that the medium-run effects are largely robust to using alternative price ranges for assigning treatment and control groups. However, treatment effects and in particular pre-trends become noisier, in particular when increasing the distance between treatment and control.

The price analyses has to be interpreted with caution, as the medium-run policy period includes a major shock to the macroeconomic environment. Interest rates shot up, and mortgage interest tripled between early 2022 and early 2023, as the European Central Bank responded to inflationary pressure related to the outbreak of the war in Ukraine. Since interest rate shocks may exert heterogeneous effects across housing market segments (Hacamo, 2024), we cannot rule out that the medium-run estimate is partly confounded by these shocks after 2022Q1, when the interest rate hike cycle began. We address such concerns by calculating a lower bound for the medium-run price response that is identified only by the nine-month period from 2021Q3 until 2022Q1. Based on this subperiod, we estimate an ITT price increase of around 0.9%, which

translates into an 1.2% effect when rescaled by the difference in the realized share of tax-exempt units in both groups. We conclude that 65–90% of the tax relief capitalized into prices and eventually benefited sellers. This price increase also helps rationalize the observed behavior of existing homeowners that was documented in Section 4.1 and visualized in Figure A.7. In particular, their increased activity in the high-value segment cannot be explained by the direct effect of the tax exemption. It can, however, be driven by a desire to trade up, if the existing home can be sold at a higher price.

5. Theoretical Framework

We now turn to the efficiency and welfare implications of the young homebuyer relief programs. Section 5.1 presents the basic model setup. Section 5.2 derives the utility consequences of the starter tax exemption. Section 5.3 discusses implications for efficiency and distributional aspects.

5.1. Basic Setup

Consider an economy with a large set of individuals indexed by $i \in \mathcal{I}$ at a given point in time. Individuals differ in preferences, characteristics, incomes, and endowments (including real estate). They receive utility from living in a housing unit, $h \in \mathcal{H}$, and consuming a non-housing composite good, c. Housing units are indivisible and differ in dimensions like size, location, quality, and housing tenure, i.e., whether they are owner-occupied or offered for rent. We denote the annual cost of renting or owning a unit, such as rent or maintenance payments, by $r(h)$. We assume exogenous disposable incomes consisting of earnings y_i and a uniform lump-sum tax rebate R.

Utility is summarized by the function $u_i : \mathbb{R}_+ \times \mathcal{H} \to \mathbb{R}$. We allow $u_i(c, h)$ to be heterogeneous across individuals, and assume it is concave and differentiable in c. As is standard in assignment models, we do not impose differentiability with respect to housing since we treat units as indivisible (see, e.g., Määttänen and Terviö, 2014, and Landvoigt, Piazzesi, and Schneider, 2015).

Household Problem. Individual i maximizes utility subject to the budget constraint:

$$y_i + R \geq c_i + r(h_i), \tag{4a}$$

which applies to both renters and homeowners, provided they don't buy or sell a unit. Buyers have to finance the price of their dwelling, $p(h)$, and real-estate transaction taxes with tax rate $\tau > 0$, net of potential subsidies $s \geq 0$:

$$y_i + R \geq c_i + r(h_i) + (1 + \tau - s \cdot \mathbb{1}_{i \in \mathcal{J}})p(h). \tag{4b}$$

We model the tax exemption as an offsetting subsidy rather than a reduction in the transfer tax (which would be mathematically equivalent, of course), since the starter tax exemption only applied to a subset of individuals $\mathcal{J} \subseteq \mathcal{I}$. It also keeps the model more generic and applicable to other forms of homebuyer relief.

Sellers' budget constraint entails additional income and is given by:

$$y_i + R + p(k) \geq c_i + r(h_i), \tag{4c}$$

where $k \in \mathcal{H}$ denotes the unit individual i sells out of her endowment. Tax revenues are redistributed lump-sum such that: $\sum_i R = \sum_{k \in \mathcal{S}} (\tau - s) p(k)$ with $\mathcal{S} \subseteq \mathcal{H} \cup \{\varnothing\}$ referring to the set of transacted units. Individuals take prices $p(\cdot)$ and $r(\cdot)$, income y_i, and transfer R as exogeneously given.

Equilibrium. Given the little structure we imposed on the economy, we cannot derive the equilibrium. In what follows, we follow the sufficient-statistics literature and assume that individuals have optimized their utility and that the economy is in equilibrium (see, e.g., Chetty, 2009, and Kleven, 2021). Observed choices thus reflect individually optimal bundles $\{(c_i^*, h_i^*)\}_{i \in \mathcal{I}}$. We then focus the analysis on small perturbations in the tax and subsidy code. An alternative approach in assignment models of the housing market is to impose more structure on the problem in order to characterize the equilibrium (see, e.g., Määttänen and Terviö, 2022).

5.2. Subsidies for Homebuyers

Against the backdrop of the institutional setting and the empirical results from Section 4, we are interested in the utility consequences of a small increase in the purchase subsidy, $\mathrm{d}s > 0$. By the envelope theorem, utility changes $\mathrm{d}u_i/\mathrm{d}s$ are governed solely by price changes, whereas behavioral responses have no first-order effect on welfare as individuals are already at their optimum. We argue this is a plausible approximation in the context of the Dutch starter tax exemption, since we do not find any compelling evidence for extensive margin effects on homeownership.

As in Levy (2024), it is useful to distinguish between different subpopulations:[10]

1. Never takers are individuals who do not buy or sell units, neither with nor without the subsidy. They do not benefit from the subsidy but must finance the subsidy's costs—the reduction in tax revenues—through changes to the lump-sum rebate. Their money-metric utility change can thus be expressed as:

$$\Delta c^{\mathrm{NT}} := \frac{\mathrm{d}u_i/\mathrm{d}s}{\partial u_i/\partial c} = \frac{\mathrm{d}R}{\mathrm{d}s} \qquad \forall i \in \mathcal{I}^{\mathrm{NT}}, \tag{5a}$$

[10] We deliberately focus on mutually exclusive population groups to simplify the exposition. In reality, people may, of course, buy and sell units around the same time. We take this into account when quantifying welfare effects in Section 6.

which is negative as long as the economy is on the bright side of the Laffer curve.

2. We next consider always buyers, i.e., individuals who would have bought and switched home to the same unit also absent a reform. A small increase in the subsidy s affects always buyers' money-metric utility by:

$$\Delta c^{\text{AB}} := \frac{\mathrm{d}u_i/\mathrm{d}s}{\partial u_i/\partial c} = p \cdot \mathbb{1}_{i\in\mathcal{J}} - (1+\tau-s)\frac{\mathrm{d}p}{\mathrm{d}s} + \frac{\mathrm{d}R}{\mathrm{d}s} \qquad \forall i \in \mathcal{I}^{\text{AB}}, \tag{5b}$$

where the first term on the right-hand side of the equation represents the direct gain from the subsidy and only accrues to eligible buyers. The second term represents price effects that may partially offset the direct effect. The third term gives the change in the lump-sum transfer, which is, again, likely negative.

The intuition is that always buyers would have chosen bundle (c', h') absent a reform and can now afford $(c' + \Delta c^{\text{AB}}, h')$, which still yields the utility from living in unit h' but gives higher composite good consumption. This difference in consumption, Δc^{AB}, measures their money-metric utility gain.

3. For always sellers, money-metric utility effects depend only on the change in transaction prices and the effect on the lump-sum rebate:

$$\Delta c^{\text{AS}} := \frac{\mathrm{d}u_i/\mathrm{d}s}{\partial u_i/\partial c} = \frac{\mathrm{d}p(k)}{\mathrm{d}s} + \frac{\mathrm{d}R}{\mathrm{d}s}. \qquad \forall i \in \mathcal{I}^{\text{AS}}. \tag{5c}$$

The sign of this expression is theoretically undetermined as the former term is likely positive and the latter likely negative. The larger the price effect, i.e., the capitalization of the subsidy in sales prices, the bigger the gain for always sellers.

4. New buyers wouldn't trade under the old regime but reconsider with the (increased) subsidy. Their change in money-metric utility is bounded by the polar cases of always buyers and never takers, and therefore in between Δc^{NT} and Δc^{AB} (Milgrom and Segal, 2002). Always buyers prefer purchasing even in the absence of the policy, which implies $u(c', h') > u(c, h)$. In contrast, never buyers prefer their initial bundle (c, h) over any (c', h') even at the new, lower prices. Both groups are inframarginal with and without the subsidy.

New buyers are in between: The marginal person would already switch to h' for a tiny compensating payment ϵ, such that $u(c', h') < u(c, h)$ but $u(c'+\epsilon, h') > u(c, h)$. Her corresponding money-metric utility gain from the subsidy is $\Delta c^{\text{AB}} - \epsilon$, and thus almost as much as for always buyers. The last switcher considers buying only for a large compensation of $\Delta c^{\text{AB}} - \epsilon$, thereby gaining marginally higher money-metric utility than never-takers. She was inframarginal before and becomes the marginal person with the subsidy. The subsidy largely goes into making her consider h'.

5. Finally, new sellers are bounded by never takers and always sellers $(\Delta c^{\text{NT}}, \Delta c^{\text{AS}})$.

These utility consequences are governed by few empirical objects ("sufficient statistics"): (i) The pass-through of the subsidy on prices, $\mathrm{d}p/\mathrm{d}s$, (ii) its impact on tax revenues, $\mathrm{d}R/\mathrm{d}s$, and (iii) its impact on transactions and the number of switchers.

5.3. Efficiency and Distribution

We assume that policymakers would want to support young homebuyers whenever such a subsidy raises social welfare:

$$\mathrm{d}W/\mathrm{d}s := \sum_i g_i \frac{\mathrm{d}u_i/\mathrm{d}s}{\partial u_i/\partial c} = \sum_i g_i \Delta c_i > 0, \tag{6}$$

with exogeneously given, fixed social marginal welfare weights denoted by g_i (Saez and Stantcheva, 2016). Equation (6) captures both efficiency and equity concerns.

Efficiency. Efficiency, or aggregate welfare, is given by the sum $\sum_i \Delta c_i$ over all individual Δc_i terms defined above (as in Määttänen and Terviö, 2022). This measure corresponds directly to the deadweight loss logic of the Harberger triangle. Raising the subsidy s allows more buyers and sellers to realize mutually beneficial transactions, thereby reducing the distortion caused by the real-estate transfer tax τ. The larger the groups of new buyers and new sellers, respectively, and the larger the transaction volume, the greater the aggregate welfare gain.

A standard measure for this efficiency gain is the tax base elasticity, which we estimated in Section 4.3. However, treating $\sum_i \Delta c_i$ as a measure of social welfare gains requires the restrictive assumption $g_i = 1$.

Distribution. In the absence of behavioral responses, the above formulas reduce to pure resource redistribution between individuals. Whereas such redistribution would not increase the overall surplus in society, it may still affect social welfare when individuals differ in their welfare weights g_i. These distributive welfare effects are governed by price changes, $\mathrm{d}p/\mathrm{d}s$, and thus tax incidence.

