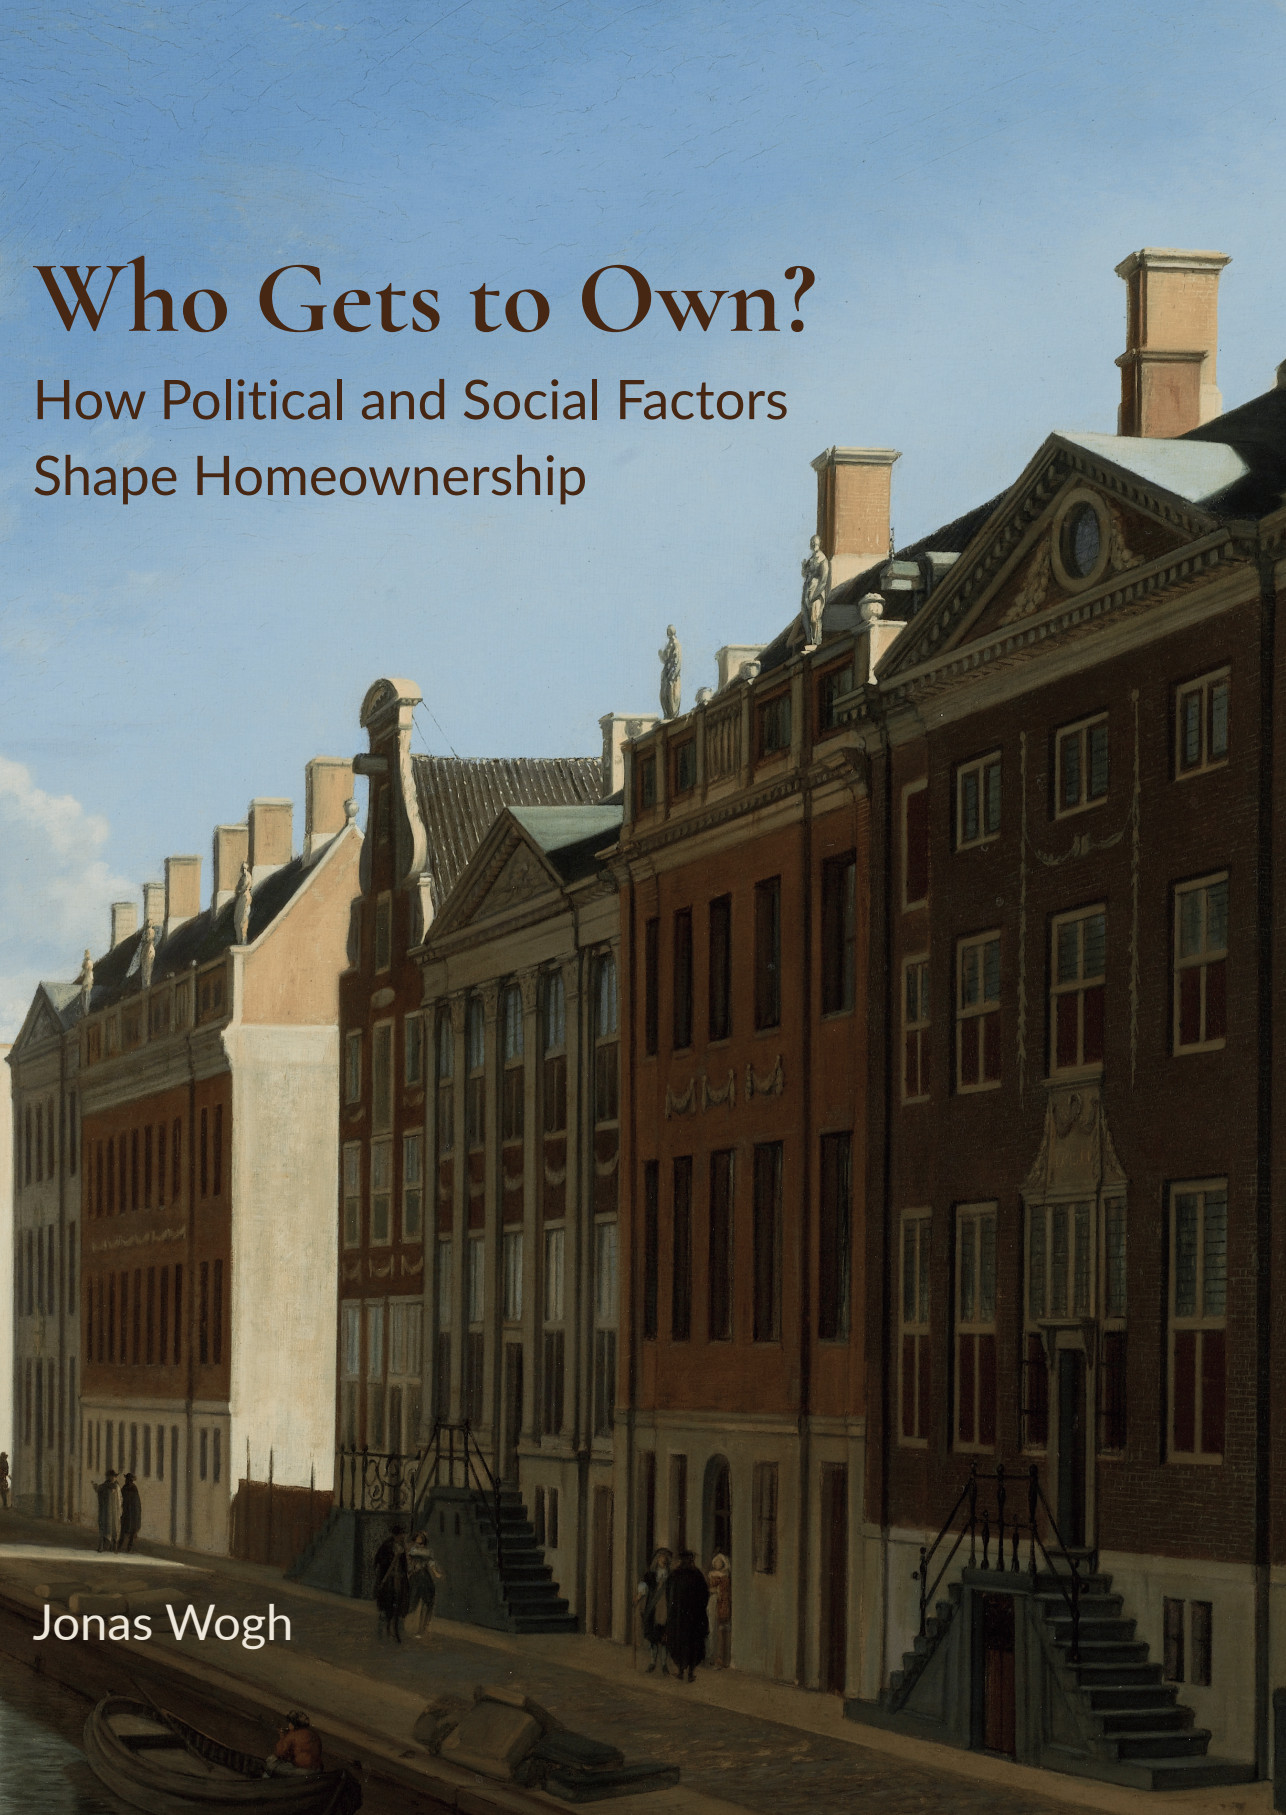
Who Gets to Own?

How Political and Social Factors Shape Homeownership



Jonas Wogh

Propositions

1. Home purchase subsidies make households speed up purchases that they were going to make anyway (Chapter 2)
2. In most developed economies, there is little room left to promote homeownership with demand-side policies (Chapter 2)
3. The inflow of parental wealth to the housing market is highly responsive to tax incentives (Chapter 3)
4. Homeowners who receive a parental transfer use it to smooth consumption, rather than just accumulate long-run wealth (Chapter 4)
5. Migrant homeowners are willing to forgo a sizable financial return to sell their home to a buyer from the same origin country (Chapter 5)
6. The most striking empirical findings do not require advanced econometric methods
7. Both activities that comprise a PhD – research and teaching – make clear the scale of the attention crisis we are in as a society
8. Writing is painful, yet unwritten ideas or results barely have any worth
9. Often the most interesting research questions in economics are those that go beyond economics
10. Interest in a research topic is nurtured by researching the topic

Doctoral thesis

WHO GETS TO OWN? HOW POLITICAL AND SOCIAL FACTORS SHAPE HOMEOWNERSHIP

Jonas Wogh

2026

© Jonas Wogh, 2026.
All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior written permission of the author.
This book was typeset by the author using LaTeX.

Cover Gerrit Berckheyde, A View of the Golden Bend in the Herengracht, 1671/72
Courtesy of the Rijksmuseum Amsterdam
Print Proefschrift Maken | proefschriftmaken.nl

WHO GETS TO OWN? HOW POLITICAL AND SOCIAL FACTORS SHAPE HOMEOWNERSHIP

Dissertation

To obtain the degree of Doctor at Maastricht University,
on the authority of the Rector Magnificus, Prof. Dr. Jan Smits,
in accordance with the decision of the Board of Deans,
to be defended in public
on 5 November 2026, at 16:00

by

Jonas Stefan Wogh

Supervisors

Prof. Dr. Jaap Bos, Maastricht University
Prof. Dr. Olivier Marie, Erasmus University Rotterdam

Assessment Committee

Prof. Dr. Dirk Brounen (Chair), Maastricht University
Prof. Dr. Asaf Bernstein, University of Colorado
Prof. Dr. Piet Eichholtz, Maastricht University
Prof. Dr. Kim Peijnenburg, Tilburg University

Parts of the research in this dissertation are enabled through financial support. I gratefully acknowledge funding from ODISSEI, MORSE and Maastricht University GSBE.

Acknowledgments

Seen from afar, doing a PhD might seem like a lonely endeavor. For four long years—five in my case—the candidate works mostly in isolation, experiencing a daily rollercoaster of emotions that people can hardly relate to, no matter how close they are. Upon closer inspection, it turns out that this description is perfectly accurate.

In my case, however, this is not the full picture. Whenever I escaped the solitude of my office, I did so to engage with people who made sure that this long journey never felt lonely. Conversations with my advisors, co-authors, colleagues, and other academics continuously inspired me and provided the motivation to keep going. Meanwhile, shared time with family and friends provided all the laughs and moments of affection that tend to get lost during solitary research work. Both of these ingredients are vital for staying sane during the PhD journey, and I therefore owe my graduation to all of you.

In 2020, I was finishing courses between Covid lockdowns and anxiously contemplating my career prospects during what I thought of as the worst labor market in history. Jaap, it was at that time that you suggested that I do a PhD under your supervision. This caught me off guard, since I had never thought myself capable of doing research, but you easily convinced me by promising that doing a dissertation would offer intellectual depth that no other career path could match. In hindsight, I am truly thankful for your persistence and for seeing something in me that I hadn't seen in myself at the time. With this in mind, I can look past the fact that you thereby pushed my graduation into a labor market that turned out much worse than the one I feared in 2020.

You chose to team up with Olivier as my second advisor, based on his track record in the field we were planning to do work on. While the choice of topics turned out quite differently, this was a perfect match nevertheless. Jaap, you allowed me to grow at my own pace by supporting each of the steps I took to become an independent researcher, while always being there when I needed you. I am most thankful for your seemingly endless willingness to discuss all aspects of academic life with me. Olivier, you are a force of nature, providing the constant flow of new ideas that keeps research projects alive. Daunted at first by your pace, I increasingly learned to match it (or, at least, adapt to it) until, at some point, I was actively looking forward to every upcoming project

meeting.

While I first thought that working with my advisors had exposed me to the whole range of approaches to doing research, I was quickly corrected as I started new collaborations. Nils, you introduced me to a level of determination that I had not seen before in academia. You taught me to work in a goal-oriented way, by breaking projects down into manageable pieces and executing them with a timeline in mind. Without this, I would still be debating the research question for our first project, rather than writing the acknowledgements for a finished dissertation. For that reason, and because you have guided me along every milestone of the PhD, I see you as a third advisor much more than a co-author.

Max, you have shown me the extent to which perseverance and attention to detail are rewarded in research projects. I have easily spent more time in meetings with you than with all my other co-authors combined. While some of those meetings probably could have been shorter, had you not distrusted every standard error I ever reported, I genuinely appreciate the amount of time we dedicated to discussing our project. Not only did I learn an enormous amount from you, but I also came to realize that exploring a topic together and building up expertise by challenging each other's ideas is how doing research is most rewarding.

Many others deserve gratitude for giving me generous advice on specific projects or general questions. Piet, Dirk, and Juan, you were those whom I consulted most frequently — my questions were mostly half-baked, while your answers still provided great insight. Thomas, I enjoyed chatting with you on a regular basis, since your always-open door was usually too inviting not to stop by for a few minutes. Among the many others whose advice I am thankful for are Peiran, Peter, Pomme, Rob, and Rachel.

Besides co-authors and advisors, there is a large group of people who made working and living in Maastricht enjoyable. Flavio and Samara, while our *trio infernal* was short-lived due to our itinerant lifestyle, I cherish the unforgettable moments we shared together anywhere between Maastricht and Melbourne. Flavio, moving in together has proven to be a wonderful decision, even though my bad influence has eroded your cooking skills over the years. I am proud of the parties we threw together, and convinced that the emotional support we gave each other over the years was crucial to help both of us finish. Sam, you brought unprecedented glamour to Maastricht and it was clear from

the start that you were not going to stay for long. The friendship we managed to maintain nonetheless is one of the happiest achievements of the past few years, and some of our joint experiences were surreal in every sense of the word. I look forward to my next trip Down Under.

Hugo and Robin, sharing an office together was the best possible start, even though the total absence of productivity in our *WWF headquarter* is the reason I now tell incoming candidates that it's totally normal to not get any work done in the first half year of the PhD. Linde, never have sassiness and compassion been combined quite the way you combine them. In addition, it is remarkable how you quietly became a great researcher, while we seemingly spent most of our time during the PhD venting and joking around. Martijn, I think it's a shame that we were not part of the same cohort, as I can only imagine how much fun it would have been to travel to conferences together. Your humor saved many lunch breaks, but it took me a while to realize that you are much more than an aspiring comedian. This finally became clear when you visited me in Boulder for some unforgettable hikes (sorry!), thoughtful discussions and remarkable nights out. Anna, our friendship developed toward the end of both of our PhDs—at a time when it was most needed. I appreciate the long conversations we had in our kitchen while Flavio cooked for us, as well as the yoga sessions that you made me join week after week.

Alex, you brought the easy-going calm that our department desperately needed, and I could always count on leaving every chat with you with a smile on my face. Stefan, your no-nonsense attitude to academic life was most refreshing and I was fascinated by the sophisticated workouts that you explained to me at UM Sports week after week. Bin, you are easily the kindest person I have ever met, and I admire the way you overcame every hurdle that was put in front of you during your PhD while remaining calm and composed. Dongxiao, you were my favorite conference buddy, though I can't shake the suspicion that you took secret pleasure in watching me suffer through Chinese food that you promised would not be spicy at all. Many more fellow PhDs and lecturers deserve to be mentioned here, including Stefany, Philibert, Frederique, Daniele, Enver, Kimon, and André. Finally, I am immensely thankful to Francien and Monique, who did so much more than just their jobs, helping me with endless patience whenever I faced any administrative problem.

Two colleagues from Maastricht to whom I owe a special mention are Paulo and Marlene. Paulo, I believe that you are the most important person at the

department. While much of the academic world is dominated by self-seeking careerism, you stand for the opposite: a kind willingness to sacrifice your own time in order to let others grow, not asking for anything in return. I cannot count the number of times that I have abused this by running into your always-open office to bombard you with questions. In addition, starting the days with a coffee together while geeking out on obscure historical facts remains one of my favorite parts of the PhD journey. Marlene, I have great respect for how quickly you have established yourself in a position that was quite challenging, something I witnessed first-hand as your office mate. That you were willing to share your grant to fund an extension of my PhD—basically asking for nothing in return—has been incredibly generous and it has helped me greatly. I hope you know that I will always remain grateful for it.

Without a doubt, my research visit at the University of Colorado has been the pivotal experience of my PhD. Asaf, I owe this almost exclusively to you, since you demonstrated to me so graciously how beneficial it is to be the dumbest person in the room. Week after week, you shared your knowledge, allowing me to get a taste of how research at the frontier of the field is done. Being able to return to Boulder for my post-doc feels like coming full circle. Of all the other great people whom I have met during my time in Boulder, those who most deserve a special mention are Mattie and Armin. You are the main reason why this time has not only been academically fulfilling, but also non-stop fun. Thank you for welcoming me right from the start and for all the special moments we shared hiking in the Rockies, cheering for the Buffs, and super-sizing ourselves at C4C.

Much of the credit for helping me stay sane during those past years belongs to my friends from Munich, who mean so much to me. Though most of you aren't living next door anymore, the memories from these last years alone are too numerous to count: almost becoming Mormon in Utah with Marc; indulging in culture and wine in Frankfurt with Michi; celebrating carnival in our new hometown with Andrea; raving to techno beats somewhere in eastern Germany with Domi; turning nights into days with the Kraxel-crew; and so many more. I am grateful that we keep those friendships alive and active regardless of the distance.

The final words belong to those who have given me the most. To Theresia and Michael, who have been the best parents I could have ever asked for. We don't really do pride, but I hope that you will be proud of yourselves for instilling in me all the traits and values without which I could not have done this. And

to Julia, whose unconditional love is the greatest gift. We bring out the best in each other, and I love the life that we are creating for ourselves.