To examine the distributional aspect, we simulate the utility consequences of the starter tax exemption for each household in the Netherlands and calculate average money-metric utility changes for various quantiles of the income distribution. This allows us to shed light on the second dimension of social welfare, namely, distributional effects among buyers, sellers, and never-takers across the income distribution. Household-level utility effects can further be aggregated to social welfare changes $\mathrm{d}W/\mathrm{d}s$ using any arbitrary set of desired weights g_i.

6. Welfare Simulation

We now use the estimated empirical parameters to simulate the welfare consequences of exempting young homebuyers from transfer taxes. Relying on our detailed administrative data, we simulate utility changes due to the policy at the individual level, and analyze them along the distribution of household income. First, we explain how the simulation is implemented, drawing on the theoretical framework from section 5. Then, we implement the baseline simulation, using the estimated parameters from section 4 and discuss sensitivity analyses.

6.1. Empirical Implementation

We base the welfare simulation on the 2022 snapshot of the administrative data, including all Dutch-born individuals aged 18 or older. This represents the status quo after the introduction of the starter tax exemption. Therefore, utility changes due to the policy are expressed as money-metric compensating variation. For each individual in the sample, we identify household-level disposable income, which we normalize to account for household size based on the OECD equivalence scale.

First, individuals are classified into the various groups introduced in section 5. Based on their observed housing market activity during the one-year period from 2021Q3 until 2022Q2, individuals are labeled as buyers, sellers (or both) or never takers.[11] Next, we introduce two additional distinctions. First, we distinguish among buyers between those above and below the age of 35, as only the latter benefit from the tax exemption. Second, we identify all transactions—both purchases and sales—below €400,000, the policy's price limit. To simulate new transactions generated by the policy, we first split the sample of all young buyers into quartiles, based on their normalized household income. Then, we randomly select a subset of young buyers in each quartile as new (i.e., policy-induced) buyers based on the estimated treatment effects by income quartile. For instance, in the first quartile, the policy is estimated to increase purchase probabilities by 32.5%.[12] Therefore, we label 75.5% ($= 1/(1+0.325)$) of all observed young buyers in this quartile as always buyers, and the remaining 24.5% as new buyers. Conversely, we label as new sellers those individuals who sell to a new buyer, whereas all other sellers are labeled as always sellers.

Individual-level utility changes are calculated based on equations (5a) – (5c), using the reduced form estimates from section 4. In the baseline scenario, we set $\frac{\mathrm{dp}}{\mathrm{ds}} = 0.012$, which equals the estimated policy's medium-run price effect, adjusted for the 2022 interest rate shock. To obtain $\frac{\mathrm{d}R}{\mathrm{d}s}$, we calculate the total foregone transfer tax revenue due to the

[11] We select transactions during this window, as they are most representative of the medium-run effect of the policy, i.e., least affected by timing responses or confounding factors due to interest rate changes.

[12] As shown in column (4) of Table 1, the estimated monthly effect in the medium-run in the first quartile equals 0.064%. This corresponds to a 32.5% increase relative to the average monthly purchase probabiliy in 2020Q2 among individuals in the first income quartile, which was equal to 0.197%.

exemption of (always) buyers, and divide it by the size of the adult population. This corresponds to a per-person tax rebate change of −€20.[13] the For new buyers and sellers, we draw utility changes from a uniform distribution, bound between the utilities of never takers and that of always buyers or sellers, respectively. This follows the intuition, outlined above, that individuals who change their behavior due to the subsidy must have a lower utility gain than those for which the subsidy is a pure windfall. We assume positive assortative matching (Chade, Eeckhout, and Smith, 2017): for that new buyers with the smallest utility gain, the corresponding seller is also assigned the smallest utility gain etc. For individuals who are both buyers and sellers, we add the respective utility changes due to these two actions. After calculating utility changes for each individual in the sample, we sum over all (adult) members of a household. Then, we aggregate to different quantiles of the standardized household income distribution.

6.2. Simulation Results

On average across the entire population, we find that the transfer tax exemption for young homebuyers induces welfare gain that is positive, but smaller than €1 in magnitude. This is unsurprising given that, since transfer taxes are distortionary, their removal creates a number of welfare enhancing-transaction that would not have occured otherwise. However, due to the detailed administrative data, we are able to simulate utility changes across the entire population, allowing us to identify who gains and loses from the policy.

The results of the this simulation are visualized in Figure 4. In all four panels, individuals are ordered along the x-axis based on their standardized household income, which we aggregate into ventiles (i.e., five-percentile groups). Panel (a) shows the baseline results, revealing a hump-shaped pattern. Low-income individuals lose utility due to the tax-exemption for young homebuyers.[14] With rising income, the policy becomes increasingly more beneficial and its positive impact peaks around the 70th percentile. Beyond that, the utility gain declines again, and for those at the top of the distribution, average gains turn negative.

An even stronger increase in prices due to the subsidy would not change this pattern qualitatively. This is shown in panel (b), where we contrast the baseline with a second simulation, based on a 1.8% price increase (in grey). Compared to the baseline, a higher price effect would—perhaps unintuitively—make the tax exemption slightly more desirable for individuals with lower incomes, at the expense of those with very high incomes. This is due to the fact that with a higher price effect, some of the policy's

[13] New buyers are not considered in the calculation of foregone tax revenue, as they would not have made a purchase absent the subsidy.

[14] Individuals in the lowest income percentiles appear to deviate from the overall trend. This is likely due to mismeasured household incomes, or driven by households with little income but high wealth (e.g., wealthy pensioneers). Evidence for this is seen in panel (c): at the bottom of the distribution, the share of observed homesellers increases.

negative effect on the tax rebate is recaptured via higher tax-receipts from ineligible homebuyers. As a consequence, never takers—i.e., those who neither buy or sell—would have a smaller utility loss in this scenario.

To better understand what drives the hump-shaped effect of the policy on individual-level utility, we separately examine the different societal groups. Intuitively, the average utility change per income percentile depends only on two factors: the share of (in-)eligible buyers, sellers and never takers in this percentile, and their respective utility change. We disentangle these two components in panels (c) and (d) of Figure 4. Panel (c) reveals that the share of individuals who participate in the housing market is steadily increasing with income. Conversely, the bottom of the income distribution comprises almost exclusively of never-takers. These are not directly affected by the policy, but nonetheless have to finance it via the tax rebate, as captured by equation (5a). With rising incomes, the share of buyers and sellers increases. In addition, the distinction between younger and older buyers is noteworthy, as only the former group benefits from the subsidy. Across the distribution, there are somewhat more younger buyers than older ones. Only among those with very high incomes, this pattern reverses, as the share of older buyers increases sharply, whereas the share of younger buyers decreases. These trends are relevant since—as shown in panel (d)—the policy's utility consequence varies drastically by group. Whereas sellers and young homebuyers benefit, older buyers suffer from higher prices. This is most pronounced among older buyers with lower incomes, who are most likely to buy lower value units, which are subject to the price increase. In contrast, buyers and sellers with higher incomes tend to be less affected, as their activity is focused on higher-value units, which are unaffected by the policy.

In sum, the results of our welfare simulation suggest that, while the *starter tax exemption* did increase the overall economic surplus marginally, it mainly led to a redistribution of resources. Those who gain from the policy are in the upper-middle part of the income distribution, whereas those with the lowest incomes lose most. We conclude from this that—under any common set of welfare weights—subsidizing young homebuyers is unlikely to promote social welfare.

7. Conclusion

Support schemes for young and first-time homebuyers are common in most developed economies, but their impact on overall homeownership, efficiency and welfare remains unclear. In this study, we analyze the effects of policies that directly support young homebuyers by reducing the upfront costs of their purchases. This, we argue, is the typical way in which most real-world policies operate. We make three contributions. First, we empirically estimate the reduced-form effects of supporting young homebuyers, using quasi-experimental variation in the Dutch housing market, where homebuyers under 35 have been exempt from transfer taxes since 2021. Second, we outline a simple

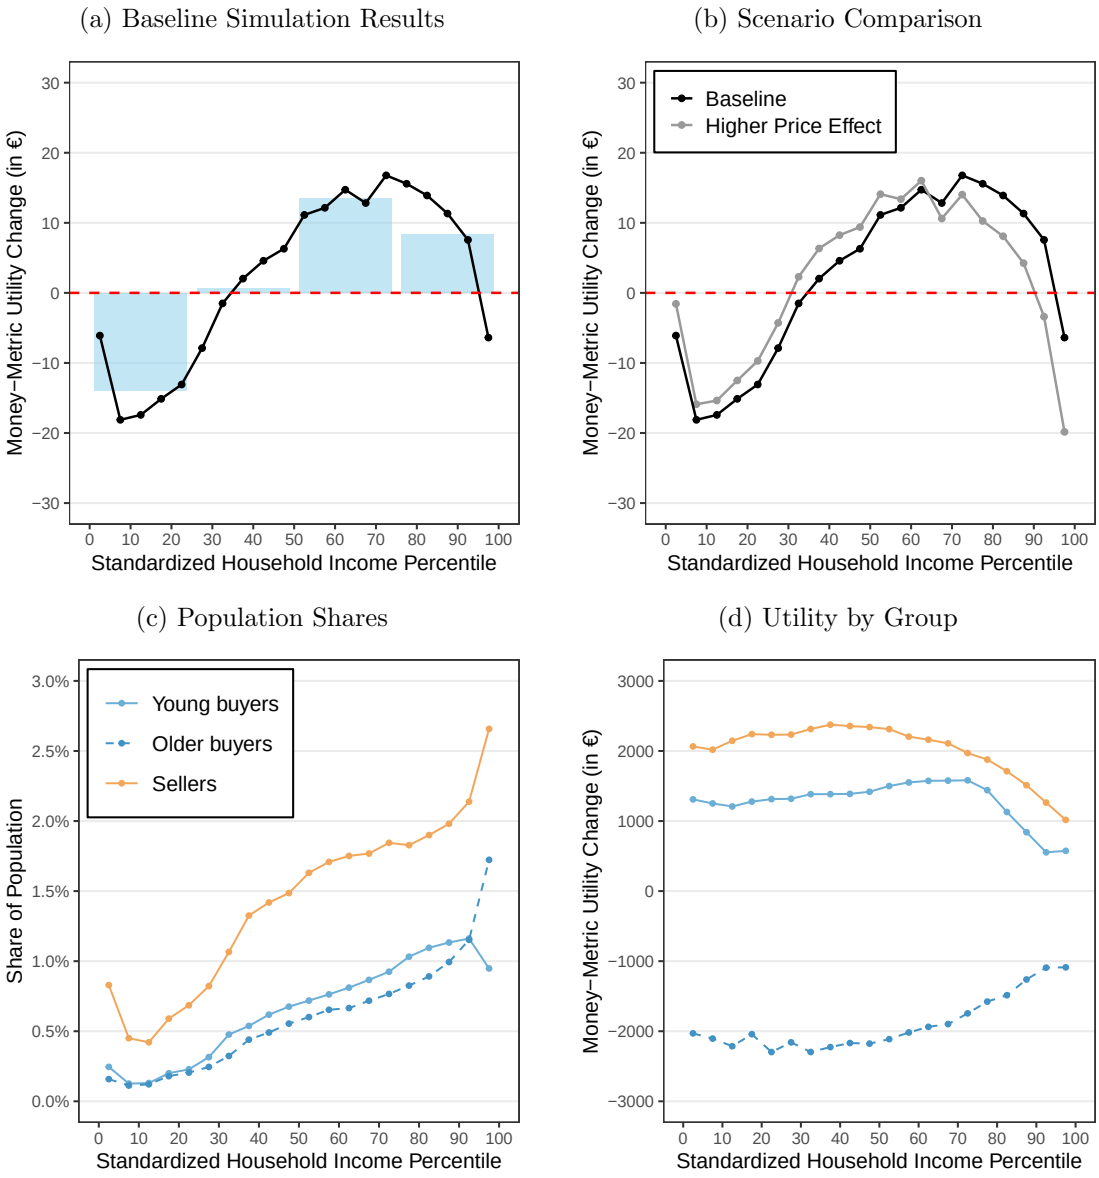


Figure 4: Welfare Simulation of the *Starter Tax Exemption*

Notes: This Figure shows the result of simulating the welfare effects of the *starter tax exemption* for the Dutch adult population, along the household income distribution. Panel (a) visualizes the main simulation result. Average utility changes due to the policy (in €) per ventile are shown in black, and averages per quartile are represented by blue bars. In panel (b), the baseline simulation is contrasted with a higher price effect (1.8% increase) scenario in grey. Panel (c) shows the share of sellers, young buyers and older buyers along the income distribution. Panel (d) shows average utility changes for the same groups. For never takers, which are omitted from the figure, utility changes always equal to $\frac{dR}{ds} = -€20$.

assignment model with tenure choice and indivisible housing to estimate the welfare consequences of home purchase subsidies. These are governed by three key empirical parameters (sufficient statistics): changes in aggregate prices, changes in tax revenues induced by the policy, and the number of additional buyers and sellers. In contrast, behavioral changes due to a subsidy have no first-order utility effect. This is the case in the Dutch setting, since the tax relief mainly led to short-term timing responses. Third,

we combine theory and empirics, and simulate the welfare consequences of the Dutch starter tax exemption for all Dutch households.

The policy had a marked effect on the Dutch housing market. In the short run, purchases by eligible young homebuyers doubled relative to their pre-policy level. However, more than half of this spike is explained by bunching around the introduction date. In the medium run, purchases remain 13% above their pre-policy level. However, this additional activity did not translate into increased homeownership rates. There are two reasons for this. On the one hand, many additional transactions were made by existing homeowners who used the tax exemption to trade up, likely stimulated by higher sale prices on their previous homes. On the other hand, at least four of five additional transactions created by the policy resulted from bunching around the 35th birthday, which marks the eligibility cutoff. This is revealed by an RD analysis, which shows that the strongest increase in purchases is driven by individuals who are about to turn 35. In contrast, when examining homeownership as an outcome directly, we can rule out even modest effects of the policy only two years after its introduction. Turning to aggregate outcomes, we find that most of the subsidy has been capitalized into higher prices. In an auxiliary difference-in-differences approach using a repeated cross-section of property transactions, we find that the price of units that qualified for the tax exemption rose by 1.2–1.8% more than that of unaffected units. This suggests that most of the subsidy's incidence fell on sellers rather than buyers. Overall, the policy change had a small positive effect on economic surplus by removing a distortionary transfer tax for young buyers. However, its predominant effect was to redistribute resources towards individuals in the upper middle class, who were most likely to benefit directly from the exemption. For any common set of welfare weights, it is therefore unlikely to be socially desirable. Finally, we conduct additional empirical tests which provide insights into homebuyers' intensive-margin response to purchase subsidies. Most importantly, we find a strong tendency of eligible buyers to remain below the price limit. One side effect of this is that, due to the policy, young homebuyers select into lower-value units in lower-value neighborhoods.

Our study contributes to a better understanding of support policies for young homebuyers. The empirical results suggest that, even though the policy we analyze offers substantial financial support for young homebuyers, it ultimately fails to achieve its stated goal of promoting (earlier) access to homeownership.

References

Berger, D., N. Turner, and E. Zwick (2020). "Stimulating Housing Markets". *Journal of Finance* 75.1, pp. 277–321.

Best, M. C. and H. J. Kleven (2018). "Housing Market Responses to Transaction Taxes: Evidence from Notches and Stimulus in the UK". *Review of Economic Studies* 85.1, pp. 157–193.

Bos, J. W., N. Kok, and J. Wogh (2024). "Bidder Beware: Intergenerational Wealth Transfers and Overpayment in Housing Markets". Mimeo.

Carozzi, F., C. A. Hilber, and X. Yu (2024). "On the Economic Impacts of Mortgage Credit Expansion Policies: Evidence from Help to Buy". *Journal of Urban Economics* 139, p. 103611.

Chade, H., J. Eeckhout, and L. Smith (2017). "Sorting through search and matching models in economics". *Journal of Economic Literature* 55.2, pp. 493–544.

Chetty, R. (2009). "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-form Methods". *Annual Review of Economics* 1, pp. 451–488.

DNB (2021). *Onze Hoge Hypotheekschulden – Risico's en Oplossingen*. https://www.dnb.nl/actuele-economische-vraagstukken/woningmarkt/onze-hoge-hypotheekschulden-risico-s-en-oplossingen/#:%7E:text=We%20lenen%20ook%20veel%20geld,een%20externe%20site)%20aan%20hypotheekschuld..

Dolls, M., C. Fuest, C. Krolage, and F. Neumeier (2025). "Who Bears the Burden of Real Estate Transfer Taxes? Evidence from the German Housing Market". *Journal of Urban Economics* 145.103717.

Floetotto, M., M. Kirker, and J. Stroebel (2016). "Government Intervention in the Housing Market: Who Wins, Who Loses?" *Journal of Monetary Economics* 80, pp. 106–123.

Flood, S., M. King, R. Rodgers, S. Ruggles, J. R. Warren, D. Backman, E. Breton, G. Cooper, J. A. Rivera Drew, S. Richards, D. Van Riper, and K. C. Williams (2025). *IPUMS CPS, Version 13.0 [dataset]*. Minneapolis, MN.

Francke, M., L. Hans, M. Korevaar, and S. Van Bekkum (2023). "Buy-to-Live vs. Buy-to-Let: The Impact of Real Estate Investors on Housing Costs and Neighborhoods". Mimeo.

Glaeser, E. L. and J. D. Shapiro (2003). "The Benefits of the Home Mortgage Interest Deduction". *Tax Policy and the Economy* 17, pp. 37–82.

Gobillon, L. and D. Le Blanc (2008). "Economic Effects of Upfront Subsidies to Ownership: The Case of the Prêt à Taux Zéro in France". *Journal of Housing Economics* 17.1, pp. 1–33.

Goodwin, K. and L. Zumpano (2011). "The Home Buyer Tax Credit of 2009 and the Transition to Homeownership". *Journal of Housing Research* 20.2, pp. 211–224.

Gruber, J., A. Jensen, and H. Kleven (2021). "Do People Respond to the Mortgage Interest Deduction? Quasi-Experimental Evidence from Denmark". *American Economic Journal: Economic Policy* 13.2, pp. 273–303.

Hacamo, I. (2024). "Interest Rates and the Distribution of House Prices". *Review of Economic Studies*. Accepted.

Hembre, E. (2018). "An Examination of the First-Time Homebuyer Tax Credit". *Regional Science and Urban Economics* 73, pp. 196–216.

Hilber, C. A. and T. M. Turner (2014). "The Mortgage Interest Deduction and Its Impact on Homeownership Decisions". *Review of Economics and Statistics* 96.4, pp. 618–637.

Kleven, H. (2021). "Sufficient Statistics Revisited". *Annual Review of Economics* 13 (1), pp. 515–538.

Kopczuk, W. and D. Munroe (2015). "Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate Market". *American Economic Journal: Economic Policy* 7.2, pp. 214–257.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015). "The Housing Market(s) of San Diego". *American Economic Review* 105.4, pp. 1371–1407.

Levy, A. (2024). *Regulating Housing Quality: Evidence from France*. Mimeo.

Milgrom, P. and I. Segal (2002). "Envelope Theorems for Arbitrary Choice Sets". *Econometrica* 70.2, pp. 583–601.

Määttänen, N. and M. Terviö (2014). "Income Distribution and Housing Prices: An Assignment Model Approach". *Journal of Economic Theory* 151, pp. 381–410.

– (2022). "Welfare Effects of Housing Transaction Taxes: A Quantitative Analysis with an Assignment Model". *Economic Journal* 132 (644), pp. 1566–1599.

Poterba, J. and T. Sinai (2008). "Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income". *American Economic Review* 98.2, pp. 84–89.

Saez, E. and S. Stantcheva (2016). "Generalized Social Marginal Welfare Weights for Optimal Tax Theory". *American Economic Review* 106.1, pp. 24–45.

Slemrod, J., C. Weber, and H. Shan (2017). "The Behavioral Response to Housing Transfer Taxes: Evidence from a Notched Change in DC Policy". *Journal of Urban Economics* 100, pp. 137–153.

Tweede Kamer (2020). *Nota naar aanleiding van het verslag bij wetsvoorstel 35576 (Wet differentiatie overdrachtsbelasting), nr. 6*. Parliamentary document (Kamerstuk 35576, nr. 6), Vergaderjaar 2020–2021. Ontvangen 20 oktober 2020; accessed Dec 22, 2025. URL: https://zoek.officielebekendmakingen.nl/kst-35576-6.pdf.

Warden, P. and A. Kadzielawski (2025). *Down Payment Assistance Programs & Grants by State 2026*. Accessed: December 11, 2025. The Mortgage Reports. URL: https://themortgagereports.com/33553/complete-guide-to-down-payment-assistance-in-the-usa.