Jonas Wogh
Cologne, August 2026

Contents

1 Introduction 1
- 1.1 The Rise of Homeowner Societies . . . 2
- 1.2 Homeownership in the Netherlands . . . 4
- 1.3 Access to Homeownership Today . . . 6
- 1.4 The Economics of Homeownership . . . 7
- 1.5 Homeownership Beyond Economics . . . 10
- 1.6 Structure of This Dissertation . . . 11

2 Supporting Young Homebuyers 15
- 2.1 Introduction . . . 17
- 2.2 Institutional Setting . . . 22
- 2.3 Data . . . 24
- 2.4 Reduced-Form Evidence . . . 25
- 2.5 Theoretical Framework . . . 40
- 2.6 Welfare Simulation . . . 45
- 2.7 Conclusion . . . 49
- A Appendix . . . 52

3 Frontloading Wealth: How Parental Transfers Shape Housing and Savings 67
- 3.1 Introduction . . . 69
- 3.2 Institutional Setting and Data . . . 73
- 3.3 Parental Transfers for First-Time Home Purchases . . . 77
- 3.4 Home Purchase Transfers and Future Wealth Accumulation . 90
- 3.5 Conclusion . . . 108
- A Appendix . . . 110

4 Bidder Beware: Intergenerational Wealth Transfers in the Residential Housing Market 121
- 4.1 Introduction . . . 123
- 4.2 Institutional Setting and Policy Design . . . 125
- 4.3 Identification Strategy . . . 128
- 4.4 Empirical Results . . . 135
- 4.5 Spillovers to Local Housing Markets . . . 141
- 4.6 Conclusion . . . 148
- A Appendix . . . 150

5 Ethnic Matching in the Housing Market 153
5.1 Introduction 155
5.2 Data and Institutional Background 160
5.3 Ethnic Matching in Housing Transactions 163
5.4 Price Differentials in Within-Ethnicity Transactions 171
5.5 Underlying Mechanism 181
5.6 Conclusion 189
A Appendix 192

Bibliography 197

Summary 211

Samenvatting 213

Impact 215

About the author 217

List of Figures

1.1 Wealth Comparison: Owners vs. Renters 8

2.1 The Effect of the *Starter Tax Exemption* on Home Purchases . 26
2.2 Testing Crowding Out by Comparing Market Segments 30
2.3 Regression Discontinuity Evidence on the *Starter Tax Exemption* 32
2.4 Welfare Simulation of the *Starter Tax Exemption* 50
2.5 Unadjusted Home Purchase Rates Among Younger and Older Individuals 52
2.6 Homeownership Rates and Growth in the US 52
2.7 Income and Liquid Assets: 2019 vs. Current 54
2.8 The Effect of the Starter Tax Exemption by Liquid Asset Quartiles 55
2.9 The Effect of the Starter Tax Exemption by Income Quartiles . 56
2.10 The Effect of the Starter Tax Exemption on First-Time vs. Repeat Purchases 57
2.11 Crowding Out by Buyer Type and Housing Market Segment . . 58
2.12 Donut RD Evidence on the *Starter Tax Exemption* 60
2.13 Quarterly RD Evidence on the *Starter Tax Exemption* 61
2.14 The Impact of the *Starter Tax Exemption* on Homeownership . 62
2.15 Intensive-Margin Responses to the *Starter Tax Exemption* . . . 64
2.16 Effect of *Starter Tax Exemption* on House Prices 65

3.1 Timeline of the Special Exemption Limit for Housing-Related Transfers 75
3.2 Parental Transfers Around Major Life Events 80
3.3 Net Parental Transfers Around Major Life Events 82
3.4 Parental Bank Deposits Around Major Life Events 83
3.5 Transfers for Home Purchases Along the Distribution of Liquid Assets 88
3.6 Heterogeneity in Net Transfers Around First-Time Home Purchases 89
3.7 The Causal Relationships Around Transfers, Housing, and Wealth 94
3.8 Housing Market Outcomes Around the Policy Reform 98
3.9 Transfers and Wealth Accumulation (IV Estimates) 104
3.10 Parental Transfers Around Major Life Events (Combined Model) 110
3.11 Household Bank Deposits Around Major Life Events 111

3.12 Parental Transfers Around Major Life Events: Event Studies Removing Observations After Second Event 112
3.13 Home Purchase Transfers Across Years 113
3.14 Regular (Non-Tax-Exempt) Transfers Around Major Life Events 114
3.15 Transfers Around First-Time Home Purchase (by Parental Assets Quartile) 115
3.16 Simulated Transfer Wealth From Reduced Mortgage 116
3.17 Transfers and Wealth Accumulation (OLS Estimates) 117
3.18 Debt Service Cost Over Time: Interest Expenses and Principal Repayment 118

4.1 Parental Liquid Assets and Wealth Transfers 135
4.2 Share of Transfer Recipients over Time 143
4.3 Price Spreads in Exposed Markets 145
4.4 Price Spreads in Exposure Deciles 147
4.5 Tax-Assessed Values in Exposed Markets 148
4.6 Price Spreads in Exposed Markets: Excluding Transfer Recipients 151

5.1 Share of Homebuyers with Migration Background 161
5.2 Matching Preferences by Ethnic Group 170
5.3 Preferences for Ethnic Matching Across Groups 172
5.4 Price Discounts by Ethnic Group 180
5.5 Price Discounts in Arm's-Length Transactions 186
5.6 Placebo Test: Random Allocation of Matching Dummy 194
5.7 Price Discounts and Neighborhood Composition 196

List of Tables

2.1 The Effect of the *Starter Tax Exemption* on Home Purchases . 28
2.2 Extensive Margin vs. Timing Responses to the *Starter Tax Exemption* 34
2.3 Intensive-Margin Responses to the *Starter Tax Exemption* . . . 36
2.4 The Effect of the *Starter Tax Exemption* on House Prices . . . 40
2.5 Monthly Difference-in-Differences Estimates 53
2.6 The Importance of Timing Responses for the 2021Q1 Effect . . 59
2.7 Intensive-Margin Responses to the *Starter Tax Exemption* . . . 63

3.1 Summary Statistics 78
3.2 Transfers and Housing Outcomes 101

4.1 Summary Statistics: Treated vs. Untreated Buyers 130
4.2 Effect of the Wealth Transfer on Bidding Outcomes 137
4.3 Robustness Analysis 139
4.4 Underlying Mechanisms 142
4.5 IV First Stage: Predicting Wealth Transfers 150

5.1 Summary Statistics 164
5.2 Cross-Tabulation of Buyer and Seller Migration Background . . 165
5.3 Ethnic Matching and Transaction Outcomes 178
5.4 Ethnic Matching and Subsequent Returns 182
5.5 Price Discounts: Underlying Mechanism 185
5.6 Ethnic Matching Discounts: Trust vs. In-Group Favoritism . . 189
5.7 Ethnic Matching and Price Differentials 193

1

Introduction

In the spring of 1835, John Batman and his company of grazier colonists set sail from Van Diemen's Land (nowadays known as Tasmania) to the southern coast of the Australian mainland. Their group, which was later extended to a formal business association, embarked on their trip with a clear objective: to acquire a large tract of land from the elders of the Wurundjeri and the Boonwurrung people, two Aboriginal nations who had been living in these lands for tens of thousands of years. After a few meetings between Batman's group, which involved seven Aboriginal interpreters, and the elders of the two groups, a treaty was signed. It settled the purchase of around 600,000 acres of land—on which the city of Melbourne was founded subsequently—in exchange for a number of everyday items, including scissors, knives and handkerchiefs (Attwood, 2015).

What strikes the modern-day reader of this episode in Australia's dark history as a particularly unfair trade agreement actually reflects a more fundamental issue. While the people of the Aboriginal nations thought of themselves as *belonging* to the land they inhabited, they did not claim to be *owning* it. In fact, the concept of land ownership—or any form of private ownership—did not exist for them. Instead, humans were seen as the custodians of the land, with an obligation to care for it (NARvis, 2021).

1.1 The Rise of Homeowner Societies

The Aboriginal relationship to land was completely at odds with the views of John Batman and other settlers who, rooted in their European origin, viewed private property as a core tenet of civil society. This centrality of private property was not always the case, however, and reflected fundamental intellectual shifts in the early-modern period. While rights over land had existed in medieval Europe under the feudal system, they were understood as privileges that could only be granted by the monarch, whose authority rested on divine sanction (Baker, 2019).[1] This was challenged by European philosophers of the early-modern era, most prominently by John Locke, who argued that property rights were natural rights that existed independently of government sanction and required no royal approval (Locke, 1690). Locke saw the government's sole responsibility as protecting individual property rights that had existed prior to it (West, 2003).

Locke and others laid the intellectual foundation that promoted the view of land and property as assets that could freely be traded, even though this remained the exclusive privilege of a small elite of noblemen, clerics, and, increasingly, wealthy citizens (Habakkuk, 1940). Only in the 19th century did land ownership start to emerge as a broad societal phenomenon. In the still nascent United States, the Homestead Act of 1862 offered 160 acres of previously federal land to applicants who were willing to cultivate it (Arrington, 2012). In the more urbanized societies of Western Europe, promoting homeownership among the working class was suggested as a countermeasure to the social hardship caused by the industrial revolution (Kohl, 2018). However, substantial differences emerged across countries during that time, determined by their respective institutional features. In the United States, homeownership rates in large cities had reached 25% by 1890, whereas they remained at 13% in Germany by 1912 (Kohl, 2014).

It is only in the second half of the 20th century that homeownership became a mass phenomenon. During this time, the existing differences between the United States and Europe were further reinforced. In the United States, the transition to a *homeowner society* was accomplished at a rapid pace in the postwar period. Since the housing stock remained intact during the war, rising demand in the following years led to an unprecedented increase in homeown-

[1] In pre-medieval societies, property rights varied considerably, ranging from state-controlled land in ancient Egypt, to individual ownership in Mesopotamia, communal ownership in tribal societies, and sophisticated private property law in Rome.

ership. This was the result of the postwar economic upturn, which led to the emergence of a broad middle class, and targeted policies aimed at promoting homeownership, especially among returning veterans and their families (Fetter, 2013). By the late 1960s, homeownership rates in the U.S. had almost reached 65%, an increase of nearly 20 percentage points since the end of the war. Since then—over a period of more than five decades—ownership rates have not deviated from this level by more than two percentage points. The sole exception to this occurred prior to the global financial crisis, when a large-scale expansion of mortgage credit caused homeownership rates to approach 70%, a short-lived increase that was painfully reversed immediately after the collapse of the housing market.

In Western Europe, which saw many of its major cities destroyed, housing policy in the postwar period was focused on reconstruction to provide shelter, rather than promote ownership (Diefendorf, 1993). The subsequent rise in homeownership came later and more unevenly than in the United States, reflecting the institutional differences across countries. A common feature across most Western European countries was large-scale investment into social housing in the aftermath of the war. However, differences in housing policy emerged thereafter, which had a lasting effect on homeownership rates. These divergent paths reflected not only existing institutional differences, but variation in political ideology with respect to homeownership. In the United Kingdom and, to a lesser extent, in Scandinavian countries, homeownership was increasingly viewed as the desirable tenure form. As a result, starting in the 1970s and 1980s, governments in those countries supported mortgage market liberalization and targeted policies to support ownership. One example is the UK's 1980 "Right to Buy" scheme, introduced by the Thatcher government, which allowed social tenants to purchase their occupied units at a significant discount (Arundel and Ronald, 2021).

In contrast, German-speaking countries remained *nations of tenants* (Kohl, 2014), where ownership rates grew much more slowly. The reasons for this divergence are not necessarily found in different attitudes toward neoliberal policies, but rather in weaker support for homeownership as the ideal tenancy form. For instance, while the German government did liberalize social housing in the late 1980s, it did not follow the UK's example of transforming social housing units into owner-occupied units (Cooper and Kurzer, 2023). Finally, Southern European countries can clearly be defined as homeowner societies, in which ownership rates were consistently high throughout the 20th century, rather than undergoing the rapid postwar expansion of some European neigh-

bors (Kohl, 2017). This different pattern reflects the structural organization of Southern European welfare systems, in which extended family networks, rather than the state or formal mortgage markets, served as the primary channel for housing provision, making owner-occupancy the default tenure long before it became a policy objective elsewhere (Allen, 2006).

1.2 Homeownership in the Netherlands

Though the Netherlands can be described as a homeowner society today—with ownership rates around 60%—this is a relatively recent trend. Throughout the first half of the 20th century, ownership rates did not exceed 30%. Instead, the defining institutional feature of the Dutch housing market was a growing social housing sector, offering rental units below market prices. This originated in the 19th century, when wealthy capital owners established housing associations, with the aim of developing affordable rental units for the urban working class. Initially, these were private institutions, funded by shareholders who only retained a small amount of the earnings and reinvested the rest into additional construction. In the early 20th century, the government entered this sector by offering low-interest loans and subsidies to social housing associations, in exchange for increased regulatory oversight. In addition, municipal governments started establishing housing associations themselves, constructing new units and helping to establish social housing as a central pillar of the Dutch housing market, rather than a fringe institution reserved only for the poor working class (Beekers, 2012). Despite this, social housing units accounted for less than 15% of the total housing stock throughout the first half of the 20th century (Van Kempen and Priemus, 2002).

With around a quarter of its building stock destroyed or damaged, the Dutch housing market was ravaged by World War II. As in other European countries, the postwar reconstruction period led to an extraordinary expansion of the social housing sector, driven by massive public intervention. Of the two million dwellings completed between 1945 and 1970, the majority were social units, which raised the share of households living in social housing to more than 35%, a historically high share (Van Deursen, 2023). Since then, in line with broader trends towards market liberalization after 1980, government intervention in the social housing sector was gradually reduced, allowing housing associations greater independence and market orientation.[2] Nonetheless, social housing

[2]The liberalization of the Dutch social housing sector was not uniform. While the 1980s

remains an essential part of the Dutch market today, still accounting for almost 30% of all residential units (Van Deursen, 2023).

At the same time, the second half of the 20th century saw a steady rise in homeownership. From relatively low levels in the postwar period, ownership rates increased gradually from around 30% in 1960 to almost 60% in the early 2000s. Broadly, this was driven by the same combination of rising household incomes and increasing market liberalization observed in other countries as well. Beginning in the late 1980s, the financialization of mortgage markets further supported the expansion of homeownership. Financial liberalization and increased competition among lenders led to greater availability of credit, higher loan-to-value ratios, and longer maturities, easing borrowing constraints for households. These developments, combined with generous mortgage interest tax deductibility, made owner-occupied housing both more attractive and more accessible, thereby contributing to the rise in homeownership (Elsinga et al., 2016).

By 2008, aggregate mortgage debt levels had reached close to 100% of GDP, the highest level of all European countries (Ronald and Dol, 2011).[3] Securitization of mortgage loans was also widespread. In addition, many homeowners had low or even negative home equity—with outstanding mortgage debt exceeding the current value of their home—since loan-to-value (LTV) ratios exceeding 100% were permitted at the time. Consequently, the Dutch market was seriously exposed to the effects of the global financial crisis, which had originated in the U.S. mortgage market. Between 2008 and 2013, nominal (real) house prices declined by more than 20% (30%) (Boelhouwer, 2017). This worsened the balance sheets of many homeowners. While default rates remained at relatively low levels—due to strict recourse laws in the Dutch market—transaction volumes almost halved, from 200,000 in 2007 to 110,000 in 2013. The government's reaction to this crisis was twofold. On the one hand, it enacted a series of macroprudential measures to ameliorate household indebtedness, including tighter mortgage lending standards (Rouwendal and Petrat, 2022) and increased tax incentives for amortizing mortgages (Bernstein and Koudijs, 2024). On the other hand, it stimulated housing market activity

and 1990s saw a retrenchment of subsidies and greater financial autonomy for housing associations, later reforms reintroduced some regulation. From the 2000s onwards, policies tightened the targeting of social housing and imposed regulatory constraints, partly in response to fiscal pressures and EU state-aid rules.

[3] In contrast, mortgage-debt-to-GDP ratios in the UK, France and Germany were around 80%, 50% and 40%, respectively.

by reducing real estate transfer taxes. Partly due to these fiscal measures, and due to expansionary monetary policy, the Dutch housing market subsequently recovered and house prices rose consistently for the following decade after 2013.

1.3 Access to Homeownership Today

In recent years, accessibility has become a central concern across developed economies that prioritize homeownership, including the Netherlands. More than a decade of rising house prices following the global financial crisis, combined with stricter mortgage market regulation, has threatened the ability of younger households to become homeowners (Vangeel et al., 2023). For instance, the average age of first-time homebuyers in the Netherlands increased by more than a year between 2010 and 2020 (Tweede Kamer, 2020).[4]

This has led governments to introduce targeted purchase subsidies designed to lower the barriers for first-time buyers. These take various forms—including transfer tax exemptions, interest-free loans, and outright cash grants—but typically aim at reducing the upfront down payment required for a home purchase. Often, they are targeted not only at first-time buyers, but are also subject to additional conditions related to buyer income or age, as well as property price or location. For instance, the Spanish government offers cash grants of up to €10,000 for homebuyers under 35 who purchase units located in municipalities with fewer than 10,000 inhabitants, in order to promote homeownership and revitalize rural areas.

At the same time, the burden of enabling homeownership has increasingly shifted toward the family, even in countries with broad mortgage availability, with parental wealth playing an ever more prominent role in determining who can and cannot access the market. The most direct form in which parental wealth can facilitate homeownership is through direct cash transfers. In a U.S. survey, 36% of young homebuyers reported relying on transfers and inheritances in 2024 (Redfin, 2024). Other forms of parental support include mortgage co-signing agreements (Allen et al., 2026) or direct intra-family sales (Colmsjö, 2024). However, an obvious concern with the growing reliance on family wealth is its highly unequal distribution.

[4] Part of the increase in the average age at housing market entry may also be driven by demand-side factors, including changing societal norms regarding fertility and marriage.

1.4 The Economics of Homeownership

The most common argument in favor of homeownership is a financial one. Owning a home is widely seen as a key strategy for accumulating wealth. Indeed, Margaret Thatcher made an economic case for promoting homeownership in a speech introducing the 1980 Right to Buy policy: "The government believes that [homeownership] should be fostered. It reflects the wishes of the people, ensures a wide spread of wealth through society, encourages personal desire to improve and modernize one's home, enables people to accrue wealth for their children, and stimulates the attitudes of independence and self-reliance that are the bedrock of a free society" (House of Commons, 1980).

At first glance, this argument seems to find support in the data. This is illustrated in Figure 1.1, which compares median net wealth—defined as the difference (in €) between all assets and all outstanding debt—for Dutch households who own a home compared to those who do not. It shows that, even though homeowners' wealth varies depending on the housing market environment, it far exceeds that of renters, most of whom own no more than a few thousand euros. However, wealth differences between owners and renters, regardless of their magnitude, do not prove that homeownership promotes wealth accumulation in a causal way. There are two main reasons for this. One is reverse causality: since home purchases require an upfront investment, it is likely that the sample of existing homeowners is skewed towards those who were wealthier prior to their purchase. The second relates to unobservable differences between homeowners and renters. It is conceivable that the same character traits or preferences that induce people to build up savings in general also make them more likely to become homeowners. The distinction between these confounding factors and a true causal effect of homeownership is very subtle, but highly important from the perspective of a social planner when evaluating policies to stimulate homeownership.

By abstracting from some real-world complexities, the financial benefit of homeownership can be evaluated in a simple framework. Consider two households who live in homes that are exactly identical. However, whereas one household owns their unit—financed by taking out a mortgage—the other household rents theirs. Both households begin with identical wealth D, which one household uses as a down payment on a home of value $H = D + L$, where L denotes the size of the mortgage. The other household invests D in a liquid equity investment, such as an index fund or a mutual fund. The owner's

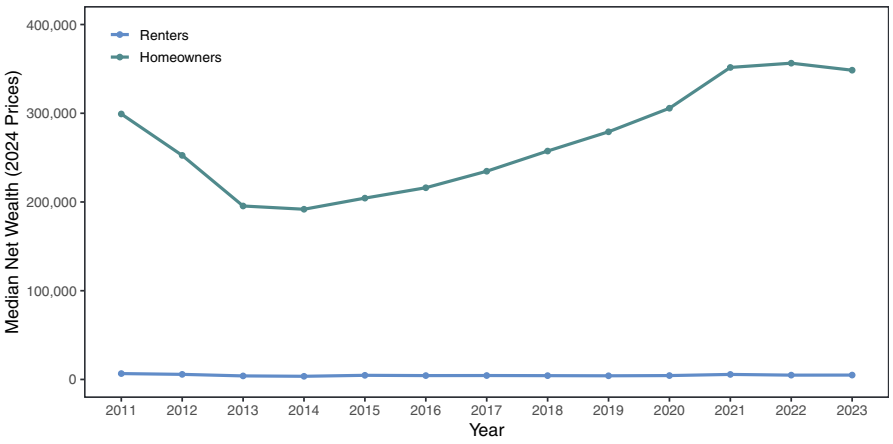


Figure 1.1: Wealth Comparison: Owners vs. Renters

Notes: This figure compares the median net wealth of homeowning vs. renting households in the Netherlands between 2011 and 2023. Net wealth is the sum of all assets minus all liabilities, including mortgage debt. All values are inflation-adjusted and expressed in 2024 prices.

mortgage is assumed to be interest-only, with interest rate r_m, such that her total housing cost is $r_m L$.[5] The renter pays rent R per period. To simplify the exposition of wealth dynamics, assume that both households earn a periodic income that is exactly equal to the owner's housing cost $r_m L$. Wealth after T periods is then straightforward to derive. The owner's wealth equals the appreciated house value net of the outstanding mortgage,

$$W_O = H(1 + r_g)^T - L, \tag{1.1}$$

while the renter's wealth equals the compounded value of his initial equity investment plus additional periodic investments, which are financed by reduced housing cost relative to the owner:

[5]Interest-only mortgages only require principal repayment at the end of the term. In practice, most modern-day mortgages are partly or fully amortizing, meaning that the outstanding debt amount decreases over time. While this increases the complexity of the model, it leaves the fundamental intuition unchanged.

$$W_R = D(1+r_f)^T + \Delta \frac{(1+r_f)^T - 1}{r_f}. \tag{1.2}$$

Here, $\Delta \equiv r_m L - R$ is the periodic housing cost differential between owning and renting and r_f denotes the return of the equity investment. The owner outperforms financially if and only if $W_O > W_R$, a condition that depends mainly on three factors: house price growth r_g, the mortgage rate r_m, and the financial market return r_f. If house prices grow at a faster rate than equity prices ($r_g > r_f$), homeowners tend to accumulate wealth more quickly than renters and vice versa. Similarly, high mortgage interest rates will increase the annual housing costs for owners, relative to renters.

This stylized model omits a number of factors that are important in practice, including maintenance costs, the tax deductibility of interest payments and property taxes. In a more realistic model, the relationship between house prices and rents would depend on these factors (Himmelberg et al., 2005). Nonetheless, the simple model is able to capture a key insight: from a financial perspective, owning a home does not necessarily dominate renting. In fact, renters may be able to build wealth faster than owners, if they achieve a higher return on their equity investment. Historically, which of the two assets yields higher returns has been widely debated, but recent evidence suggests that their long-run average returns are broadly similar (Jordà et al., 2019).

In this simple comparison of owners and renters, all initial wealth D is invested—either as a down payment, or in a liquid financial investment. In practice, this is often not the case. While owners must commit their initial wealth, because down payments are required by mortgage lenders, renters have greater flexibility about how to allocate their savings. This flexibility is associated with lower saving rates and underinvestment in equity markets, despite the fact that these assets have historically delivered returns comparable to or exceeding those of housing. Arguably, the key mechanism behind higher wealth accumulation of owners is therefore not the superior performance of housing investments, but the fact that mortgages—especially those that amortize over time—require households to build equity that they might otherwise not accumulate (Bernstein and Koudijs, 2024).

1.5 Homeownership Beyond Economics

The preceding discussion suggests that the economic case for homeownership may be less clear-cut than is often assumed. However, proponents of homeowner societies often allude to additional, less tangible, benefits that arise from owning a home. This ideological framing is captured in a speech given by President George W. Bush advocating governmental support to promote homeownership among racial minorities: "I believe there is such a thing as the American Dream. [...] Owning a home is a part of that dream, it just is. Right here in America if you own your own home, you're realizing the American Dream" (Bush, 2002).

Motivated by this notion—which is shared by legislators from other countries and with different political orientations—a large amount of academic research has examined the non-financial benefits of homeownership. Broadly, most of these benefits are related to two aspects: residential stability and investment incentives (i.e., "skin in the game"). On the one hand, transaction costs associated with buying and selling homes imply that homeowners tend to remain in their homes longer than renters.[6] This may in turn open up more economic opportunities in local neighborhoods due to stronger social ties (Rohe et al., 2002) and improve child development (Haurin et al., 2002). However, these early studies do not rely on (quasi-)experimental variation, raising concerns about their causal interpretation.

On the other hand, the fact that homebuyers invest into an immovable property provides an incentive to maintain and upgrade both the property itself, as well as the immediate neighborhood. Consistent with this, homeownership has been associated with higher investments into property and neighborhood maintenance (DiPasquale and Glaeser, 1999), increased voting behavior (Hausman et al., 2022) and lower crime rates (Disney et al., 2023).

Critics of homeownership tend to use the same two arguments—residential stability and investment incentives—but highlight their negative implications. For the former, the most common concern is that ownership may have adverse labor market consequences, as "locked-in" homeowners have lower labor market mobility (Oswald, 1999). For the latter, critics argue that increased political participation due to homeownership entails serious negative side-effects,

[6] An additional factor contributing to this is *mortgage lock-in*, whereby households with negative home equity are unable to move because they cannot afford the down payment required to purchase a new home.

as existing homeowners have a strong incentive to limit local housing supply by voting against new construction (Koster and Rouwendal, 2024).

Due to a variety of private and societal outcomes that may be affected by homeownership, and the near impossibility of generating causal estimates, a conclusive answer about its net benefits will likely remain elusive. By extension, there cannot be a definitive argument about the merit of homeowner societies vis-à-vis nations of tenants. Nevertheless, access to homeownership remains on top of the political agenda across developed economies.

1.6 Structure of This Dissertation

The preceding historical overview takes a bird's-eye view, one in which homeownership is shaped by long-run economic growth, financial institutions and historical path-dependency. Consequently, this macro-perspective leaves little room for current legislators to promote homeownership in the aggregate. Empirically, it appears to be vindicated by the fact that in recent decades, homeownership rates in developed economies remained largely constant, in spite of active government intervention. By looking only at averages, however, one cannot understand a distribution. Similarly, by narrowly focusing on aggregate homeownership rates, one cannot answer the question that is pressing in many modern-day housing markets: *who gets to own?* That question is at the heart of my dissertation.

It starts from the observation that while aggregate homeownership rates may be stable, various societal groups' market access can still fluctuate—and therefore be subject to outside influence.[7] The focus of the dissertation is therefore to understand how various external forces shape access to owner-occupied housing. It is the result of four empirical studies of the Dutch housing market, most of which I have conducted in collaboration with other researchers. The core idea that binds together the different chapters is that owner-occupied homes are highly unusual assets, which can hardly be fully understood through the lens of individual-level financial optimization. Put simply: the decision to buy or sell a home is not made in a vacuum. Instead, it is influenced by various agents beyond buyers and sellers; in turn, the behavior of buyers and sellers is

[7] For instance, while homeownership rates among young households have declined in the Netherlands, those among retiree households have increased. These trends may offset each other from an aggregate perspective, but have important societal repercussions nonetheless.

influenced by various factors beyond financial motives. Some of these factors will be investigated in this dissertation.

In chapter 2, I study political influence in the housing market. Concretely, I analyze the effect of a fiscal subsidy that was introduced by the Dutch government in order to promote access to homeownership among young households. In chapters 3 and 4, I evaluate the role of parental wealth. In chapter 3, I show that wealthy parents make substantial financial transfers to help their child make a first-time purchase, and study the downstream financial consequences of these transfers. In chapter 4, I zoom in on the direct effects of parental transfers on housing market outcomes, by studying their relationship with bidding behavior. Chapter 5 studies the role of ethnicity in housing market transactions. I find that homebuyers and sellers are not indifferent to the identity of their counterparty, and that sellers exhibit a strong preference for finding a co-ethnic buyer. Finally, a summary of all chapters, together with a discussion of their broader impact, is given at the end of the dissertation

2

Supporting Young Homebuyers

This chapter is co-authored with Max Löffler (Maastricht University).

Abstract

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, benefiting 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.

2.1 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

[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.

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 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-differences 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% 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 et al. (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 et al. (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, there are 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 et al. (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 et al. (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 homeownership, 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 et al. (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 holding 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 2.2 details the institutional background. Section 2.3 introduces the data used. Section 2.4 presents the empirical results. Section 2.5 introduces the theoretical framework to evaluate the utility consequences of subsidies for young homebuyers. Section 2.6 simulates welfare consequences, based on the reduced-form empirical parameters. Section 2.7 concludes.

2.2 Institutional Setting

2.2.1 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 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 intergenerational wealth transfers related to housing (Bos et al., 2024), mortgage assistance schemes, and restrictions on buy-to-let transactions in certain municipalities (Francke et al., 2023).

2.2.2 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 2% 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 was used for owner-occupied living, whereas the applicable tax rate for buy-to-let transactions increased concurrently from 6% to 8%, 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 2% 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, in-

tending 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.

2.3 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

[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.

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.

2.4 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 2.4.1, we first analyze the policy's impact on home purchase behavior. Section 2.4.2 presents regression discontinuity analyses of timing responses. We evaluate intensive-margin responses and bidding behavior in Section 2.4.3. Finally, we study the policy's effect on aggregate house prices in Section 2.4.4.

2.4.1 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 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 2.1, we plot the difference in monthly purchase probabilities be-

[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.

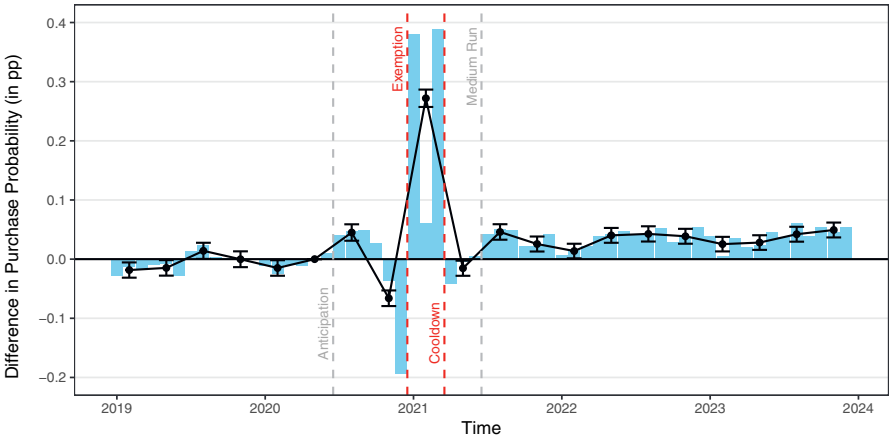


Figure 2.1: The Effect of the *Starter Tax Exemption* on Home Purchases

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 (2.1). 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.

tween 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 fewer

[4]Before the reform, the younger group is around 0.15 pp more likely to purchase a home in any given month. Appendix Figure 2.5 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 2.6 for corresponding evidence based on US data).

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:

$$Y_{it} = \sum_{k \neq 20Q2} \beta_k \cdot \mathit{Treat}_i \cdot 1_{(t=k)} + \gamma_t + \delta \cdot \mathit{Treat}_i + \varepsilon_{it}, \tag{2.1}$$

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 2.1 (in black). To summarize the policy's impact, we further aggregate the estimates for the five subperiods indicated in the figure. Column (1) of Table 2.1 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

Table 2.1: 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 (2.1). 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: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

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 reestimate Equation (2.1) via sample splits for each population quartile.

The resulting estimates are reported in columns (2)–(5) of Table 2.1. We first note that pre-trends become somewhat unstable when cutting along these margins.[6] After 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

[6]Note that, despite these precisely estimated small deviations before the reform, the quarter-by-quarter estimates exhibit clear trend breaks 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 temporal stability of these population quartiles relative to their 2019 distributions, to Appendix A.2.

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 2.1. 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 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 (2.1) 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

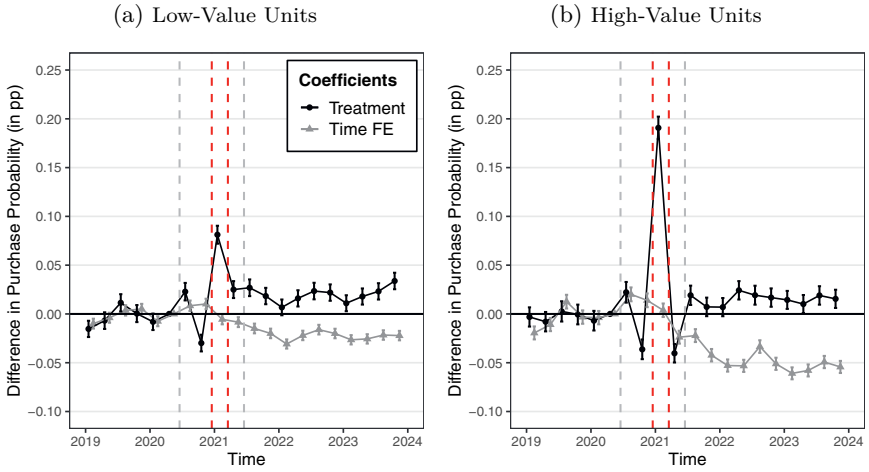


Figure 2.2: Testing Crowding Out by Comparing Market Segments

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 (2.1). However, in contrast to Figure 2.1, 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 38–42. Vertical bars indicate 95% confidence bounds.

in Section 2.4.4.

In Figure 2.2, 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.

2.4.2 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.

The first piece of evidence on extensive margin responses comes from the split between first-time and repeat purchases. As documented in Table 2.1, 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 2.4.1, 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% 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 2.3 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 subsidy is introduced and the price limit is in place, we

[7] The details of this calculation can be found in Appendix A.4, Table 2.6.

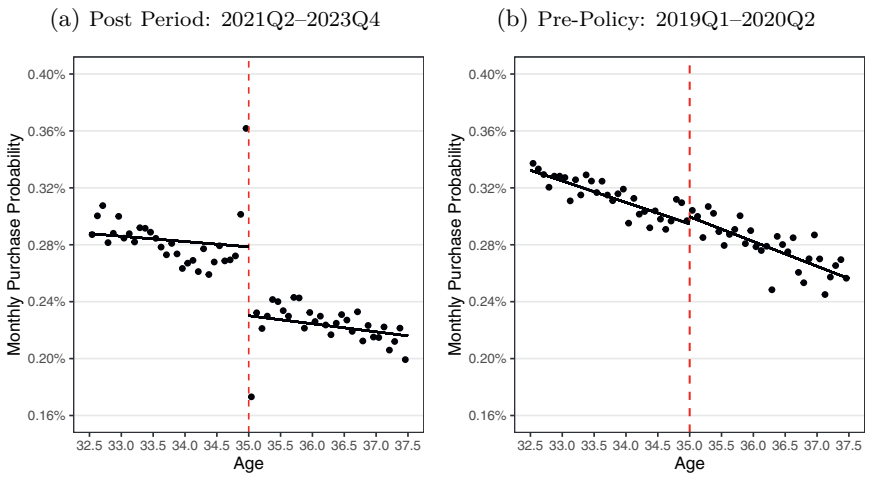


Figure 2.3: Regression Discontinuity Evidence on the *Starter Tax Exemption*

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 2019Q1 until 2020Q2. The solid black lines represent estimates from a linear RD model as specified in Equation (2.2). Corresponding estimation results are provided in Table 2.2, panel A. Complementary visualizations for regression-discontinuity models that exclude specific age-in-month groups around the threshold are presented in Appendix Figure 2.12 ("donut RDs").