A. Additional Reduced-Form Analyses

A.1. More Details on Home Purchases

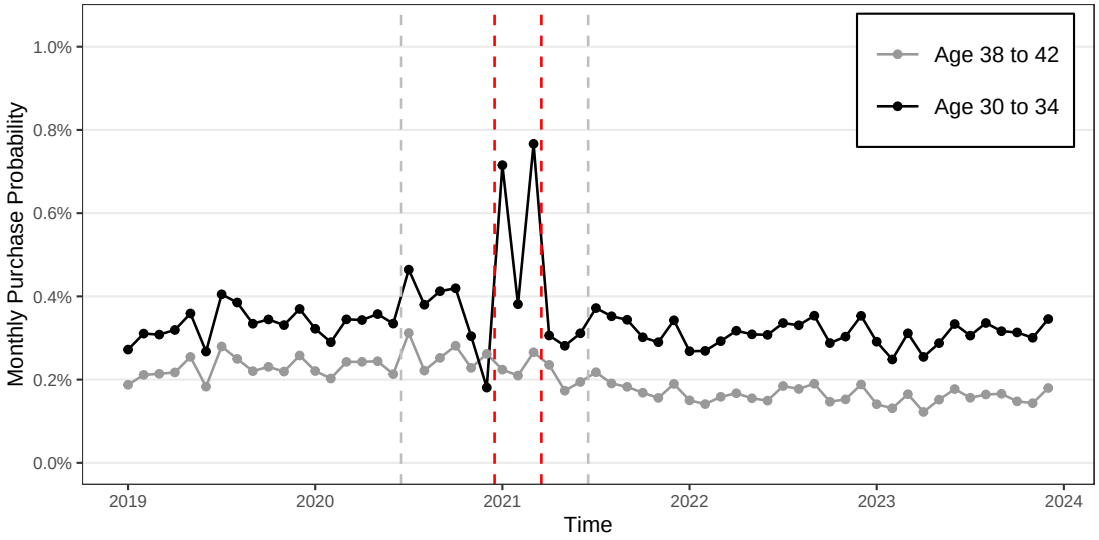


Figure A.1: Unadjusted Home Purchase Rates Among Younger and Older Individuals

Notes: This figure shows the share of individuals aged 30–34 and 38–42 buying a home in a given month between January 2019 and December 2023, including both first-time and repeat purchases. This unadjusted time series is the basis for the difference-in-difference model visualized in Figure 1.

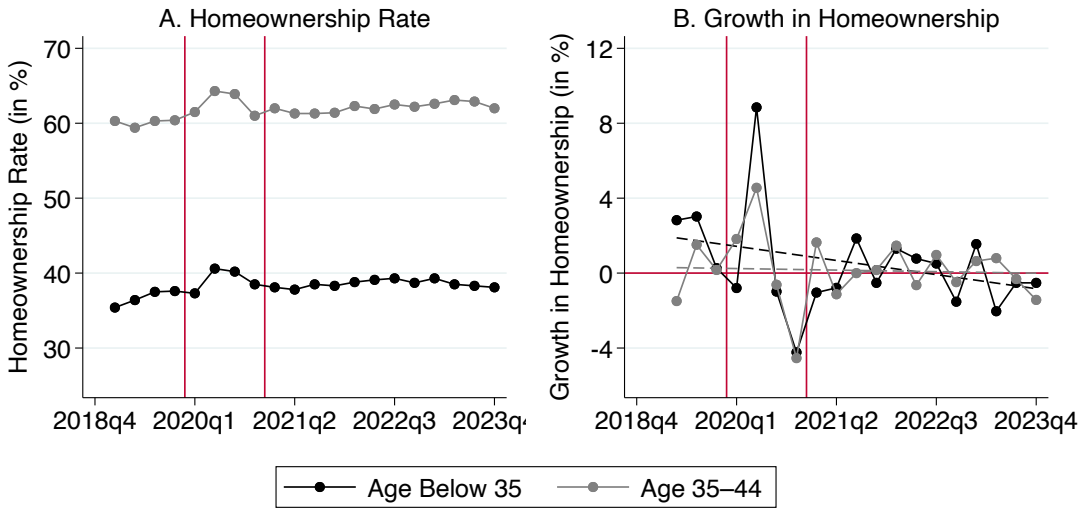


Figure A.2: Homeownership Rates and Growth in the US

Notes: This figure depicts trends in homeownership among young individuals in the US over the period 2019–2023. Panel A contrasts homeownership rates among individuals under 35 with those aged 35–44. Panel B shows quarter-to-quarter growth rates in homeownership for the two groups. The underlying timeseries data comes from the Current Population Survey (IPUMS CPS, 2025).

Table A.1: Monthly Difference-in-Differences Estimates

	Raw Data		Estimated		Raw Data		Estimated
Month	Younger 30–34 (1)	Older 38–42 (2)	Normalized Difference (3)	Month	Younger 30–34 (1)	Older 38–42 (2)	Normalized Difference (3)
January 2019	0.272	0.188	–0.027	July 2021	0.372	0.218	0.042
February 2019	0.311	0.211	–0.012	August 2021	0.352	0.191	0.050
March 2019	0.308	0.214	–0.017	September 2021	0.344	0.183	0.049
April 2019	0.319	0.217	–0.010	October 2021	0.302	0.168	0.022
May 2019	0.359	0.255	–0.007	November 2021	0.290	0.156	0.022
June 2019	0.267	0.183	–0.027	December 2021	0.343	0.190	0.041
July 2019	0.405	0.280	0.014	January 2022	0.268	0.150	0.006
August 2019	0.385	0.250	0.024	February 2022	0.269	0.141	0.016
September 2019	0.334	0.220	0.002	March 2022	0.292	0.159	0.022
October 2019	0.345	0.231	0.002	April 2022	0.317	0.167	0.038
November 2019	0.331	0.219	0.000	May 2022	0.309	0.155	0.042
December 2019	0.370	0.258	0.000	June 2022	0.308	0.149	0.046
January 2020	0.322	0.221	–0.010	July 2022	0.336	0.185	0.040
February 2020	0.290	0.203	–0.025	August 2022	0.331	0.178	0.041
March 2020	0.345	0.243	–0.009	September 2022	0.354	0.190	0.052
April 2020	0.343	0.243	–0.011	October 2022	0.288	0.147	0.029
May 2020	0.357	0.244	0.002	November 2022	0.303	0.153	0.039
June 2020	0.335	0.213	0.010	December 2022	0.353	0.188	0.053
July 2020	0.464	0.312	0.040	January 2023	0.291	0.141	0.038
August 2020	0.380	0.222	0.047	February 2023	0.248	0.131	0.006
September 2020	0.412	0.252	0.049	March 2023	0.311	0.165	0.035
October 2020	0.419	0.281	0.026	April 2023	0.255	0.122	0.021
November 2020	0.304	0.228	–0.036	May 2023	0.288	0.152	0.024
December 2020	0.181	0.262	–0.193	June 2023	0.334	0.177	0.045
January 2021	0.715	0.224	0.380	July 2023	0.306	0.156	0.038
February 2021	0.381	0.210	0.060	August 2023	0.336	0.164	0.060
March 2021	0.767	0.266	0.389	September 2023	0.316	0.166	0.039
April 2021	0.306	0.235	–0.041	October 2023	0.313	0.148	0.054
May 2021	0.281	0.173	–0.004	November 2023	0.301	0.143	0.045
June 2021	0.311	0.194	0.005	December 2023	0.346	0.180	0.054

Notes: This table reproduces monthly difference-in-differences estimates from Figure 1 (blue bars). Columns (1) and (2) of each row reproduce the raw monthly purchase probability per group, as shown in Figure A.1. Column (3) shows the difference between the two groups, normalized to the average difference between April and June of 2020.

A.2. Heterogeneity by Income, Liquid Assets and Buyer Type

Using administrative data, we identify each individual's disposable household income and household bank balances for each year. Individuals are then grouped into quartiles based on their 2019 values for both variables. This approach ensures that the covariates are measured prior to the policy and remain fixed over time. One limitation is that some individuals may still have been living with their parents in 2019, in which case income and assets reflect those of the parental household. To assess the stability of these classifications, we compare 2019 quartiles assignments with those in later years (2020–2023) at the individual-year level. As shown in Figure A.3, most individuals remain in the same quartile over time. We therefore proceed by classifying individuals based on their 2019 income and liquid asset levels.

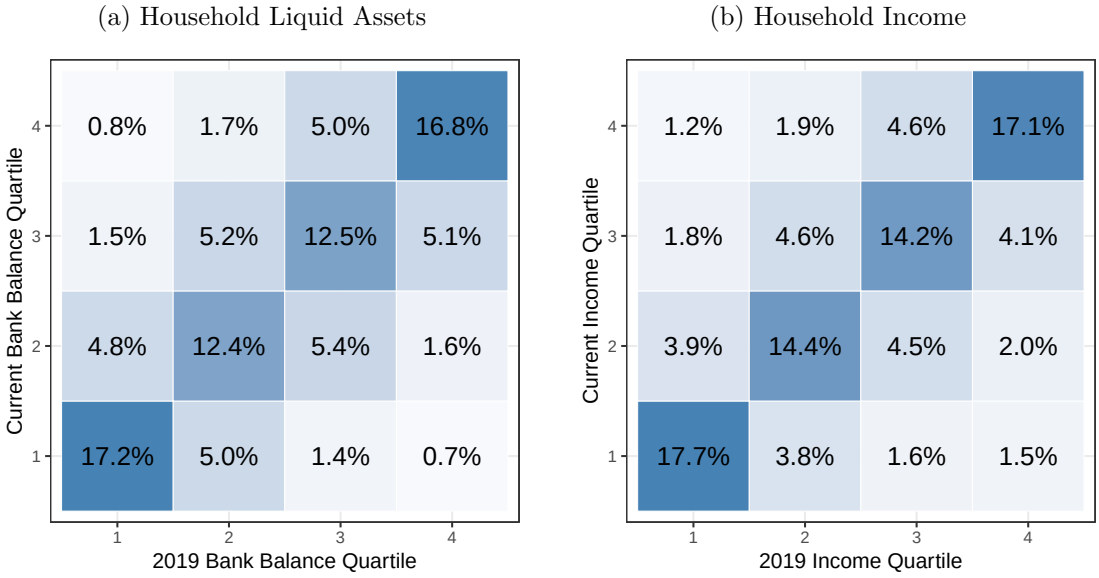


Figure A.3: Income and Liquid Assets: 2019 vs. Current

Notes: This figure shows confusion matrixes that compare individual's liquid assets (income) in 2019 to their liquid assets (income) in all later years. Cells on the diagonal from bottom-left to top-right denote individuals who remain in the same liquid asset (income) quartile as in 2019.

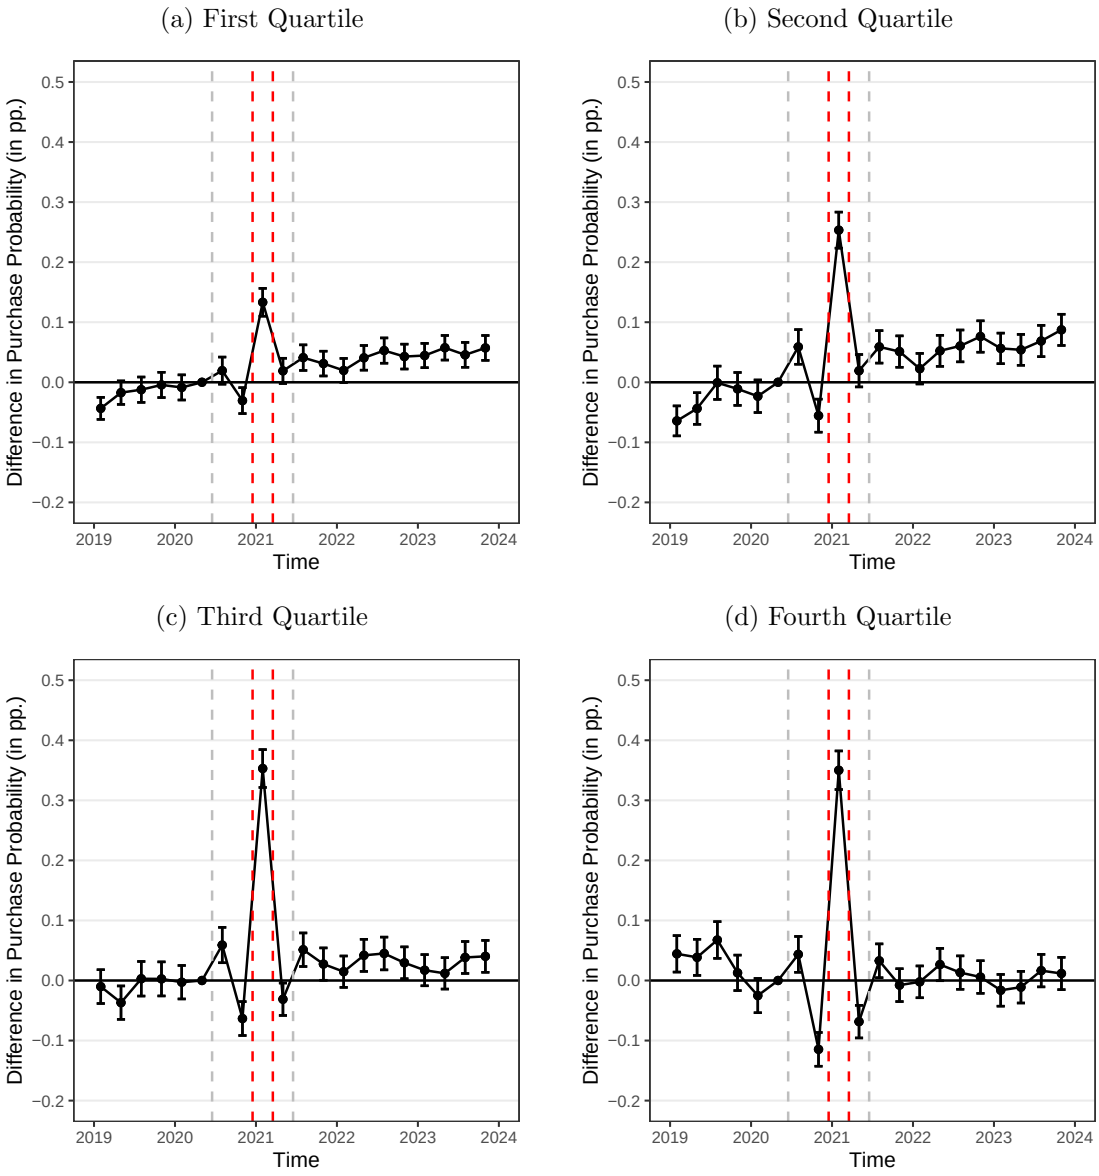


Figure A.4: The Effect of the Starter Tax exemption by Liquid Asset Quartiles

Notes: This figure shows the difference in home purchase probabilities between individuals aged 30-34 and those aged 38-42, across different household liquid asset quartiles. In each subsample, differences are normalized to their 2020Q2 level. Black vertical bars indicate 95% confidence bounds.

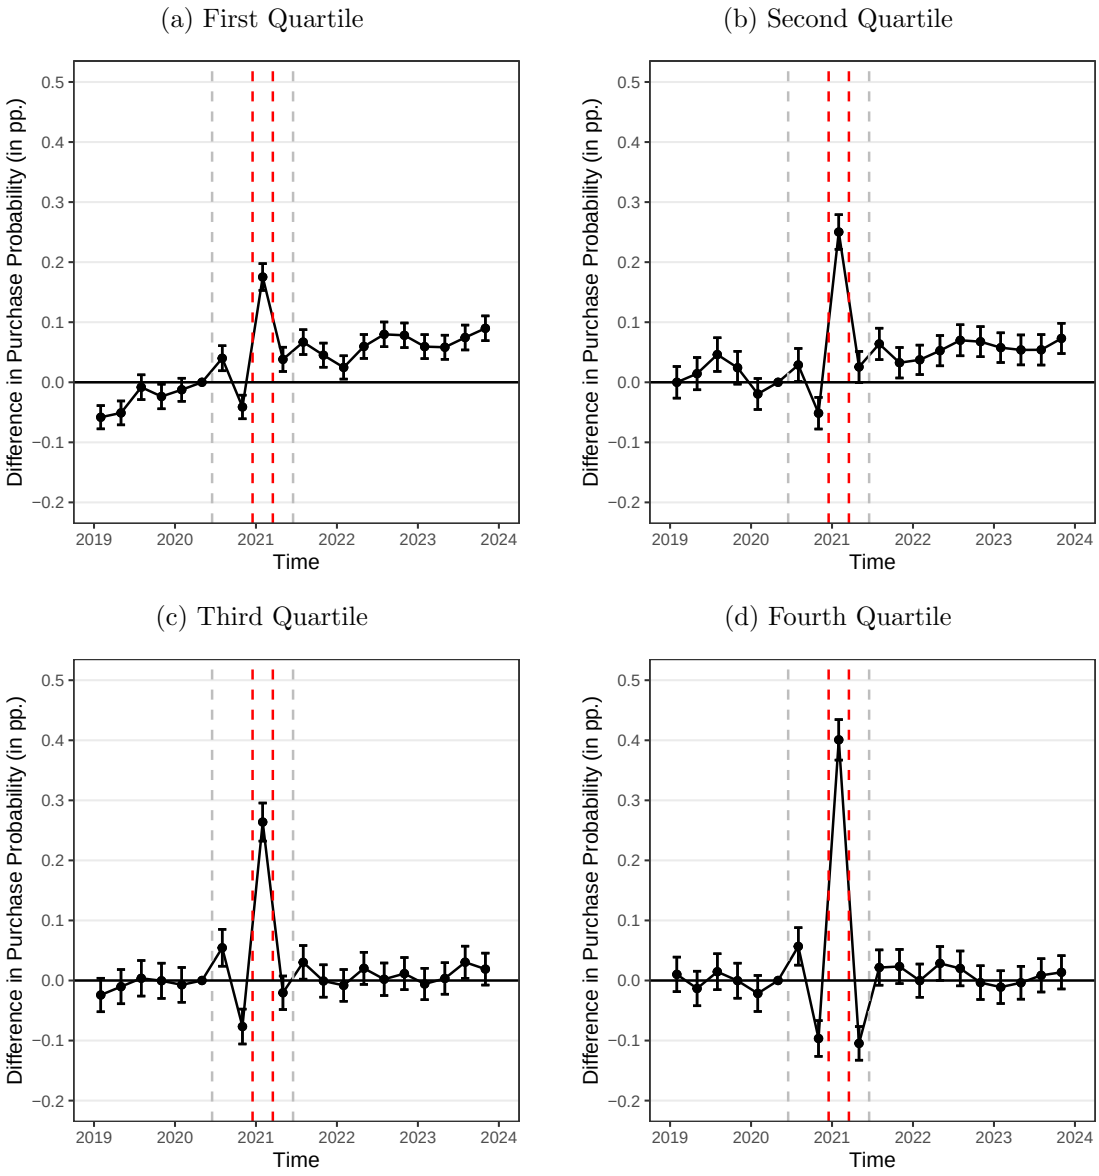
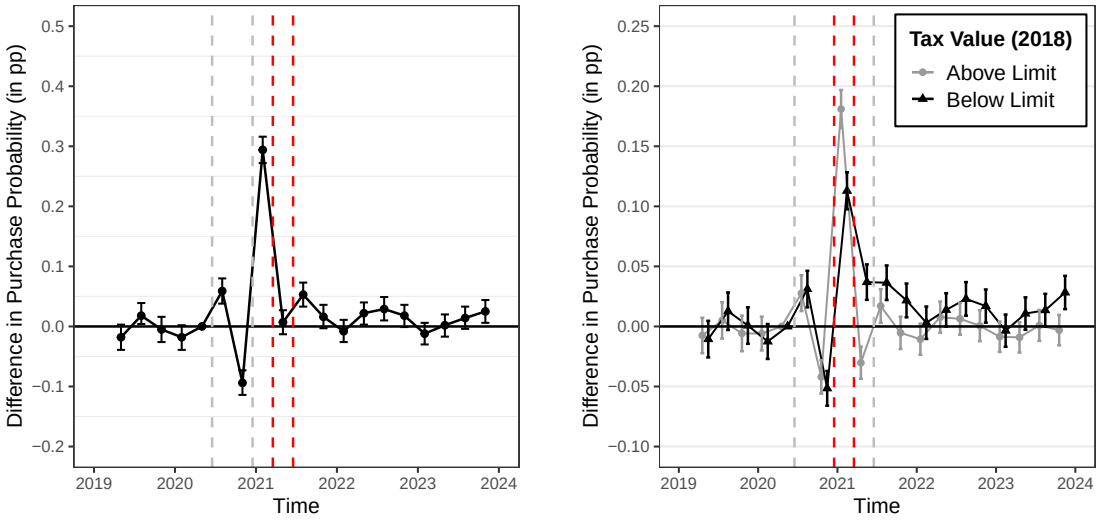


Figure A.5: The Effect of the Starter Tax Exemption by Income Quartiles

Notes: This figure shows the difference in home purchase probabilities between individuals aged 30-34 and those aged 38-42, across different household income quartiles. In each subsample, differences are normalized to their 2020Q2 level. Black vertical bars indicate 95% confidence bounds.