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:

$$Y_{it} = \alpha + \beta \cdot 1(Age_{it} < 35) + \gamma \cdot Age_{it}^{(35)} + \delta \cdot 1(Age_{it} < 35) \cdot Age_{it}^{(35)} + \varepsilon_{it}, \tag{2.2}$$

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 2.2. 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 2.3, 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 2.2, columns (2)–(4); corresponding donut RD plots are relegated to Appendix Figure 2.12). We conclude that around 80% 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 2.13 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-

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

Table 2.2: 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 (2.2), 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), recentered 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: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

month as the running variable in the annual snapshots of the homeownership register. Panel B of Table 2.2 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 exceed-

ing 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, Statistics Netherlands 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 2.14).

Definitively 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.

2.4.3 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 (2.1) for various transaction characteristics. Table 2.3 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 2.4.1), 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,

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

Table 2.3: 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 2.15. Estimated treatment effects $\widehat{\beta}_k$ are based on the difference-in-differences model specified in Equation (2.1). 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 unit's final sale price and the initial list price. Appendix Table 2.7 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: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

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 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.

2.4.4 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 2.4.1, 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 particularly problematic in the presence of bargaining around the cutoff (Slemrod et al., 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:

$$\ln P_{it} = \sum_{k \neq 20Q2} \beta_k \cdot \mathit{Treat}_i \cdot 1_{(t=k)} + \gamma_{c(i),t} + \delta \cdot \mathit{Treat}_i + \varepsilon_{it}, \tag{2.3}$$

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 2.4. 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. Of the units in the lower price range, 95% 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 the likelihood that treated (control) units will eventually transact above (below) the price limit. In columns (5) and (6) of Table 2.4, 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 analysis 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 a 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 2.4.1 and visualized in Figure 2.11. 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.

Table 2.4: 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 2.16 visualizes the corresponding quarterly estimates. Estimated treatment effects $\widehat{\beta}_k$ are based on the difference-in-differences model specified in Equation (2.3), which compares 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: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

2.5 Theoretical Framework

We now turn to the efficiency and welfare implications of the young homebuyer relief programs. The empirical analysis has revealed that additional transactions caused by the policy predominantly reflect a short-term timing response. This motivates a simple theoretical framework, in which utility changes are only governed by the policy's effect on prices, whereas behavioral responses have no first-order effect. Section 2.5.1 presents the basic model setup. Section 2.5.2 derives the utility consequences of the starter tax exemption. Section 2.5.3 discusses implications for efficiency and distributional aspects.

2.5.1 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 et al., 2015).

Household Problem.

Individual i maximizes utility subject to the budget constraint:

$$y_i + R \geq c_i + r(h_i), \tag{2.4a}$$

which applies to both renters and homeowners, provided they do not 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 $s \geq 0$:

$$y_i + R \geq c_i + r(h_i) + (1 + \tau - s \cdot 1_{i \in \mathcal{J}})p(h). \tag{2.4b}$$

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:

$$y_i + R + p(k) \geq c_i + r(h_i), \tag{2.4c}$$

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 exogenously 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).

2.5.2 Subsidies for Homebuyers

Against the backdrop of the institutional setting and the empirical results from Section 2.4, 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

[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 2.6.

thus be expressed as:

$$\Delta c^{\mathrm{NT}} := \frac{\mathrm{d}u_i/\mathrm{d}s}{\partial u_i/\partial c} = \frac{\mathrm{d}R}{\mathrm{d}s} \qquad \forall i \in \mathcal{I}^{\mathrm{NT}}, \tag{2.5a}$$

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:

$$\Delta c^{\mathrm{AB}} := \frac{\mathrm{d}u_i/\mathrm{d}s}{\partial u_i/\partial c} = p \cdot 1_{i\in\mathcal{J}} - (1+\tau-s)\frac{\mathrm{d}p}{\mathrm{d}s} + \frac{\mathrm{d}R}{\mathrm{d}s} \quad \forall i \in \mathcal{I}^{\mathrm{AB}}, \tag{2.5b}$$

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 $(c' + \Delta c^{\mathrm{AB}}, h')$, 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:

$$\Delta c^{\mathrm{AS}} := \frac{\mathrm{d}u_i/\mathrm{d}s}{\partial u_i/\partial c} = \frac{\mathrm{d}p(k)}{\mathrm{d}s} + \frac{\mathrm{d}R}{\mathrm{d}s} \qquad \forall i \in \mathcal{I}^{\mathrm{AS}}. \tag{2.5c}$$

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 would not 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 $\Delta c^{\mathrm{AB}} - \epsilon$, and thus almost as much as for always buyers. The last switcher considers buying only for a large compensation of $\Delta c^{\mathrm{AB}} - \epsilon$, 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 $(\Delta c^{\mathrm{NT}}, \Delta c^{\mathrm{AS}})$.

These utility consequences are governed by a 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.

2.5.3 Efficiency and Distribution

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

$$\mathrm{d}W/\mathrm{d}s := \sum_i g_i \frac{\mathrm{d}u_i/\mathrm{d}s}{\partial u_i/\partial c} = \sum_i g_i \Delta c_i > 0, \tag{2.6}$$

with exogenously given, fixed social marginal welfare weights denoted by g_i (Saez and Stantcheva, 2016). Equation (2.6) 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 2.4.3. 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.

2.6 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 2.5. Then, we implement the baseline simulation, using the estimated parameters from Section 2.4 and discuss sensitivity analyses.

2.6.1 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 2.5. 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 (2.5a)–(2.5c), using the reduced-form estimates from Section 2.4. 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 forgone transfer tax revenue due to the exemption of (always) buyers, and divide it by the size of the adult population. This

[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 2.1, 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 probability in 2020Q2 among individuals in the first income quartile, which was equal to 0.197%.

corresponds to a per-person tax rebate change of −€20.[13] For new buyers and sellers, we draw utility changes from a uniform distribution, bounded between the utilities of never takers and those 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 et al., 2017): for 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.

2.6.2 Simulation Results

On average across the entire population, we find that the transfer tax exemption for young homebuyers induces a 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 transactions that would not have occurred 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 this simulation are visualized in Figure 2.4. 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.

[13] New buyers are not considered in the calculation of forgone 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 pensioners). Evidence for this is seen in panel (c): at the bottom of the distribution, the share of observed homesellers increases.

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 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 nor 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 2.4. 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 consists 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 (2.5a). 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 does this pattern reverse, 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.

2.7 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 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

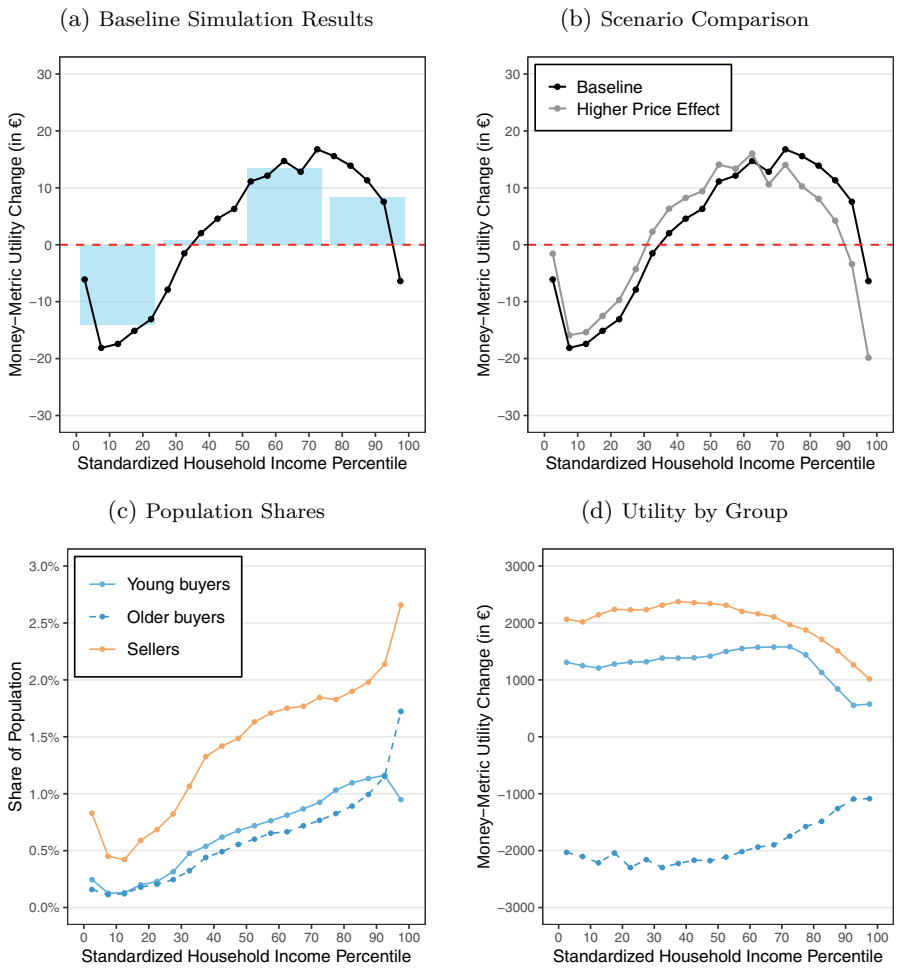


Figure 2.4: Welfare Simulation of the *Starter Tax Exemption*

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 $\frac{dR}{ds} = -€20$.

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.

A Appendix

A.1 More Details on Home Purchases

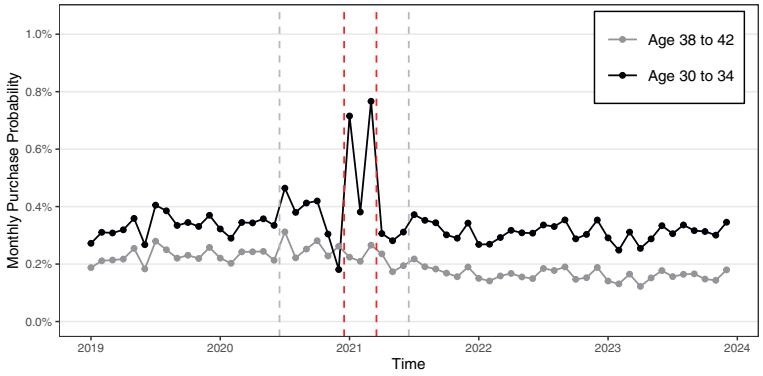


Figure 2.5: Unadjusted Home Purchase Rates Among Younger and Older Individuals

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-differences model visualized in Figure 2.1.

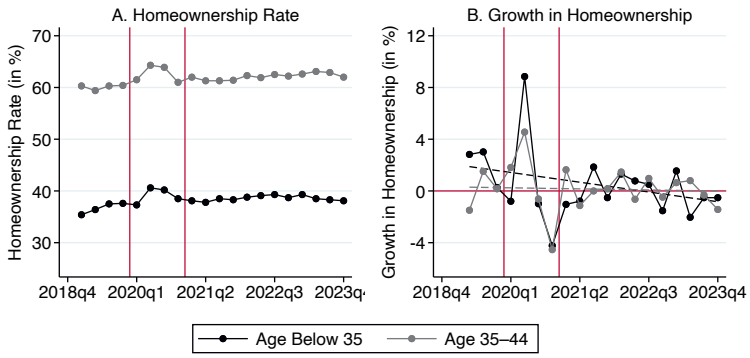


Figure 2.6: Homeownership Rates and Growth in the US

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 time-series data comes from the Current Population Survey (Flood et al., 2025).

Table 2.5: 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	0.322	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	0.324	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	0.323	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	0.325	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	0.320	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	0.325	0.213	−0.009	September 2022	0.354	0.190	0.052
April 2020	0.323	0.213	−0.011	October 2022	0.288	0.147	0.029
May 2020	0.357	0.214	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	0.218	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	0.326	0.180	0.054

Notes: This table reproduces monthly difference-in-differences estimates from Figure 2.1 (blue bars). Columns (1) and (2) of each row reproduce the raw monthly purchase probability per group, as shown in Figure 2.5. 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 quartile assignments with those in later years (2020–2023) at the individual-year level. As shown in Figure 2.7, most individuals remain in the same quartile over time. We therefore proceed by classifying individuals based on their 2019 income and liquid asset levels.

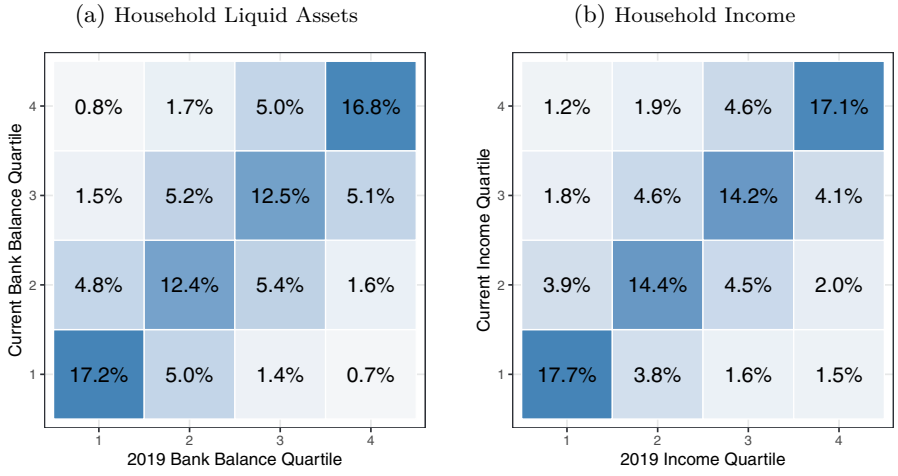


Figure 2.7: Income and Liquid Assets: 2019 vs. Current

Notes: This figure shows confusion matrixes that compare individuals' 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.

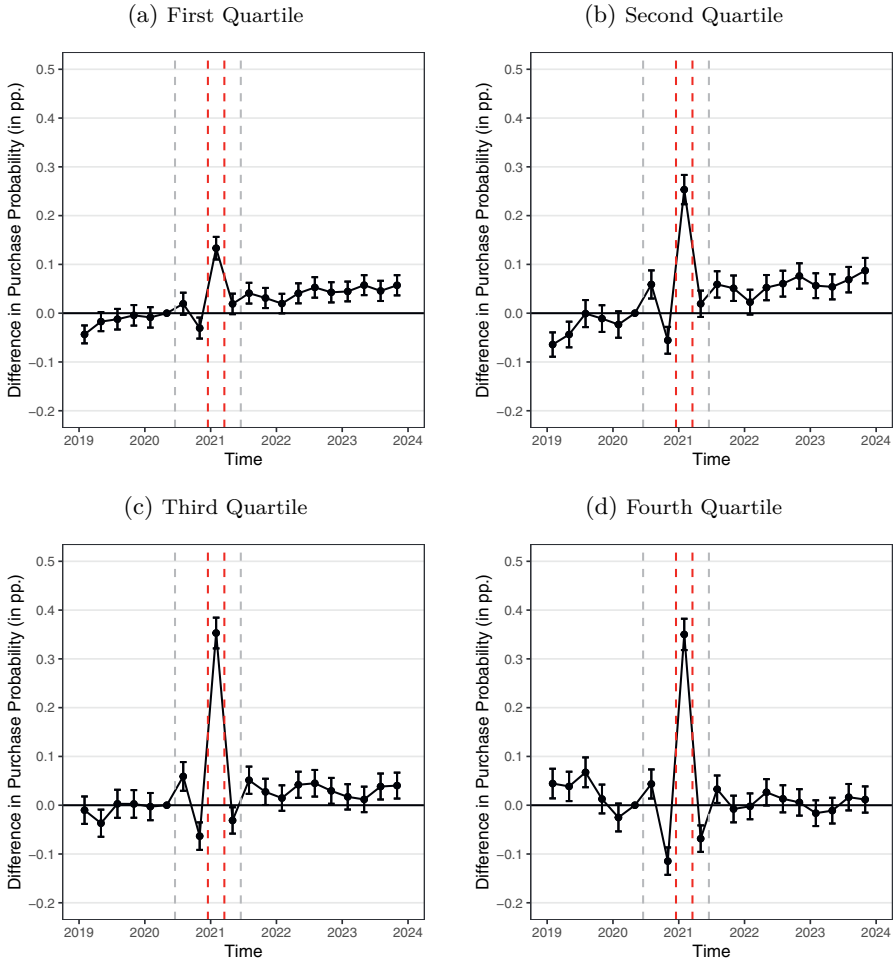


Figure 2.8: The Effect of the Starter Tax Exemption by Liquid Asset Quartiles

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.