(a) First-Time Buyers: All Purchases

(b) First-Time Buyers by Unit Type



(c) Repeat Buyers: All Purchases

(d) Repeat Buyers by Unit Type

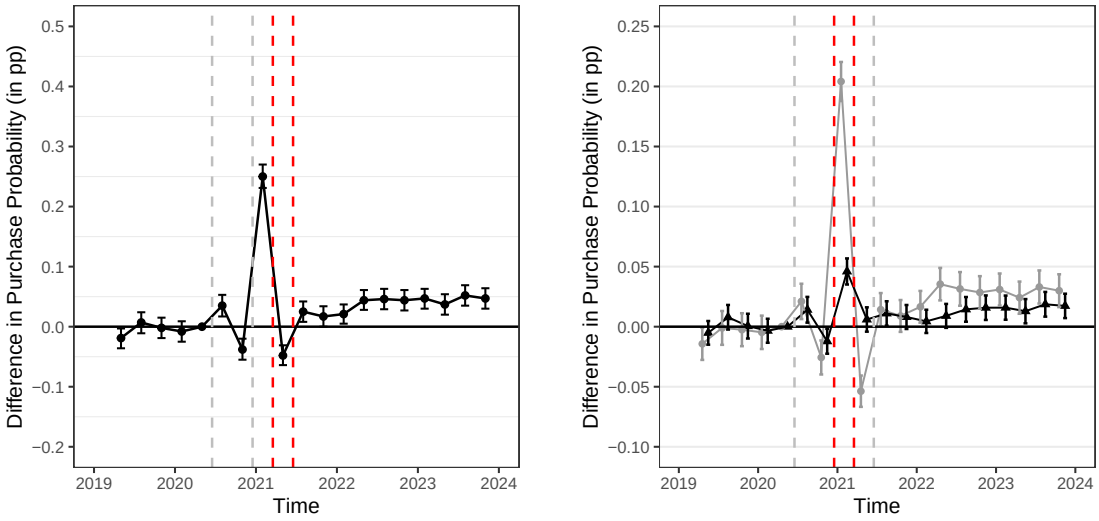


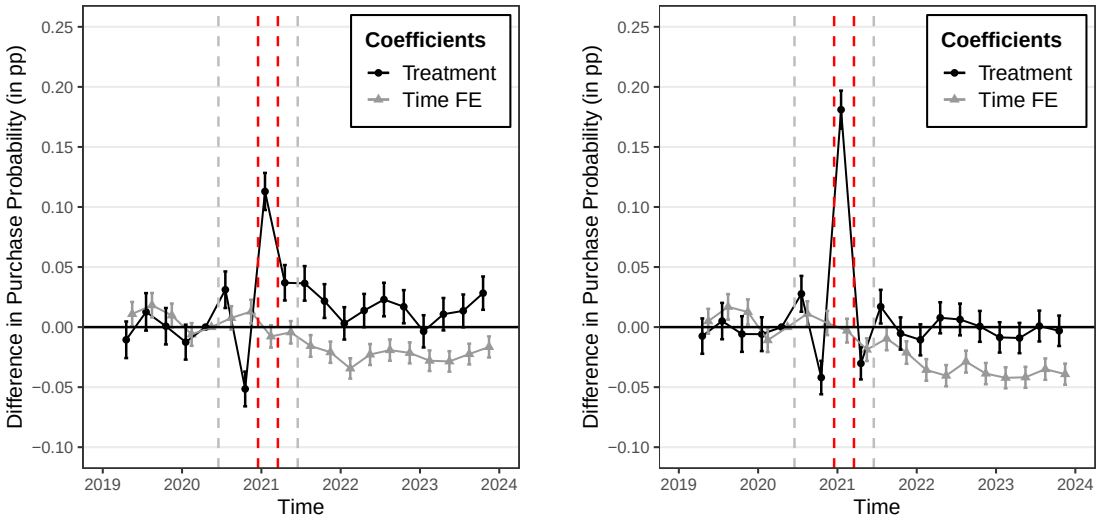
Figure A.6: The Effect of the Starter Tax Exemption on First-Time vs. Repeat Purchases

Notes: This figure shows the effect of the starter tax exemption on home purchase behavior, separately for existing homeowners (i.e., repeat buyers) and non-owners (i.e., first-time buyers). Black vertical bars indicate 95% confidence bounds. Panels (a) and (c) pool all purchases. They reveal that the medium-run increase in purchases due to the policy was largely driven by repeat buyers, consistent with Table 1. Panels (b) and (d) investigate this further, by distinguishing between purchases of low- and high-value units. The patterns are reversed for both groups. Whereas non-owners increase their purchases of low-value units in the medium run, existing owners mainly trade up to higher-value units. This suggests that their change in behavior is not a direct response to the subsidy (for which high-value units did not qualify), but is a second-order effect, whereby price increases in the lower segment allowed them to trade up.

A.3. More Details on Crowding-Out

(a) First-Time: Low-Value Units

(b) First-Time: High-Value Units



(c) Repeat: Low-Value Units

(d) Repeat: High-Value Units

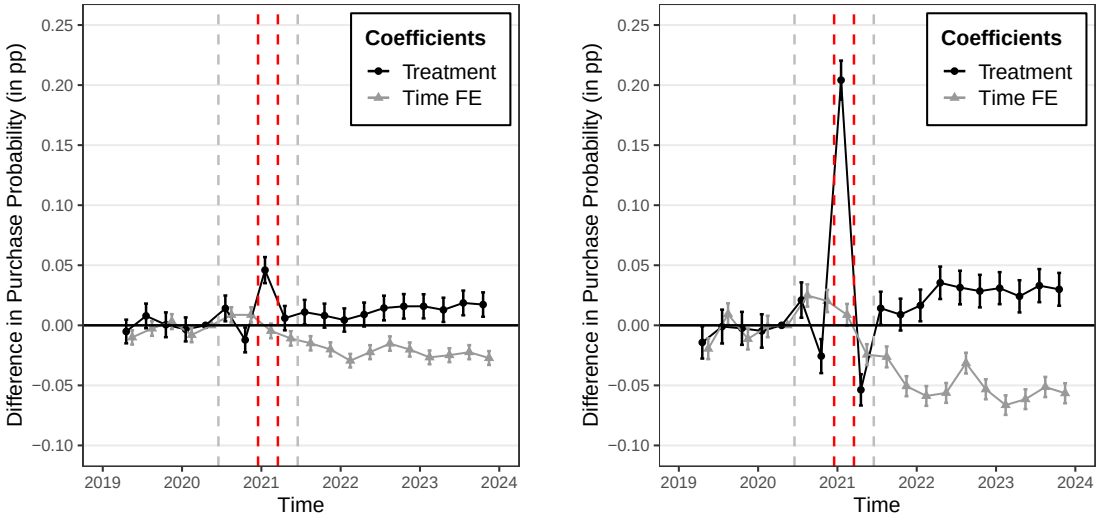


Figure A.7: Crowding-Out by Buyer Type and Housing Market Segment

Notes: This figure analyzes crowding-out more closely, by distinguishing not only by unit type (low- vs. high-value), but also by buyer type (first time vs. repeat buyers). As in Figure 2, the plots show not only estimated treatment effects, but also trace the purchase behavior of older individuals (in grey). In this group, the largest decline is found among existing owners buying high-value units (panel d). This contradicts a crowding-out interpretation of the policy, which would imply a decline in purchases among low-value units.

A.4. Additional Results for Timing and Ownership Analyses

Table A.2: The Importance of Timing Responses for the 2021Q1 Effect

	Estimate (Table A.1 column (3)) (1)	Assumed Counter-factual (2)	Implied Missing Transactions (3)	Implied Excess Transactions (4)
Pre Policy Introduction – Anticipation Period				
November 2020	–0.036	0.000	–0.036	
December 2020	–0.193	0.000	–0.193	
After Policy Introduction – Exemption Period				
January 2021	0.380	0.035		0.345
February 2021	0.060	0.035		0.025
March 2021	0.389	0.035		0.354
After Policy Introduction – Cooldown Period				
April 2021	–0.041	0.035	–0.076	
May 2021	–0.004	0.035	–0.039	
June 2021	0.005	0.035	–0.030	
			$\sum = 0.350$	

	Estimate (Table 1, column (1)) (5)	Additional Transactions ($= \sum /3$) (6)	Share of Additional Transactions (7)	Share of Retimed Transactions (8)
Exemption Period (21Q1)	0.272	0.117	42.95%	57.05%

Notes: This table illustrates our back-of-the-envelope calculation for the relative importance of time reoptimization around policy reform dates vs. extensive margin responses in explaining the strong spike in purchases in the first quarter of 2021, i.e., the exemption period. Column (1) shows monthly treatment effect estimates, which underlie the blue bars in Figure 1 and are presented in column (3) of Appendix Table A.1. As a counterfactual for this exercise, we assume that the pre-policy treatment effects for November and December 2020 would be zero if young homebuyers had not postponed their purchases to the exemption period (see the upper part of column (2)). Post-policy, we assume a counterfactual difference between both groups equal to the medium-run treatment effect as given by our medium-run estimate in column (1) of Table 1 (see lower parts of column (2)). We calculate the implied missing and excess transactions due to the policy in columns (3) and (4), respectively. The sum over these two columns gives the additional home purchases compared to our assumed counterfactual scenario. We then scale this number to three months and compare it to our exemption-period estimate from column (1) of Table 1. Columns (7) and (8) present the relative importance of additional vs. retimed transactions, respectively.

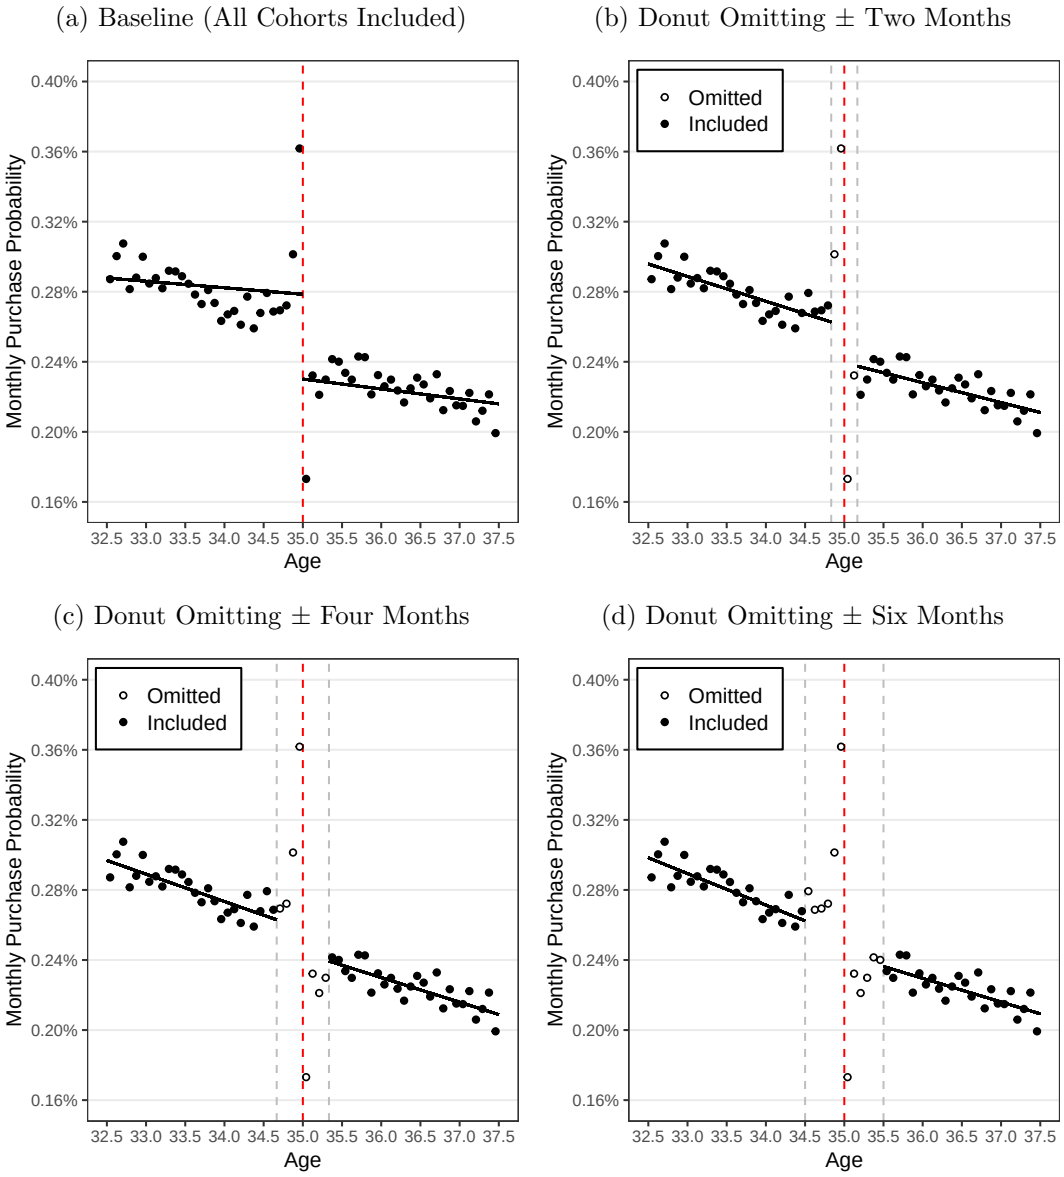


Figure A.8: Donut RD Evidence on the *Starter Tax Exemption*

Notes: This figure shows the result of estimating the linear RD model from equation (2), by excluding individuals close to the age cutoff (donut RD). Panel (a) replicates the baseline post-policy RD from Figure 3. In panels (b) – (d), two, four and six months are excluded from either side of the cutoff, respectively. In each panel, the solid black line depicts the corresponding RD estimates. Additional details on the estimation are given in the main text.