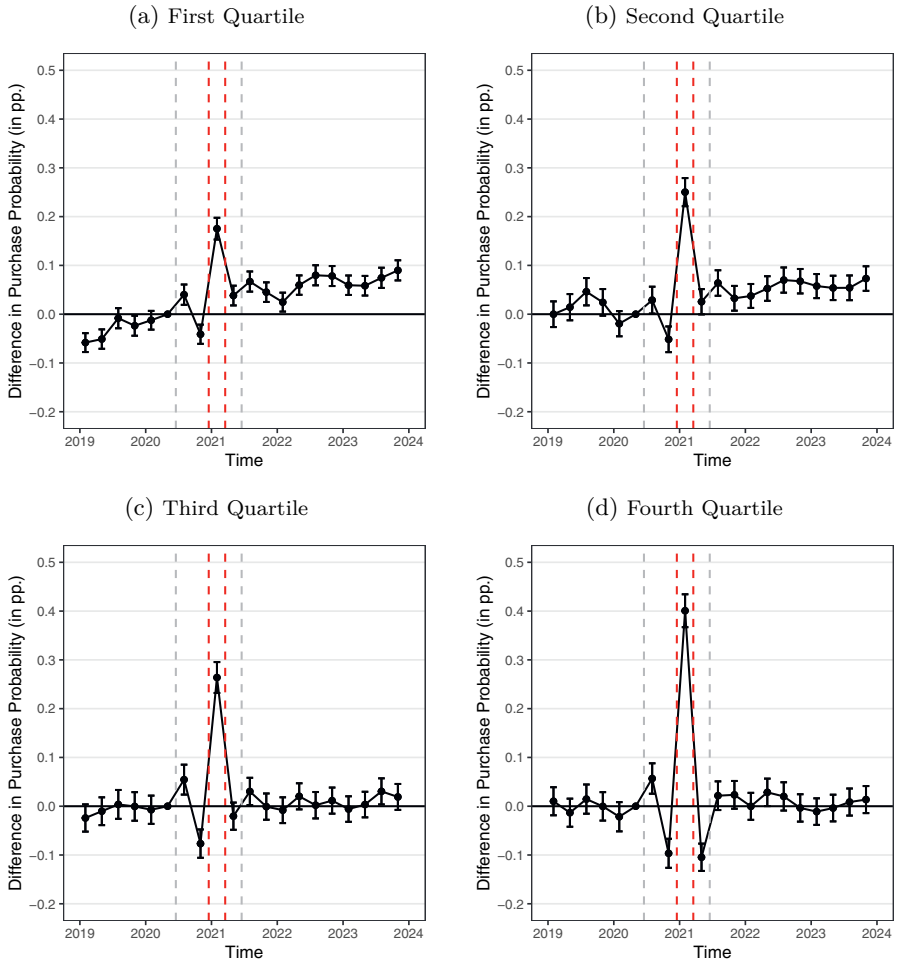


Figure 2.9: The Effect of the Starter Tax Exemption by Income Quartiles

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.

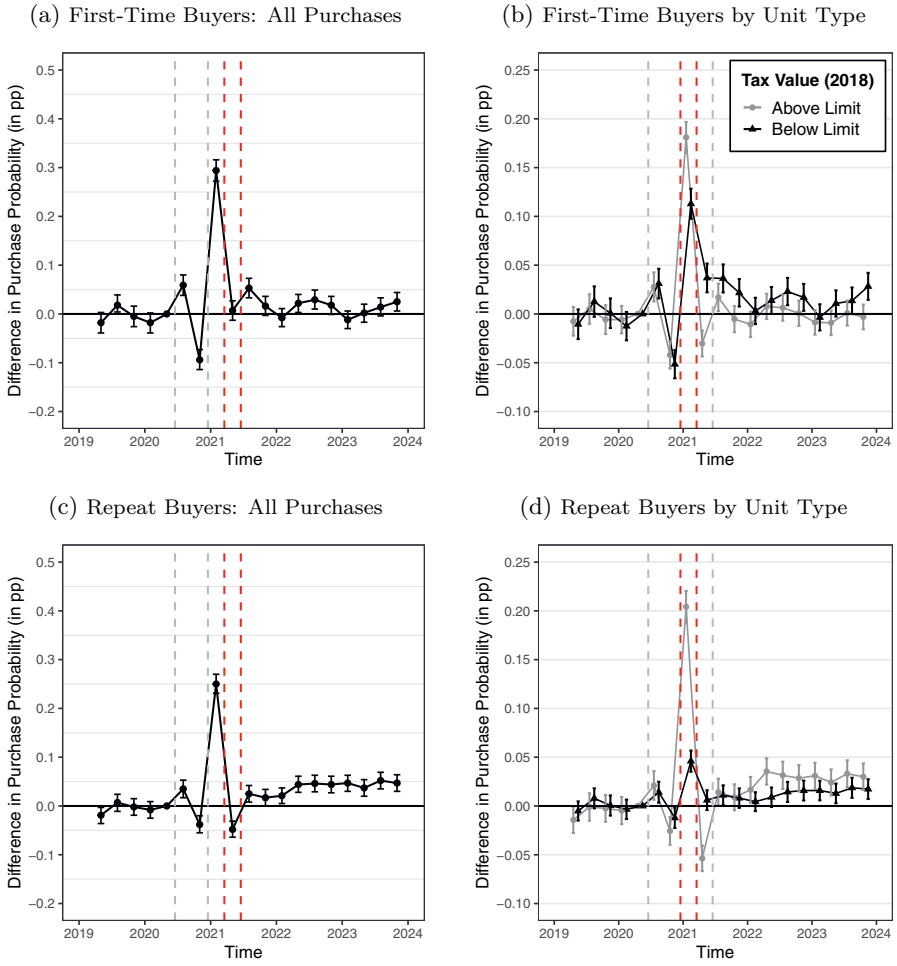


Figure 2.10: The Effect of the Starter Tax Exemption on First-Time vs. Repeat Purchases

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 2.1. 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

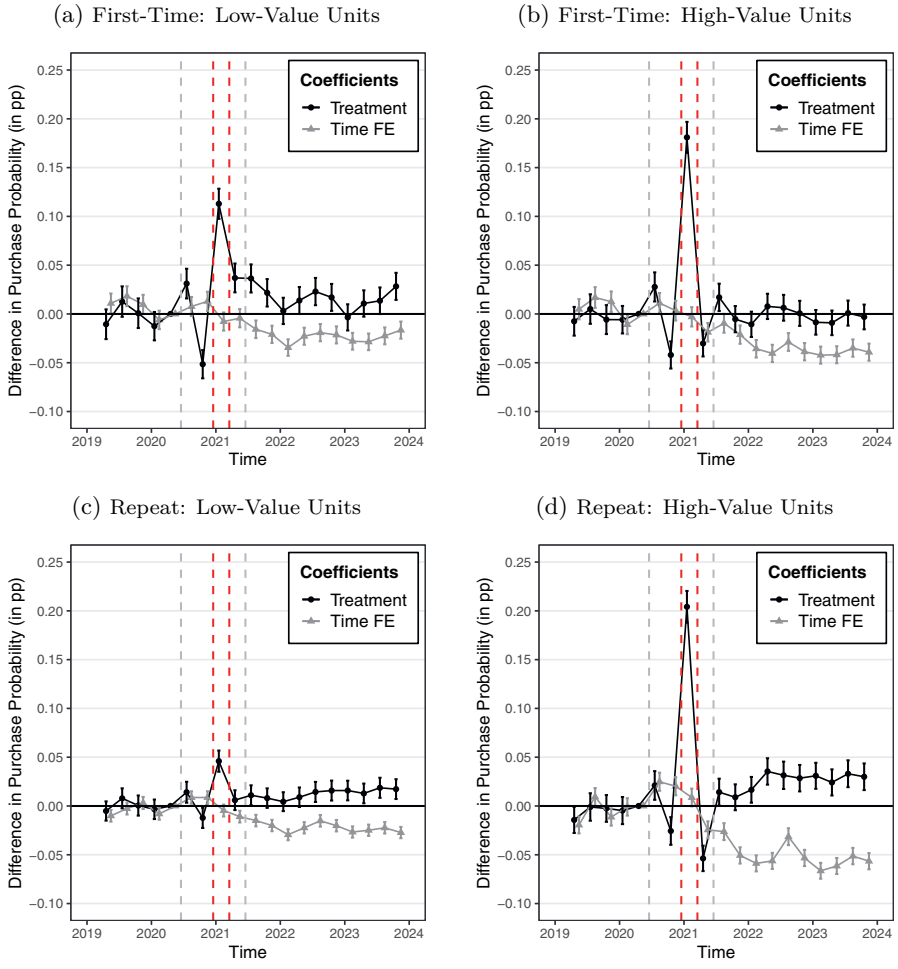


Figure 2.11: Crowding Out by Buyer Type and Housing Market Segment

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 2.2, 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 2.6: The Importance of Timing Responses for the 2021Q1 Effect

	Estimate (Table 2.5 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		0.325
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 2.1, 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 2.1 and are presented in column (3) of Appendix Table 2.5. 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 2.1 (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 2.1. Columns (7) and (8) present the relative importance of additional vs. retimed transactions, respectively.

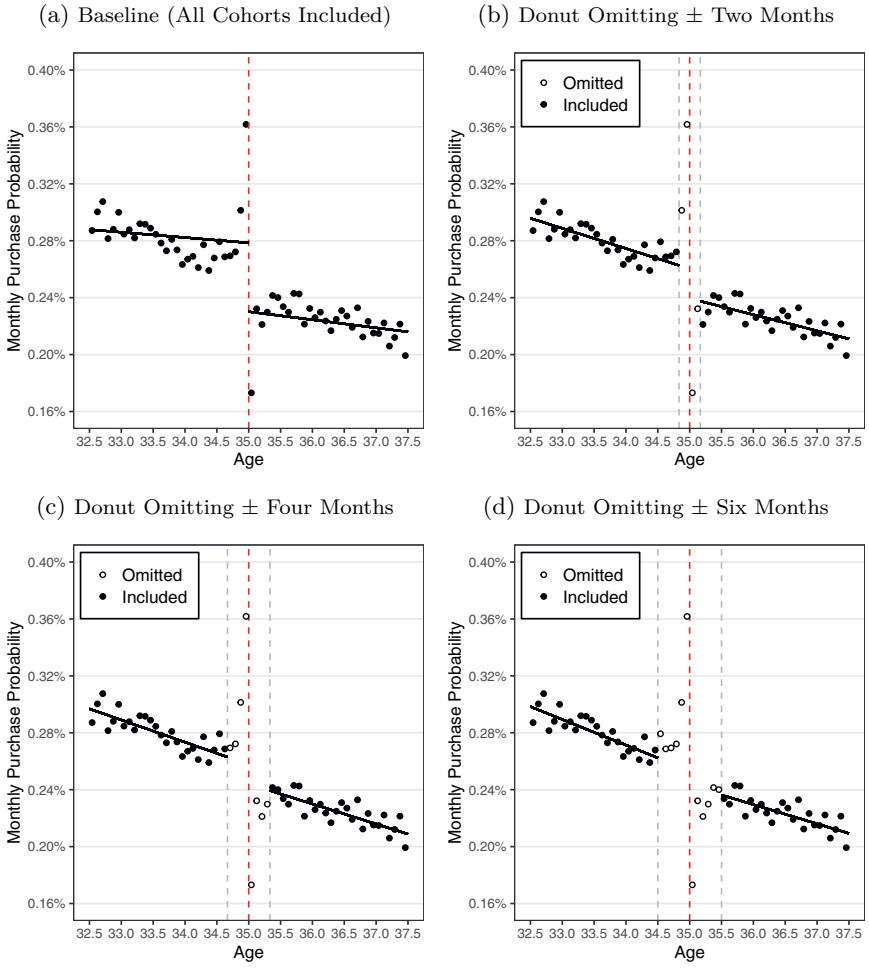


Figure 2.12: Donut RD Evidence on the *Starter Tax Exemption*

Notes: This figure shows the result of estimating the linear RD model from Equation (2.2), by excluding individuals close to the age cutoff (donut RD). Panel (a) replicates the baseline post-policy RD from Figure 2.3. 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.

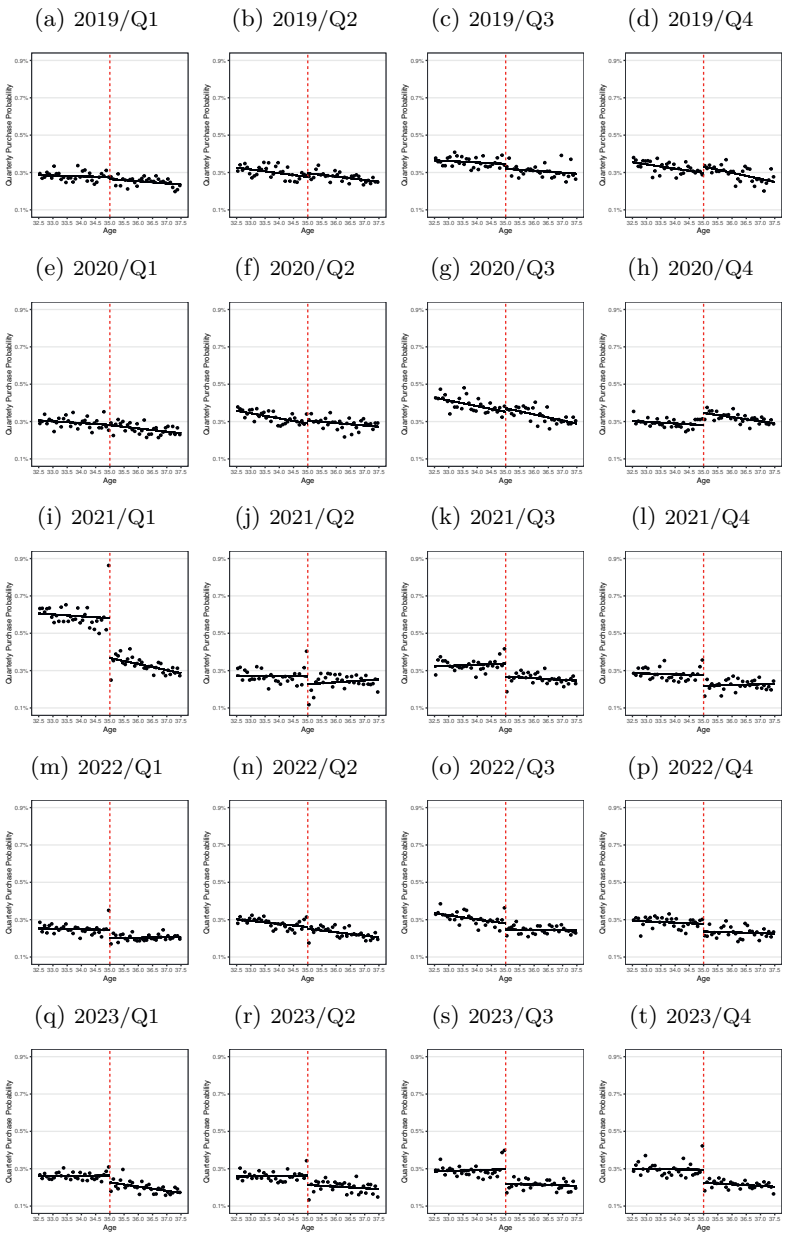


Figure 2.13: Quarterly RD Evidence on the *Starter Tax Exemption*

Notes: This figure shows the results from estimating the pooled RD model from Figure 2.3, separately for each quarter of the sample. All details on the estimation are given in the main text.

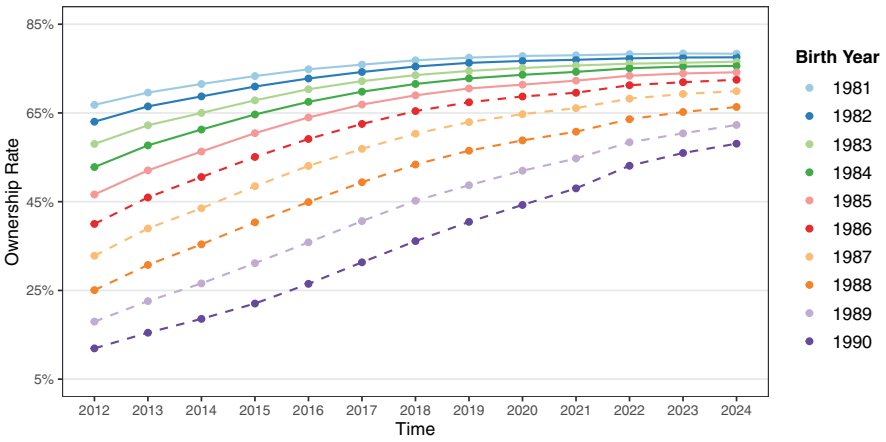


Figure 2.14: The Impact of the *Starter Tax Exemption* on Homeownership

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 2.7: Intensive-Margin Responses to the *Starter Tax Exemption*

	Purchase Amount (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 (2.1). 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. Standard errors, clustered at the individual/transaction level, are reported in parentheses (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