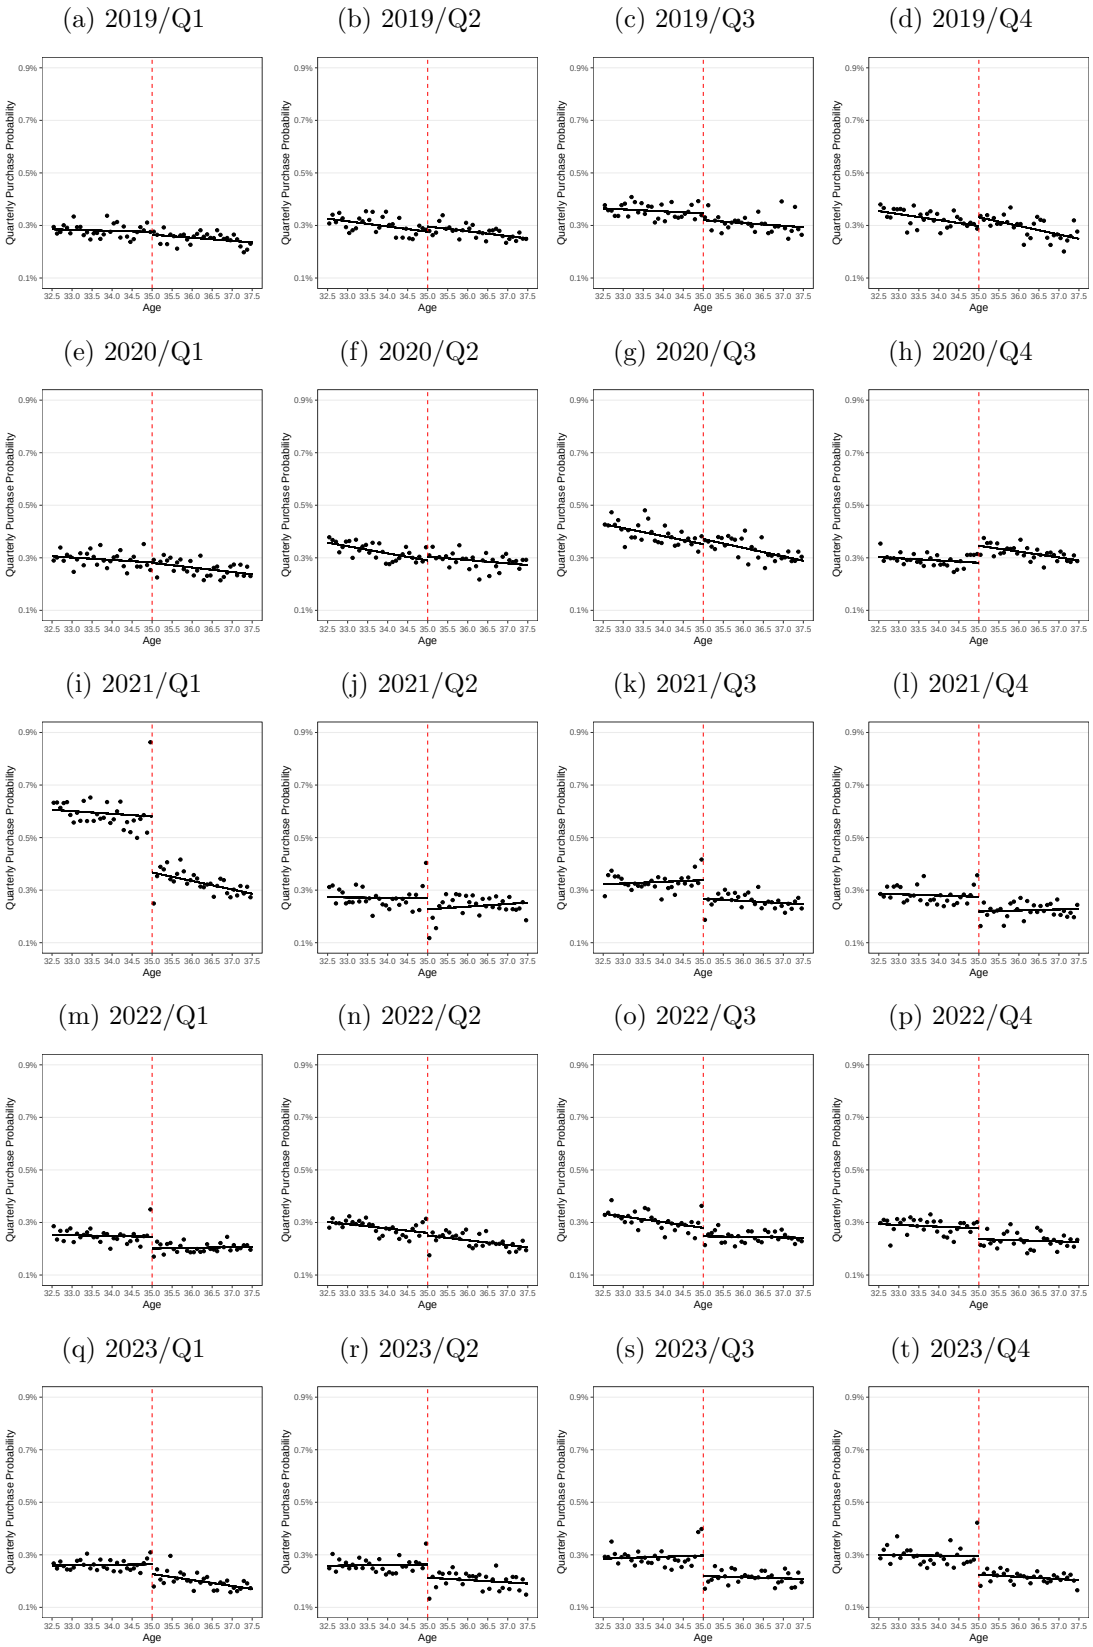


Figure A.9: Quarterly RD Evidence on the *Starter Tax Exemption*

Notes: This figure shows the results from estimating the pooled RD model from Figure 3, separately for each quarter of the sample. All details on the estimation are given in the main text.

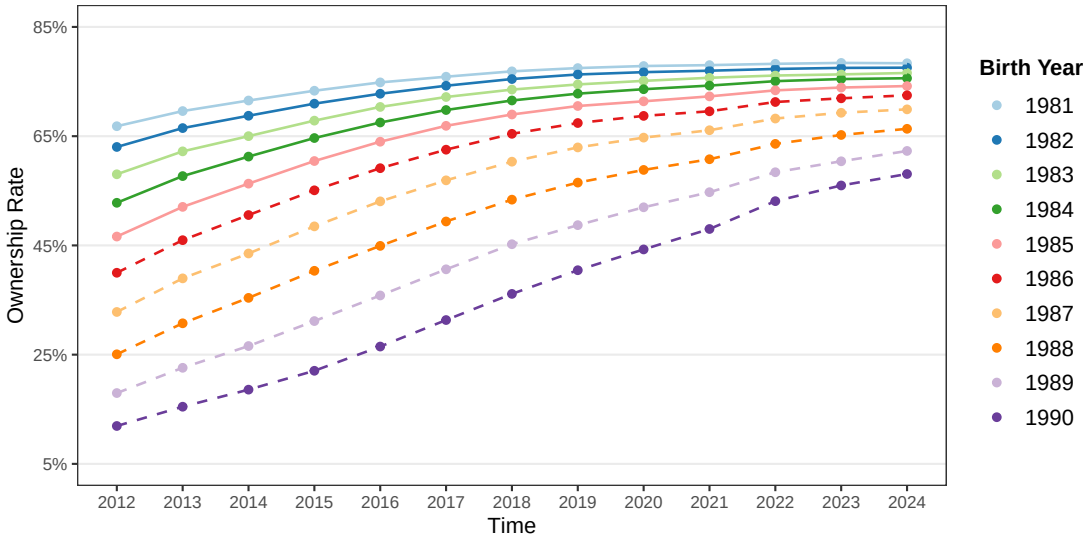


Figure A.10: The Impact of the *Starter Tax Exemption* on Homeownership

Notes: This figure descriptively tracks the life-cycle pattern of homeownership across birth-year cohorts. Dashed lines depict "ever treated" cohorts, i.e., individuals who are eligible for the tax exemption at some point in their life. Ownership information is based on household-level income records. Specifically, the plotted variable indicates whether an individual belonged to a household residing in an owner-occupied dwelling as of January 1 of a given year.