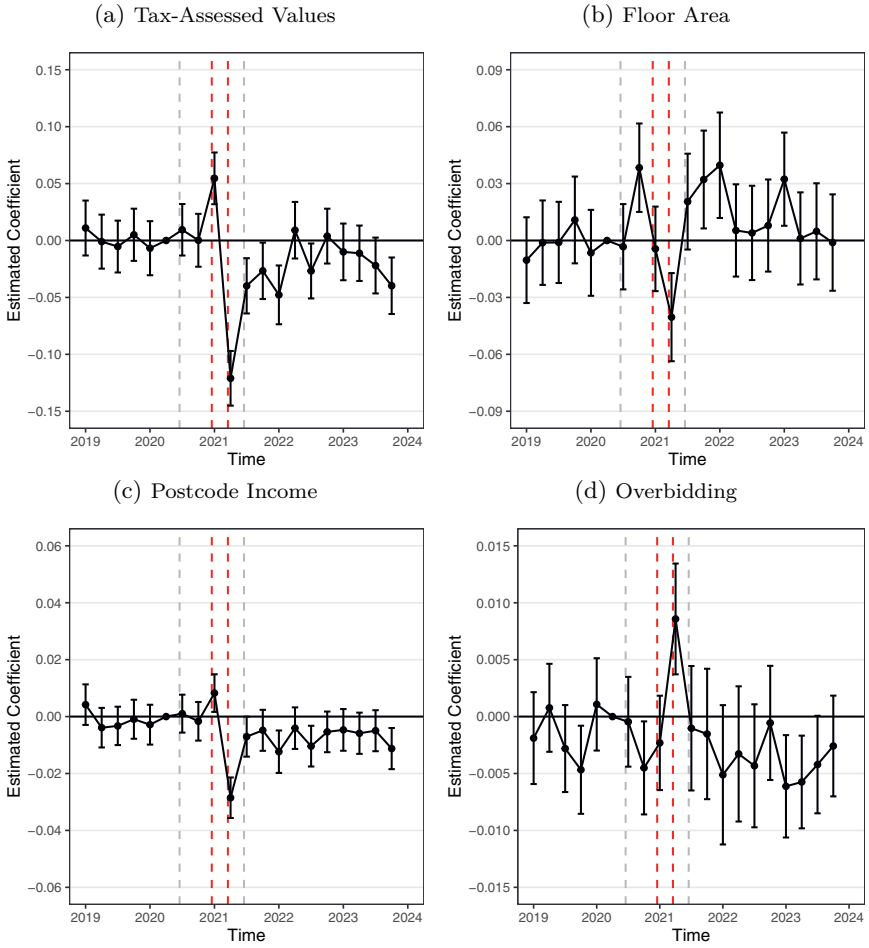


Figure 2.15: Intensive-Margin Responses to the *Starter Tax Exemption*

Notes: This figure shows the effect of the starter tax exemption on homebuyers at the intensive margin. All plots are based on Equation (2.1), 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

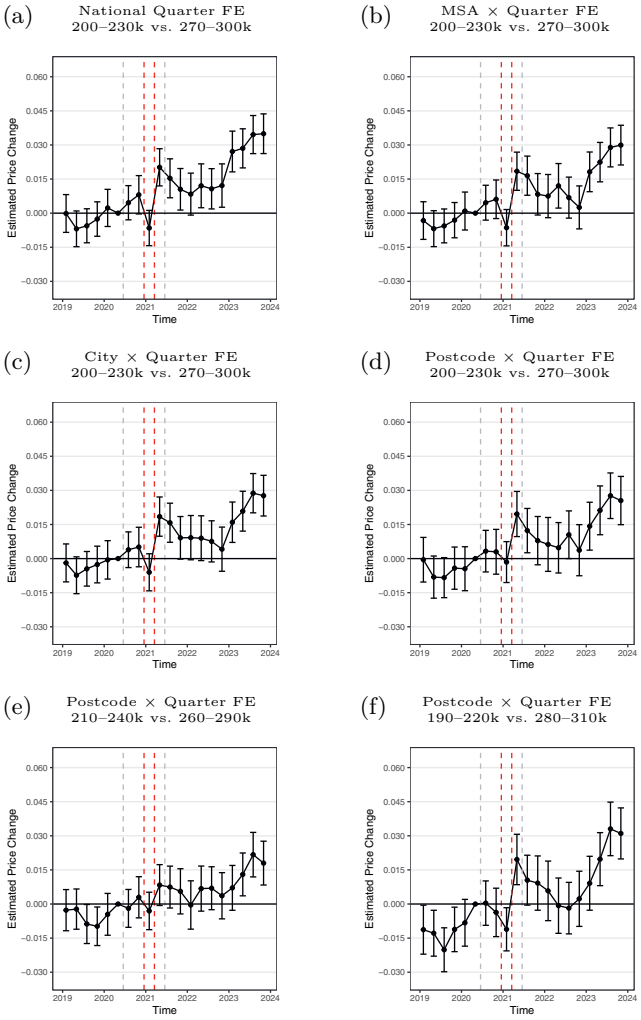


Figure 2.16: Effect of *Starter Tax Exemption* on House Prices

Notes: This figure shows the estimated effect of the *starter tax exemption* on the price of tax-exempt properties, based on Equation (2.3). 2020Q2 is the omitted quarter. The different panels correspond to columns (1)–(6) in Table 2.4. Black vertical bars correspond to 95% confidence bands.

3

Frontloading Wealth: How Parental Transfers Shape Housing and Savings

Abstract

Parents transfer substantial wealth to their children during their lifetime, yet the timing and consequences of these inter vivos transfers remain poorly understood. I show that most parental transfers occur around first-time home purchases, while transfers are largely unrelated to other life events. Because they are predominantly received by already wealthier children, these transfers increase absolute wealth inequality in the short run. Next, I identify the causal effect of parental home purchase transfers on housing outcomes and subsequent wealth accumulation, exploiting exogenous variation from a Dutch tax-exemption policy. While transfers promote wealth accumulation by enabling homeownership, recipients use them primarily to reduce mortgage debt. As a result, the initial wealth gains dissipate over time, as recipients consume the savings from lower debt service rather than reinvesting them.

3.1 Introduction

Financial transfers from parents and other relatives have long been established as a central component of household wealth (Gale and Scholz, 1994). Their relevance has been growing further in recent years, driven by demographic change and rising asset prices. Across advanced economies, the share of inherited wealth relative to national income is rising rapidly (The Economist, 2025). Yet while the academic and political debate has largely centered on inheritances, parents transfer a sizable share of their total estate while still alive (Yang and Ripoll, 2023). Crucially, *inter vivos* transfers occur earlier in life than inheritances. By affecting individuals at formative stages of the life cycle, they shape future wealth accumulation and inequality. Despite their importance, however, our understanding of parental inter vivos transfers remains limited, largely due to data constraints and a lack of credible identification strategies.

To overcome these challenges, this study leverages detailed administrative data from the Netherlands, paired with exogenous variation from a gift tax reform. The Dutch administrative records offer a number of key advantages over previous studies. They contain family network information that links individuals to their parents and provides disaggregated wealth data for both generations. They also include all reported financial transfers, with details on gift size and the identities of donor and recipient. Moreover, the data record the timing of home purchases and other major life events, enabling me to examine the incidence of parental transfers in an event study framework, rather than relying solely on cross-sectional variation. Finally, the data cover a period during which the incidence of transfers increased exogenously due to a policy change, allowing for a causal analysis of how parental transfers affect subsequent wealth accumulation.

I first document that first-time home purchases are the primary life event associated with parental transfers, whereas few transfers are made around other events. Notably, parents do not make any transfers after an involuntary job loss, suggesting that transfers are not used as a risk-sharing mechanism within extended families. While the direction of causality cannot be established in an event-study framework, the likelihood of receiving a parental transfer increases tenfold in the year of a first-time purchase and remains elevated thereafter. On average, first-time buyers in my sample receive an additional €3,200 in parental support during the year of purchase. Importantly, these are unconditional averages across all homebuyers, including those who do not receive a

transfer.

To explore how transfers are distributed across the wealth distribution, I estimate separate event studies by percentiles of parental liquid assets. Buyers below the median receive less than €1,000 on average, whereas those with parents at the 95th and 99th percentiles receive approximately €10,000 and €34,000, respectively. A similar but less pronounced gradient emerges when conditioning on recipients' own assets: transfers predominantly accrue to those with higher initial liquid assets. These patterns suggest that, at least in the short run, parental transfers reinforce rather than mitigate wealth inequality. Turning to other sources of heterogeneity, I find that larger transfers are received by female and single buyers.

Next, I investigate the impact of targeted parental transfers for first-time home purchases (*home purchase transfers*) on future wealth accumulation. Because these transfers occur early in adulthood, they can meaningfully shape life-cycle wealth trajectories. By decomposing *transfer wealth* based on how the funds are invested, I show that the implications for future wealth accumulation are not clear ex ante. On the one hand, transfer wealth may continue to grow over time if recipients invest in additional housing or liquid assets, or use the funds to deleverage costly debt—widening the wealth gap between recipients and non-recipients over time. On the other hand, if recipients use the funds primarily to finance consumption, transfer wealth will dissipate, and the initial gap will narrow in the longer run.

Empirically, testing the effect of home purchase transfers on future wealth accumulation is challenging, since a child's home purchase and the parent's transfer decision are likely jointly determined within families. I address this endogeneity by exploiting a Dutch policy reform that temporarily raised exemption limits specifically for home purchase transfers. The policy triggered a sharp, short-lived surge in such transfers: while the exemption limit was raised, 9.6% of all first-time homebuyers received a parental transfer, compared to 4.7% in the year prior. Because the reform primarily affected wealthy parents who stood to benefit from higher exemption limits, I construct an instrument that interacts the reform period with initial parental liquid assets. I use this instrument to estimate the effect of parental home purchase transfers on housing decisions and on subsequent wealth accumulation.

Home purchase transfers can affect wealth accumulation in two ways: at the extensive margin, by enabling children to buy a home at all, and at the in-

tensive margin, by influencing the type of home purchased and the associated borrowing decision. In my setting, the extensive margin response is largely predetermined, since transfers made under the special exemption were explicitly earmarked for housing-related expenditures. Nonetheless, I can use the observed spike in home purchases that is predicted by my instrument to investigate wealth accumulation in subsequent years.

2SLS estimates show that, relative to the full control group—including non-buyers—transfer recipients continue to accumulate more wealth in the years following their purchase. Seven years after the purchase, transfer recipients hold an additional €2.5 per euro of initial transfer. This is exclusively driven by the extensive margin: transfer recipients earn additional housing returns relative to those who remain renters, due to the fact that house prices appreciated considerably in the period under investigation.

However, when compared to other homebuyers without parental support, a different picture emerges. I find that, rather than buying more valuable properties or making alternative investments, recipients use parental transfers entirely to reduce their mortgage balances. As a result, they enter homeownership with lower liabilities—and therefore higher net wealth—than comparable buyers without parental support. Over time, however, this initial wealth gap narrows. After seven years, transfer recipients hold only €0.7 per euro of initial transfer, compared to buyers in the same year without a transfer. Decomposing transfer wealth into its components reveals why: rather than adopting faster amortization schedules or reinvesting savings from reduced debt service, recipients use their higher residual income (net of mortgage payments) to finance greater consumption. Taken together, these findings suggest that home purchase transfers primarily *frontload* wealth, rather than promoting sustained accumulation. In the longer run, parental support enables young homeowners to smooth consumption rather than widen wealth inequalities.

This article contributes to several fields of research. First, it contributes to the literature on intergenerational wealth persistence—a topic long of interest to economists (Charles and Hurst, 2003) but receiving renewed attention amid rising wealth inequality across countries. Two recent trends in this field are the increasing use of administrative data (Adermon et al., 2018; Boserup et al., 2016; Boserup et al., 2018; Black et al., 2020) and improved identification strategies (Ager et al., 2021; Fagereng et al., 2021; Daysal et al., 2023). One strand of this literature examines the role of inheritances and transfers in shaping wealth accumulation (Gale and Scholz, 1994; Boserup et al., 2016;

Elinder et al., 2018; Palomino et al., 2022; Nekoei and Seim, 2023; Black et al., 2025). However, whereas these studies typically provide a snapshot by examining cross-sectional wealth at a point in time, I follow transfer recipients over multiple years, allowing for a richer understanding of long-run effects and behavioral responses. In this respect, my work complements Colmsjö (2024), who studies the effects of intra-family home sales—another channel of intergenerational wealth transfer—on long-run wealth accumulation. A key advantage of my data is that I can observe all financial transfers directly and identify their effects through policy-induced variation.

A second, related field of research examines the drivers of parental transfers. While early models of parental altruism view cash transfers primarily as a risk-sharing mechanism to help children smooth consumption (Cox, 1990; Altonji et al., 1997; McGarry, 1999), the more recent literature emphasizes the importance of tax considerations (Kopczuk, 2007; Sommer, 2017; Escobar et al., 2023) and life events (Leopold and Schneider, 2011; McGarry, 2016; Andersen et al., 2020; Lee et al., 2020; Boileau and Sturrock, 2023; Kvaerner, 2023) in explaining parental transfers. However, most of these studies rely on survey data or focus narrowly on a single life event. For instance, Kvaerner (2023) studies the effect of negative health shocks experienced by parents on financial transfers. The comprehensive administrative data I use allow me to analyze multiple major life events children experience and compare their relative importance in relation to parental transfers.

Third, a growing body of evidence highlights the importance of parental wealth in the housing market. While early studies examine inter vivos transfers based on survey data (Engelhardt and Mayer, 1998; Guiso and Jappelli, 2002), more recent work combines reduced-form estimates from administrative data with life-cycle models to trace the role of parental wealth in housing markets in greater detail (Blickle and Brown, 2019; Brandsaas, 2021; Benetton et al., 2024; Wold et al., 2024; Landen Mammos, 2025). I extend this literature by examining parental wealth at an unprecedented level of detail. Specifically, I group individuals into percentiles based on their parents' wealth to study how parental support in housing transactions varies across the distribution.

Finally, I add to the empirical literature that examines marginal propensities to consume (MPC) in response to changes in housing wealth (Aladangady, 2017) and, more specifically, to changes in mortgage debt and debt service cost (Di Maggio et al., 2017; Fan and Yavas, 2020; Agarwal et al., 2022). My findings align with Ganong and Noel (2020), who show that changes in debt

service costs trigger consumption responses, whereas principal reductions that leave debt service costs unchanged have no effect. Moreover, my finding that transfer recipients forgo their initial wealth advantage by not increasing savings is in line with Bernstein and Koudijs (2024), who show that homeowners do not reduce other forms of savings when facing higher monthly costs from amortization.

The remainder of the paper is structured as follows. Section 3.2 introduces the institutional context and describes the data. Section 3.3 examines the incidence of parental transfers around first-time home purchases and contrasts this with other major life events. Section 3.4 investigates the implications of home purchase transfers for future wealth accumulation. Section 3.5 concludes.

3.2 Institutional Setting and Data

In the Netherlands, households typically aspire to enter homeownership at a relatively young age. This is facilitated by a well-developed mortgage market without down payment requirements and further encouraged by generous tax treatment, including substantial mortgage interest deductibility and low transaction costs. It is also partially driven by limited alternatives in the rental market. While the Dutch social housing sector is the largest in Europe, accounting for around a third of all housing units in the country, waiting times can be up to 10 years in the most crowded urban markets (Van Dijk, 2019). By contrast, supply in the private rental market is relatively scarce.

Rising house prices over the past decade have made access to homeownership increasingly difficult for young households. Although banks offer mortgages covering the full property value, low- and middle-income households are often constrained by tight debt-service-to-income ratios. Prices have risen most in the densely populated Randstad region (Amsterdam, Utrecht, The Hague, Rotterdam), where new construction has not kept pace with demand. As a result, the average age of first-time homebuyers increased by more than a year between 2010 and 2020, from 28.6 to 29.8 years (Tweede Kamer, 2020). These developments have generated substantial political pressure, prompting the government to introduce policies aimed at improving access to homeownership. Measures include an exemption from transfer taxes for young buyers and an outright ban on buy-to-let transactions, intended to increase the supply of owner-occupied housing (Francke et al., 2023). These policies—along

with declining house prices—have since lowered the average age of first-time buyers. Nevertheless, concerns about housing access and affordability remain central to the political debate.

3.2.1 Taxation of Parental Transfers in the Netherlands

Against this backdrop, parental wealth has become an increasingly important factor in the Dutch housing market. For children of wealthy parents, financial transfers can bridge the gap between borrowing capacity and the price of the desired home. As in most countries, the Dutch gift tax system combines marginal tax rates that rise with transfer size and a set of exemption limits.[1] Both the kinks in the tax schedule and the exemption limits are adjusted annually, roughly in line with inflation. To maintain clarity, I focus on the key rules and major policy changes, illustrating them with examples from selected years rather than documenting each annual adjustment.

The marginal tax rate for parental transfers is either 10% or 20%, depending on the size of the transfer. In 2011, at the start of the sample, the threshold at which the marginal tax rate increased was €118,708. Three exemptions were in place. First, a regular annual exemption allowed parents to transfer up to €5,030 tax-free. Second, a special one-time exemption applied to parental transfers of up to €24,144, regardless of purpose. These two exemptions remained largely unchanged over time, aside from annual inflation adjustments. By 2023, the regular and special exemption limits had risen to €6,035 and €28,947, respectively.

In addition, a separate one-time exemption exists for parental transfers explicitly intended for housing-related expenditures (such as home purchases or mortgage prepayments) or to finance expensive education. Unlike the other exemptions, this one was changed several times. In 2011, the limit was €50,030 and remained around that level until October 2013, when the government unexpectedly raised it to €100,000—but only for housing-related transfers.[2] In 2015, the limit was lowered again, before being raised for a second time in 2017. It then remained at the higher level until the end of 2022, when the

[1] To avoid confusion with other types of transfer taxes (e.g., property transfer tax), I refer to the tax on financial transfers between two individuals as gift tax, as is common in practice.

[2] The government's main motivation for raising the exemption limit for housing-related transfers was to stimulate the housing market and reduce the prevalence of underwater mortgages following the global financial crisis.

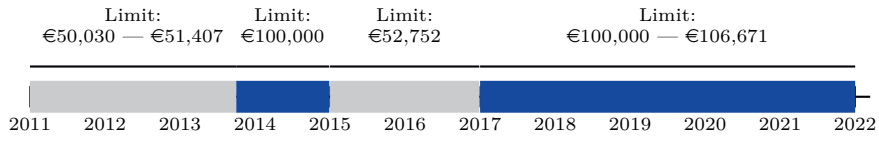


Figure 3.1: Timeline of the Special Exemption Limit for Housing-Related Transfers

Notes: This figure traces important changes to the special tax treatment of transfers for housing-related expenditures. While housing-related transfers are subject to a more generous tax exemption limit throughout the sample, this special tax exemption treatment is changed at various points. The most high-profile policy change occurred in October 2013, when the special tax exemption limit was first raised to €100,000.

special exemption was abolished altogether. Figure 3.1 illustrates the timeline of these policy changes, which are discussed in more detail in Section 3.4.

Although the special exemption limits—particularly that for housing-related transfers—allow for sizable tax-free transfers, gift tax rules in the Netherlands remain relatively strict compared to other developed economies. In the United States, a lifetime exemption of $14 million makes virtually all parental transfers tax-exempt. In Germany, no annual exemption exists, but parents can transfer up to €400,000 in a tax-free fashion every 10 years. Similarly, France allows tax-free transfers of up to €100,000 every 15 years.

3.2.2 Data Sources and Sample Construction

To study the effect of parental wealth transfers on life-cycle outcomes of young individuals, I rely on extensive administrative data provided by Statistics Netherlands (CBS). The data cover a wide range of information on every person living in the Netherlands from 2011 to 2023, reported annually. A unique feature of the Dutch administrative data is the information on family network ties, which allows me to link individuals to their parents, whose administrative records are fully available, too. These records provide detailed annual household-level balance sheet information, including a breakdown of net wealth into different assets and liabilities. One challenge includes the treatment of parents who live in separate households. To ensure a fair comparison with cohabiting parents, I average all financial variables for parents living in separate households.

For the analysis, I include all individuals in the Dutch population who were

born between 1972 and 1993 and had at least one living parent in 2011. I restrict the sample to those who are registered as household heads throughout the observation period. Since income and wealth are measured at the household level, this restriction ensures a clear separation between the resources of parents and children. A drawback of this approach is that the household head designation in different-sex couples typically refers to the male partner, introducing a gender imbalance: 67.4% of individuals in the sample are male. This imbalance represents a limitation, as parental transfer behavior, as shown below, differs by the child's gender.

Since I focus on major life events as potential triggers of parental transfers, I further restrict the sample to individuals who, in 2011, are non-homeowners, unmarried, childless, and formally employed. The resulting cohort of 133,762 individuals is followed in a balanced person-by-year panel from 2011 to 2023. While the child is the unit of observation, I also track all parental wealth variables over time. First-time home purchases are identified from changes in property ownership records, and the characteristics of purchased properties are obtained from *Kadaster*, the Dutch land registry. To contrast these purchases with other major life events, I draw on additional datasets from Statistics Netherlands (CBS). The year in which individuals marry or have their first child is inferred from annual updates to the family network data, while job loss events are identified from records of layoffs due to firm bankruptcy.

Financial transfers are recorded in a separate CBS dataset. In each year, this includes all cash transfers registered with the tax authorities that exceed the regular annual exemption limit. Importantly, transfers made under special exemption policies must also be registered, even if no tax is due. However, very small gifts below the annual exemption limit may not be captured. For each recorded transfer, both donors and recipients are identified, along with the transfer amount and payable tax. This allows me to capture, for every person and every year, the number and total amount of parental transfers received. In addition, it is recorded whether a transfer was made under one of the two special exemption limits.

3.2.3 Summary Statistics

Of all 133,762 individuals in the sample, 10,427 (7.8%) receive a parental transfer at some point between 2011 and 2023. In Table 3.1, key summary statistics are reported, separately for transfer recipients and non-recipients. All

demographic and financial variables are measured in 2011, before any transfers are received.

Unsurprisingly, the most striking difference between recipients and non-recipients relates to their parents' resources. Recipients' parents own, on average, about €219,000 in liquid assets (bank balances, stocks, and bonds), whereas non-recipients' parents hold only €53,000. Recipients themselves also hold more than twice as much in liquid assets as non-recipients, even before receiving any transfer. Recipients are also positively selected with respect to their income and education: while only around 48% of non-recipients hold at least a college degree, more than 76% of recipients do.

Turning to major life events, the statistics suggest an association with parental transfers. Transfer recipients more often experience "positive" life events (first-time home purchase, marriage and childbirth), with the largest gap observed for first-time home purchases. While only around 34% of non-recipients become homeowners throughout the sample period, more than 59% of recipients do. The pattern is reversed for the sole negative life event considered—job loss due to firm bankruptcy—which transfer recipients experience less frequently. Whereas fewer than 20% of transfer recipients experience none of the major life events that I study, more than 42% of non-recipients do. In the following section, this association will be tested in a more formal way.

3.3 Parental Transfers for First-Time Home Purchases

In the first part of the empirical analysis, I investigate the incidence of parental cash transfers around first-time home purchases, and compare this to other major life events that children experience. A key strength of my data is that it includes information on different life events within the same sample, allowing for a quantitative comparison of their relative importance. Nonetheless, given that first-time home purchases (and other life events) are likely co-determined with parental transfers, I do not attempt to infer the direction of causality. In the first part of the analysis, I show that first-time home purchases are an important—and indeed the most relevant—life event that is associated with parental transfers. I assess to what extent the dominance of home purchases is driven by the special tax treatment. Next, I document substantial inequalities in transfers along the distributions of parent and child assets. Finally, I show additional heterogeneity based on parent and child characteristics.

Table 3.1: Summary Statistics

	Transfer Recipients			Non-Recipients		
Variable	P25	Mean	P75	P25	Mean	P75
Age	27.00	30.29	33.00	27.00	31.65	36.00
Hours Worked	26.54	31.78	39.08	25.69	31.48	38.10
Household Income (€)	20,599	30,112	36,624	18,608	26,556	31,759
Household Liquid Assets (€)	4,260	29,776	34,299	1,410	13,382	14,472
Household Net Wealth (€)	−593	26,329	32,405	−3,535	6,534	11,804
Parental Household Income (€)	41,404	64,706	81,638	27,388	43,033	53,936
Parental Liquid Assets (€)	46,304	218,567	310,612	6,194	52,647	55,078
Parental Net Wealth (€)	247,755	923,758	857,997	9,866	185,288	248,705
Share Female (%)		32.79			32.54	
Share College Educated (%)		76.46			48.17	
Share Single Household (%)		73.58			77.72	
Share First Home Purchase (%)		59.48			34.12	
Share Job Loss (%)		4.77			6.50	
Share First Childbirth (%)		48.60			36.15	
Share Marriage (%)		33.38			26.12	
Share No Life Event (%)		19.70			42.39	
Individuals		10,427			123,335	

Notes: This table shows summary statistics, separately for individuals who receive a parental transfer between 2011 and 2023 (recipients) and those who do not. All financial and demographic variables are measured in 2011 levels, i.e., before any transfer is received. All prices are adjusted for inflation based on the CPI and expressed in 2015 levels.

3.3.1 The Dominance of Transfers Around First-Time Home Purchases

At the start of the sample period in 2011, all individuals are renters and have not yet experienced any of the life events analyzed in this study. While a substantial share remains in this initial state throughout the sample, others experience one or more life events at different points in time. I examine how parental cash transfers evolve around the first occurrence of each event.[3]

In the baseline analysis, I consider each life event separately and estimate staggered event studies of the following form:

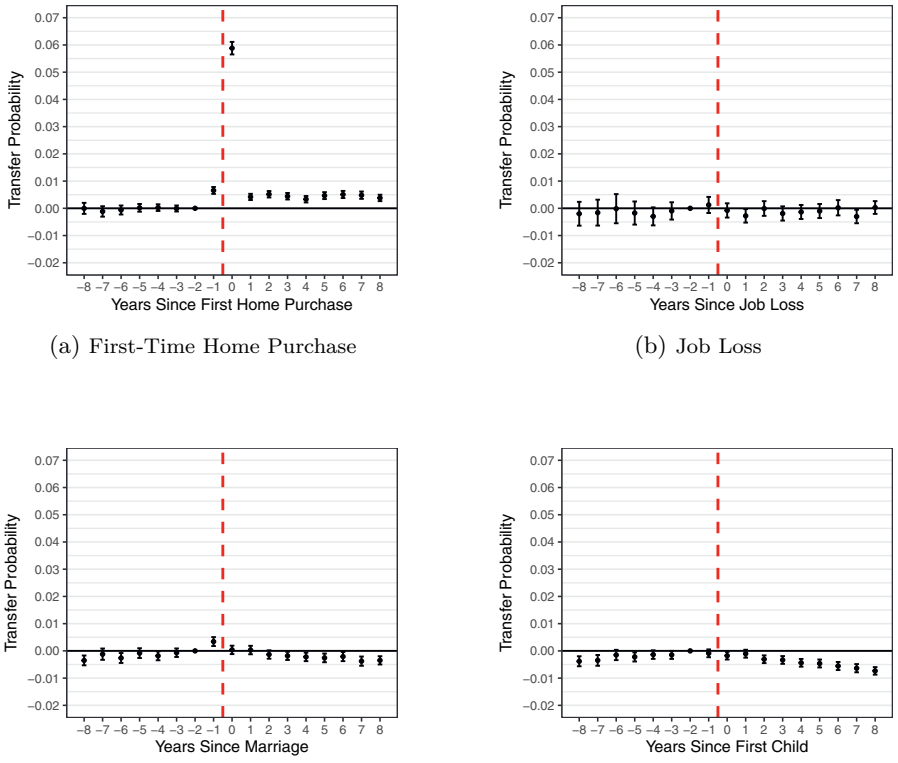
[3]Because the data begin in 2011, I cannot rule out that some individuals experienced these life events prior to the sample period (e.g., they may have previously owned a home or been married). However, given that the sample consists of young adults, such cases are likely rare.

$$Y_{it}^{(d)} = \alpha_i^{(d)} + \delta_t^{(d)} + \sum_{k \neq -2} \beta_k^{(d)} D_i^{(d)} \mathbf{1}\Big\{K_{it}^{(d)} = k\Big\} + \varepsilon_{it}^{(d)}, \tag{3.1}$$

where $d \in \{\text{Home Purchase; Job Loss; Marriage; Birth of Child}\}$. The outcome variable in each model, $Y_{it}^{(d)}$, denotes one of the measures of parental transfers introduced below. $D_i^{(d)}$ indicates whether person i ever experiences event d throughout the sample and $K_{it}^{(d)} = \left(t - T_i^{(d)}\right)$ indexes time relative to $T_i^{(d)}$, the year in which the event takes place. The coefficients $\beta_k^{(d)}$ trace the average difference in transfers between individuals who experience event d and those who never do, k years before or after the event. The coefficient for $k = -2$ is omitted, so that all dynamic effects are expressed relative to the difference two years before the event. Person fixed effects $\alpha_i^{(d)}$ control for time-invariant individual heterogeneity, and year fixed effects $\delta_t^{(d)}$ absorb common shocks. Finally, the event window is capped at eight years around the respective event and observations outside of that window are grouped in the end-cap variable (Miller, 2023).

I measure parental transfers in three ways. First, I construct a binary indicator for whether person i received any parental transfer in year t. The corresponding coefficients are visualized in Figure 3.2. All four life events are shown in panels (a)–(d), using a common vertical scale. The results are striking: while parental transfers respond strongly to first-time home purchases, they are largely unresponsive to other major life events. As shown in panel (a), pre-trends are flat in the years leading up to a first-time home purchase. In the year before the purchase, the likelihood of receiving a transfer increases by 0.7 percentage points on average, and by 6.0 percentage points in the year of the purchase itself. This represents roughly a tenfold increase relative to the incidence of transfers among individuals who experience no life event during the sample period. In subsequent years, the probability of receiving a transfer remains elevated by about 0.5 percentage points. This suggests that parents not only contribute to the initial home purchase but also continue to assist with subsequent housing-related expenses such as renovations or durable goods.

In contrast, there is no evidence of increased parental transfers following an involuntary job loss. This finding contrasts with Andersen et al. (2020), who document that parents often provide informal insurance against income shocks.



(a) First-Time Home Purchase

(b) Job Loss

(c) Marriage

(d) Birth of Child

Figure 3.2: Parental Transfers Around Major Life Events

Notes: This figure presents event study coefficients estimated from Equation (3.1), where all four major life events are modeled separately. The dependent variable is a binary indicator equal to one if a parental transfer was received in a given year. The omitted period is $t-2$, corresponding to two years before the event occurs. Standard errors are clustered at the individual level, and the error bars represent 95% confidence intervals.

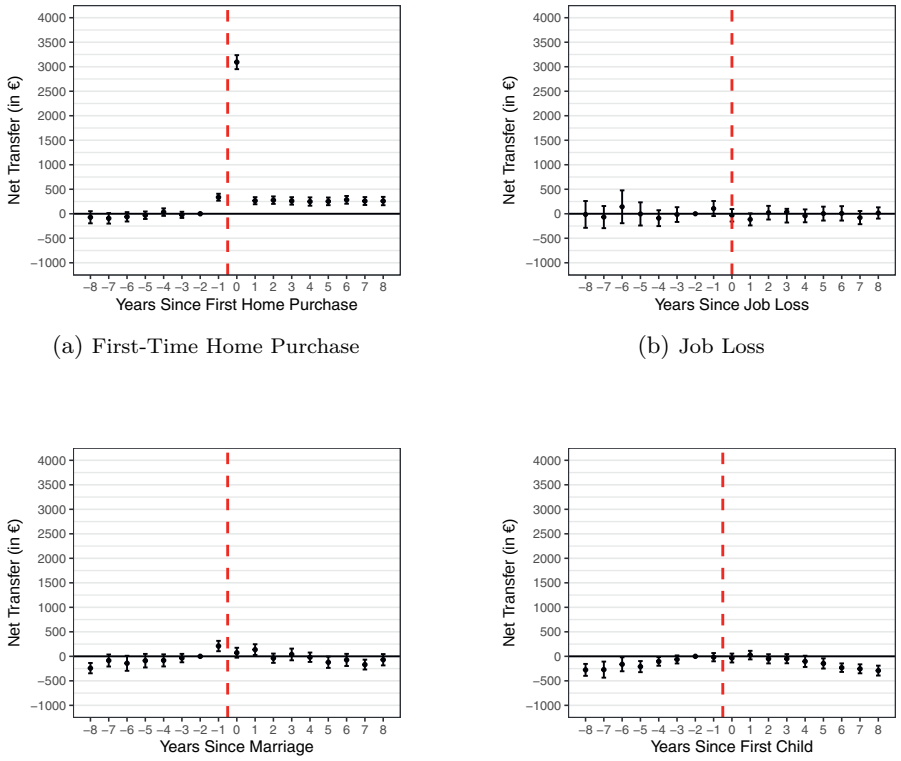
For marriage, I observe a small but statistically significant increase in transfers in the year preceding the event. Finally, there is no evidence that parents make transfers around the birth of their first grandchild. If anything, the incidence of parental transfers declines in the years after their children have children of their own.

Considering net parental transfers (in €) as an alternative outcome, a similar picture emerges. This is illustrated in Figure 3.3. For first-time home purchases, I find that the average (after-tax) transfer increases by €3,200 in the year of the purchase, and remains around €300–400 higher in the following eight years. Importantly, these are unconditional averages: for individuals who do not receive a transfer, net transfers are coded as 0 rather than NA. Thus, these numbers represent the additional financial support that the average homebuyer can expect from their parents around the time of the first purchase. Regarding other life events, I again find little to no evidence for parental support. The sole exception is marriage, before which parental transfers increase by around €200.

Next, I focus on an alternative outcome that is not based on formal transfers. There are two reasons for this: first, since cash transfers below the annual exemption threshold (around €6,000) do not need to be registered, the administrative data may understate the incidence of small transfers. Second, instead of making an outright transfer, parents may offer financial support by covering some of their child's expenditures after a life event. For instance, parents may take over (some of) the cost of their child's wedding ceremony directly, rather than making a registered cash transfer. To capture such informal support, I examine changes in parental bank deposits around each of the four life events. While both forms of support—registered transfers and direct payments—should manifest as a drop in parental deposits following the event, I interpret drops in parental deposits that do not coincide with registered transfers as evidence for informal support.

For first-time home purchases, this pattern is evident. After similar pre-trends, parental deposits fall by about 4% in the year of purchase and continue to decline in subsequent years, consistent with the sustained elevation in parental transfers documented earlier. For involuntary job loss, estimates are imprecise given the small number of such events, but there is no visible break in trend, suggesting that parents do not provide insurance even informally by covering expenses.

In contrast, parental deposits decline by nearly 3% around the time of marriage, consistent with parents helping to finance marriage-related expenses such as wedding ceremonies. Finally, around the birth of a first child, there is again evidence of informal support: although no registered transfers are observed, parental deposits decline steadily after the event. This suggests that, in addition to the intergenerational time transfers, which are established in



(a) First-Time Home Purchase

(b) Job Loss

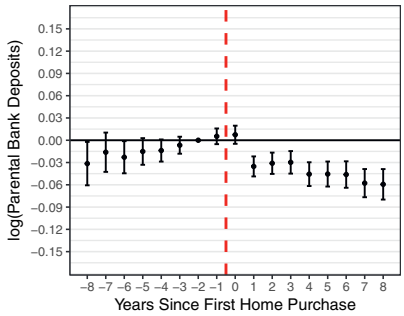
(c) Marriage

(d) Birth of Child

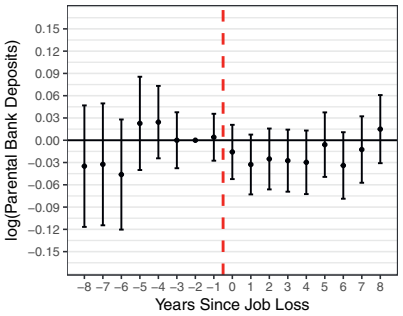
Figure 3.3: Net Parental Transfers Around Major Life Events

Notes: This figure presents event study coefficients estimated from Equation (3.1), where all four major life events are modeled separately. The dependent variable is the net transfer (in €) that a person has received in a given year. Years without transfer are coded as 0, rather than NA. The omitted period is $t - 2$, corresponding to two years before the event occurs. Standard errors are clustered at the individual level, and the error bars represent 95% confidence intervals.