A.5. Auxiliary Results for Intensive Margin Analysis

Table A.3: Intensive-Margin Responses to the *Starter Tax Exemption*

	Sale Price (1)	Tax Value (2)	Floor Area (3)	Postcode Income (4)	Over-bidding (5)
Panel A – First-Time Buyers					
Treated ×					
Pre-Policy Period (19Q1–20Q1)	−33.8	0.005	−0.025**	−0.002	0.000
	(31.1)	(0.014)	(0.013)	(0.004)	(0.003)
Anticipation Period (20Q3–20Q4)	−25.8	0.004	−0.015	0.001	−0.001
	(34.5)	(0.015)	(0.015)	(0.005)	(0.003)
Exemption Period (21Q1)	1,293.4***	0.077***	−0.016	−0.008**	−0.003
	(45.5)	(0.017)	(0.016)	(0.005)	(0.003)
Cooldown Period (21Q2)	−0.7	−0.083***	−0.056***	−0.029***	0.006**
	(39.2)	(0.017)	(0.017)	(0.005)	(0.004)
Medium Run (21Q3–23Q4)	209.4***	−0.022**	−0.19	−0.009***	−0.001
	(29.3)	(0.013)	(0.13)	(0.004)	(0.003)
Number of Transactions	13.9mn	132k	71k	136k	71k
Panel B – Repeat Buyers					
Treated ×					
Pre-Policy Period (19Q1–20Q1)	−50.6**	−0.023**	0.008	−0.003	0.001
	(30.2)	(0.012)	(0.011)	(0.004)	(0.002)
Anticipation Period (20Q3–20Q4)	−28.9	−0.009	0.033***	−0.001	0.000
	(34.7)	(0.013)	(0.013)	(0.004)	(0.002)
Exemption Period (21Q1)	1,286.7***	0.041***	0.013	0.008**	0.001
	(48.3)	(0.016)	(0.015)	(0.005)	(0.003)
Cooldown Period (21Q2)	−312.4***	−0.137***	−0.017	−0.023***	0.013***
	(38.6)	(0.017)	(0.016)	(0.005)	(0.004)
Medium Run (21Q3–23Q4)	245.1***	−0.021**	0.041***	−0.003	−0.001
	(29.2)	(0.012)	(0.011)	(0.004)	(0.002)
Number of Transactions	17.8mn	105k	45k	108k	45k

Notes: This table summarizes the estimated intensive-margin effect of the *starter tax exemption*, separately for first-time (panel A) versus repeat buyers (panel B). Estimated treatment effects $\widehat{\beta}_k$ are based on the difference-in-differences model specified in Equation (1). Column (1) is based on the baseline person-by-quarter panel dataset. Estimates in columns (2) and (4) stem from a transaction-level dataset including all purchases from the baseline panel dataset of individuals 30–34 and those 38–42. Estimates in columns (3) and (5) are based on a subset of transactions for which we have additional realtor data. Treatment is defined either at the individual level (column (1)) or at the buyer level (columns (2)–(5)). Dependent variables are: (1) an individual's monthly home purchase volume in euros, (2) the natural logarithm of a unit's tax-assessed value as of January 2019, (3) the natural logarithm of a unit's floor area, (4) the natural logarithm of median postcode income in 2018, and (5) the log ratio between a unit's final sale price and the initial list price. Appendix Table A.3 provides additional evidence from splitting the sample between first-time and repeat buyers and estimating heterogeneous effects. Standard errors, clustered at the individual/transaction level, are reported in parentheses (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

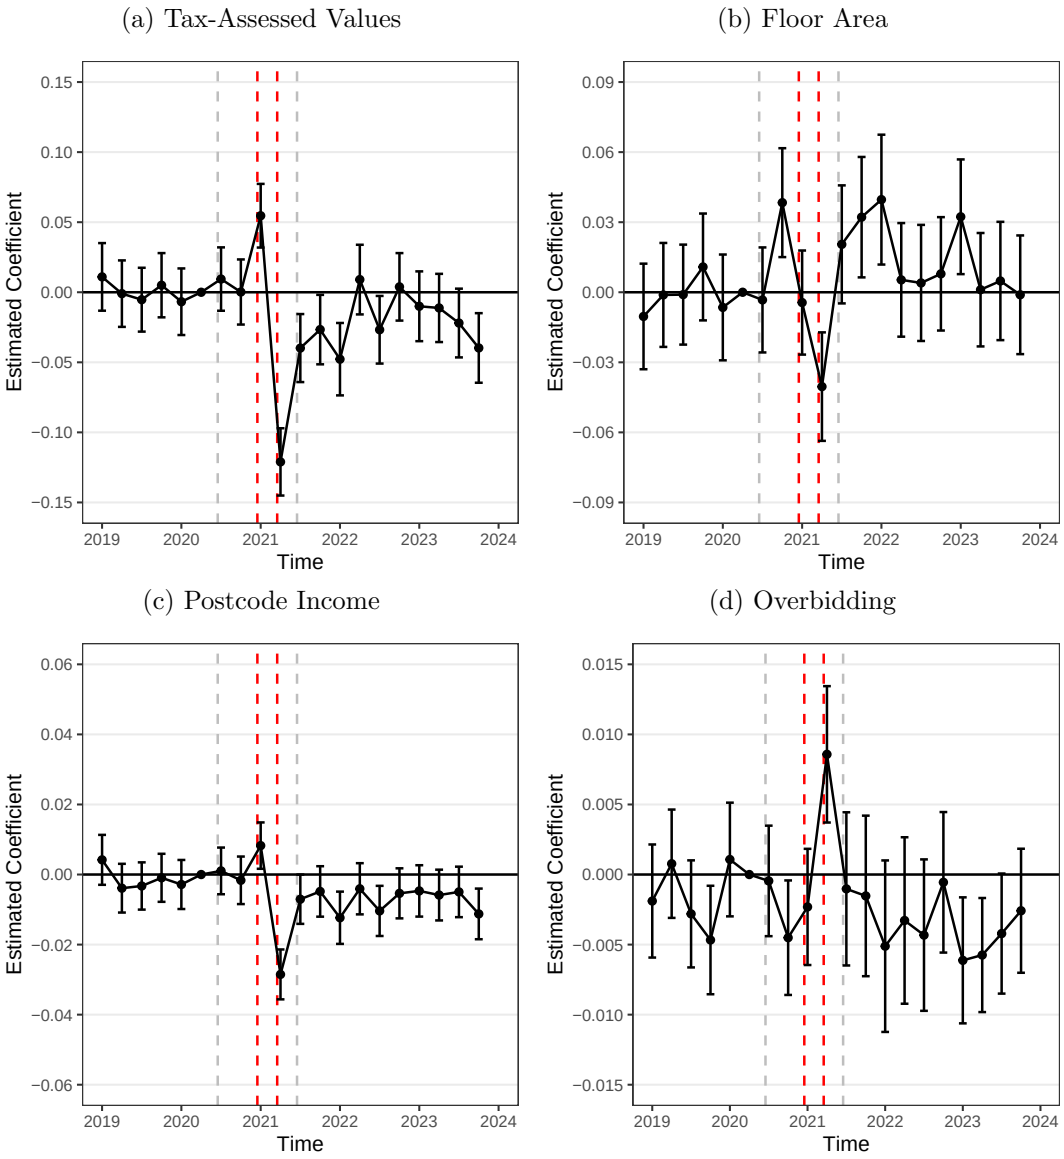


Figure A.11: Intensive-Margin Responses to the *Starter Tax Exemption*

Notes: This figure shows the effect of the starter tax exemption on homebuyers at the intensive margin. All plots are based on equation 1, estimated on the transaction sample. The dependent variables are: log of 2019 tax assessed values (panel a); log of floor size in m2 (panel b); log of postcode median income (panel c); log of sale price − log of list price (panel d). Black vertical bars indicate 95% confidence bounds.

A.6. Prices

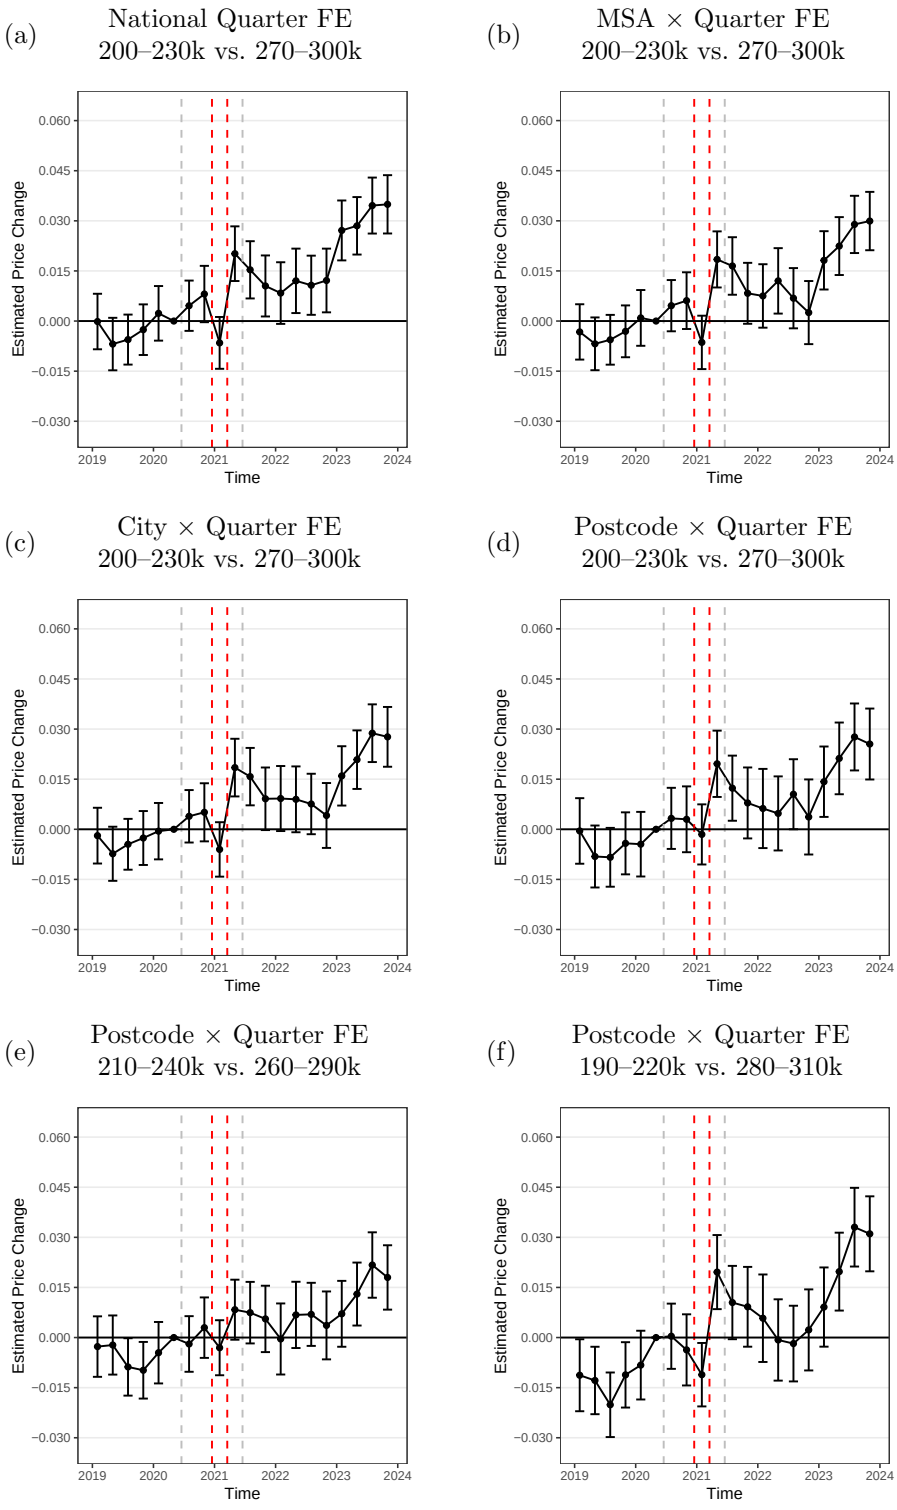


Figure A.12: Effect of *Starter Tax Exemption* on House Prices

Notes: This Figure shows the estimated effect of the *starter tax exemption* on the price tax-exempt properties, based on equation (3). 2020Q2 is the omitted quarter. The different panels correspond to columns (1) - (6) in Table 4. Black vertical bars correspond to 95% confidence bands