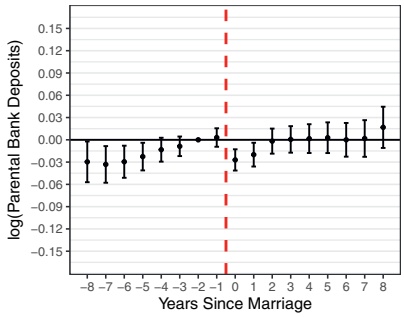
the literature (Eibich and Siedler, 2020), grandparents also provide monetary assistance by covering some of their children's expenses.



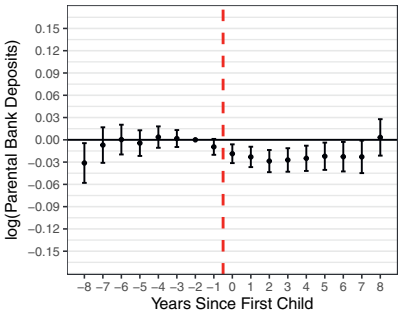
(a) First-Time Home Purchase



(b) Job Loss



(c) Marriage



(d) Birth of Child

Figure 3.4: Parental Bank Deposits Around Major Life Events

Notes: This figure presents event study coefficients estimated from Equation (3.1), where all four major life events are modeled separately. The dependent variable is the natural logarithm of parental bank deposits in a given year. The omitted period is $t-2$, corresponding to two years before the event occurs. Standard errors are clustered at the individual level, and the error bars represent 95% confidence intervals.

3.3.2 Robustness Analysis

One key issue with the baseline event study specification is that it ignores correlations between different life events. As a result, it may fail to correctly attribute the role of each individual event. For example, if marriage typically occurs shortly before or after the first-time purchase, home purchase transfers

may be mistakenly attributed to marriage. Therefore, I implement a combined model, in which all life events are considered jointly:

$$Y_{it} = \alpha_i + \delta_t + \sum_d \sum_{k \neq -2} \beta_k^{(d)} D_i^{(d)} \mathbf{1}\Big\{K_{it}^{(d)} = k\Big\} + \varepsilon_{it}. \tag{3.2}$$

In contrast to Equation (3.1), this model includes a full set of event time dummies for all four life events. Therefore, it identifies the association with parental transfers for each event, after partialling out the effect of all other events. The omitted group comprises those who do not experience any of the four events during the sample, accounting for 40.6% of all individuals. The resulting coefficients, reported in Figure 3.10, align very closely with the baseline model, suggesting that correlations in the timing of events are not a first-order concern. To investigate this further I remove, for each individual who experiences more than one life event, observations after a second event occurs. This further reduces potential confounding, as the post-event coefficients are purged of any influence from additional events. The results, shown in Figure 3.12, again confirm that formal parental transfers are primarily directed toward supporting children's first-time home purchases, while other events do not elicit a response.

One reason for modest parental support around life events other than first-time purchases is that they may not be important enough, in financial terms. To test this, I estimate how an individual's own bank balance evolves around each life event. The resulting coefficients are shown in Figure 3.11. Notably, for all events except involuntary job loss, parallel trends are clearly violated. This is unsurprising, given that first-time home purchases, marriage and childbirth are endogenous from the perspective of the individual. Nevertheless, there is a discernible change in bank deposits around the time of each event. For instance, people who have their first child accumulate bank deposits at an increased rate in the years leading up to the event, before reducing their balance quickly afterwards. This suggests that the small response of parental transfers to life events other than first-time home purchases is not due to these events being financially insignificant for the child.

Finally, I address the role of the special tax treatment of housing-related transfers in driving the baseline results. A legitimate concern is that the sharp increase in parental transfers around first-time home purchases may partly

reflect the more generous exemption threshold for housing-related gifts, as discussed in the previous section. I devise two additional tests to investigate this concern. First, I examine how transfers for first-time home purchases vary over time, related to changes in the special exemption threshold. To that end, I estimate the difference in parental transfers between those who make a first-time purchase and those who do not, separately for each year in the sample.[4] The resulting estimates are shown in Figure 3.13, with years in which the special exemption threshold was raised highlighted in blue. The pattern is clear: home-purchase-related transfers respond strongly to the tax treatment. In years when the special exemption was raised to €100,000 or more, first-time homebuyers were substantially more likely to receive parental transfers than in other years. This response is most pronounced in 2014—the year of the first major policy change—which I analyze in more detail in the following section. However, even in years when the exemption threshold was around €50,000, first-time homebuyers remained 3–4 percentage points more likely to receive a transfer compared to non-buyers.

Nevertheless, because housing-related transfers were always subject to some form of special exemption during the sample period—only the size of this exemption varied over time—this analysis cannot fully rule out tax-based explanations. Therefore, I take one step further and identify, for each transfer in the data, whether it was made under a special exemption or not. In a final robustness test, I retain only regular (i.e., non–tax-exempt) transfers and repeat the event-study analysis. This yields the most conservative estimate of the increase in parental transfers around home purchases.[5]

The resulting estimates are shown in Figure 3.14. The implications are twofold. First, the special tax treatment of housing explains a substantial share of the quantitative magnitude of the observed spike in transfers. Focusing only on regular transfers, the increase in parental transfers in the purchase year is about 0.6 percentage points, compared to roughly 6 percentage points for all transfers. Second, in qualitative terms, the dominance of housing-related transfers is not driven by tax policy. Even when focusing only on non–tax-exempt transfers, first-time home purchases remain the only life event associated with a significant increase in parental support.

[4] Concretely, I estimate the following model: $Y_{it} = \delta_t + \sum_t \beta_t \cdot D_{it} + \varepsilon_{it}$, where D_{it} indicates that person i has made a first-time home purchase in year t.

[5] Because a special exemption for housing-related transfers existed throughout the sample, this restriction retains only parents who were unable or unwilling to file for an exemption—for example, because a large transfer had already been made previously.

3.3.3 Transfers Across the Wealth Distribution

Ultimately, the main determinant of a financial transfer is the parents' available resources. Transfers can only be made by the parents who have the means to do so. The administrative data allow me to investigate in detail how transfers vary with parental wealth. To that end, I rank each person in the sample according to their parents' total liquid assets (bank deposits plus financial securities), measured in 2011. In Figure 3.15, the baseline event study is estimated separately in each quartile of parental liquid assets. It suggests that, unsurprisingly, parental home purchase transfers depend strongly on the amount of liquid assets that parents have access to. Importantly, it also suggests that the parallel trends assumption continues to hold if the sample is split by parental assets.

Building on this, I then group people into percentiles of parental liquid assets and repeat the baseline analysis within each percentile. There are vast differences in parental resources across the sample. At the median, parents' total liquid assets are equal to €17,000. By contrast, parents in the 10th percentile of the liquid asset distribution hold only €800, while those in the 90th percentile hold more than €130,000. Given these stark inequalities, financial transfers are expected to rise nonlinearly across the distribution. Given a sample size of at least 1,000 individuals in each of the 100 percentiles, there is sufficient statistical power to conduct 100 separate event studies, in order to investigate the transfer patterns across the distribution in detail. Since pre-trends are virtually flat in the baseline analysis, I focus here on estimating a single treatment effect for each percentile. Specifically, I restrict the sample to observations from the pre-period (up to two years before the first-time home purchase) and the purchase year. I then collapse the dynamic model from Equation (3.1) to a static specification, comparing transfers made in the purchase year with those made during the entire pre-period. In other words, I estimate only the size of the spike in transfers immediately surrounding the purchase, neglecting transfers made shortly before or in the subsequent years.

The estimates, shown in Figure 3.5, reveal pronounced inequalities in parental support across first-time homebuyers. Up to the 60th percentile, the average first-time buyer receives less than €2,000 in parental support, while those at the 80th percentile receive approximately €3,000 on average. At the 90th and 95th percentiles, buyers receive approximately €4,500 and €10,000, respectively. Finally, children with the wealthiest parents, at the top of the liquid asset distribution, receive more than €30,000 upon entering homeownership.

To reiterate, these are unconditional averages calculated over all people within each percentile group, including those who do not receive a transfer. This suggests that children from affluent families do not need to wait for an inheritance to access parental resources. Instead, parents with sufficient means use their child's first-time home purchase as an opportunity to pass on wealth proactively.

A similar, though less pronounced, pattern emerges when considering heterogeneity by child (i.e., recipient) assets. Although the gradient is less steep, panel (b) indicates that children with higher liquid assets in 2011 receive more transfers in the year of their first-time home purchase. This pattern largely reflects the positive intergenerational correlation in wealth. On average, children at the median of the liquid assets distribution receive around €2,000 in parental support, while those at the 90th and 95th percentiles receive around €5,000 and €7,500, respectively. The fact that transfers are predominantly received by wealthier children suggests that, at least in the short run, parental transfers increase absolute wealth inequality.

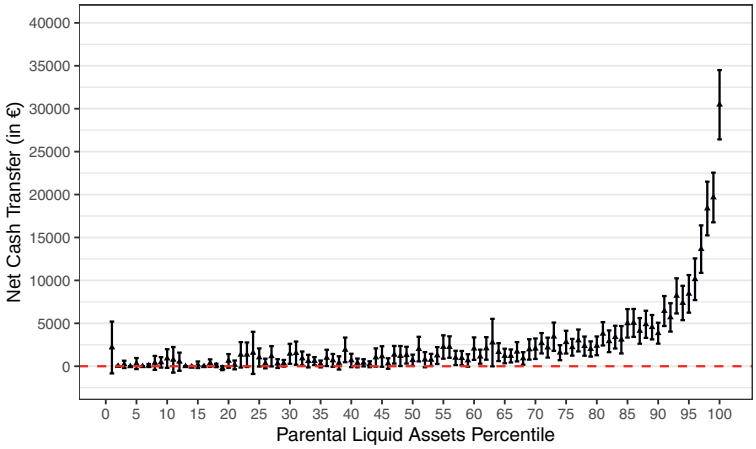
3.3.4 Other Sources of Heterogeneity

The richness of the administrative data allows me to examine several additional sources of heterogeneity, which are presented in turn. To test these, I extend the baseline dynamic event study as follows:

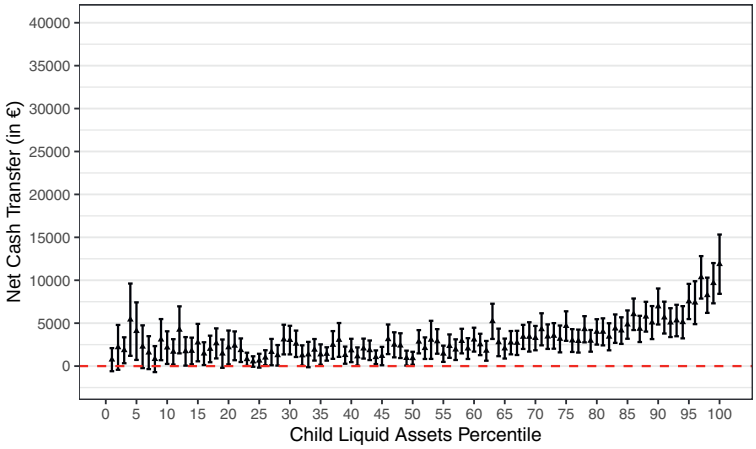
$$Y_{it} = \alpha_i + \delta_t + \sum_{k \neq -2} \beta_k \, D_i \, \mathbf{1}\{K_{it} = k\} + \sum_{k \neq -2} \gamma_k \, D_i \, \mathbf{1}\{K_{it} = k\} \; Int_i + \varepsilon_{it}. \tag{3.3}$$

The additional term, Int_i, represents different binary interaction terms. The new dynamic treatment coefficients, γ_k, indicate how patterns of parental transfers differ between the two groups being compared. Figure 3.6 presents dynamic treatment effects for the omitted baseline group (β_k) in blue, alongside the total interacted effect ($\beta_k + \gamma_k$) in green. Since the latter represents the sum of two estimated coefficients, standard errors are calculated manually, accounting for the covariance between them.

First, first-time buyers over 35 receive slightly larger transfers than those un-



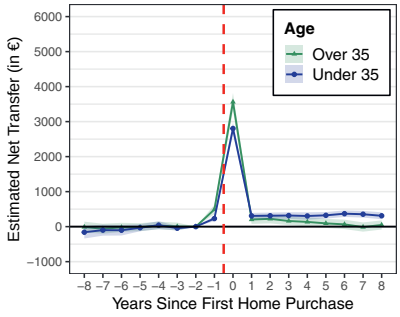
(a) Net Transfer for First-Time Home Purchase: Split by Parental Assets



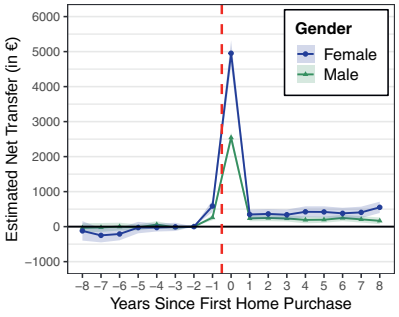
(b) Net Transfer for First-Time Home Purchase: Split by Child Assets

Figure 3.5: Transfers for Home Purchases Along the Distribution of Liquid Assets

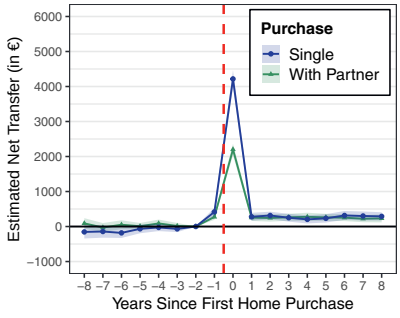
Notes: This figure presents event study coefficients from Equation (3.1), estimated separately within each percentile of the distributions of parent household assets in 2011 (panel (a)) or household assets in 2011 (panel (b)). The life event that is studied is a first-time home purchase. Rather than estimating a full dynamic model, all pre-event years are pooled together, and the plotted treatment effect reports only the spike in transfers in the purchase year ($t = 0$). Standard errors are clustered at the individual level, and the error bars represent 95% confidence intervals.



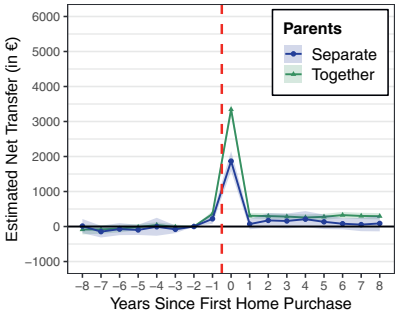
(a) Younger vs. Older Buyers



(b) Male vs. Female Buyers



(c) Single Buyers vs. Couples



(d) Parents Together vs. Separated

Figure 3.6: Heterogeneity in Net Transfers Around First-Time Home Purchases

Notes: This figure presents event study coefficients estimated from Equation (3.3), with first-time home purchases as life event. The dependent variable is the net transfer (in €) that a person has received in a given year. Years without transfer are coded as 0, rather than NA. The omitted period is $t - 2$, corresponding to two years before the event occurs. Standard errors are clustered at the individual level, and the error bands represent 95% confidence intervals.

der 35. Panel (a) of Figure 3.6 illustrates this pattern. In the transaction year, younger buyers receive an average net transfer of approximately €2,800, compared to €3,500 for older buyers. Most likely, this reflects a tendency of older parents to make more transfers, compared to younger ones. However, older buyers receive no further transfers after the purchase, whereas homeowners under 35 continue to receive approximately €300–400 per year in the subsequent years. Second, a more pronounced discrepancy emerges when comparing buyers by gender. While male household heads receive an average transfer of €2,500 at the time of purchase, female household heads receive twice that amount, approximately €5,000. They continue to receive slightly higher levels of support in the years after the purchase, although these estimated differences are no longer statistically significant. Third, buyers receive substantially larger transfers if they purchase on their own (more than €4,000), compared to purchases that are made with a partner (around €2,000). Finally, parents who do not live together make smaller transfers (around €2,000) than those who cohabit (around €3,500), potentially due to coordination challenges in providing financial support.

3.4 Home Purchase Transfers and Future Wealth Accumulation

So far, I have documented that parents transfer a substantial amount of wealth around their child's first-time home purchase and that transfer recipients tend to be wealthier to begin with. However, given that transfers for first-time home purchases are received at a relatively young age, it is important to understand what their effect is on future life-cycle outcomes. Finding that transfer recipients also accumulate more wealth in the future would suggest that the effect of transfers on wealth inequality is even stronger in the longer run. If, instead, transfer recipients consume most of what they are given, the longer-run effect of transfers on inequality would be attenuated. Which of these two patterns is found depends crucially on how transfer recipients decide to allocate their funds. This will be investigated in this section.

I will proceed as follows: first, I explain theoretically how *transfer wealth*—the additional net wealth in year t per euro of parental transfer received in year 0—evolves over time, based on the behavior of the recipient. Then, I discuss the empirical challenges that are involved in estimating the causal effect of parental transfers and introduce my strategy for overcoming these challenges, based on

quasi-experimental variation from policy-induced changes to the special tax exemption limit. Finally, I present the empirical results from estimating the future path of transfer wealth. These are based on a comparison of transfer recipients with two different counterfactual groups: the full sample, including individuals who remain renters; and a buyer-only sample, where recipients are compared to other homebuyers who do not receive parental support.

3.4.1 Decomposing Wealth From Home Purchase Transfers

Consider an individual who receives a parental transfer of size T to support her first-time home purchase. Of that amount, T_h is used to purchase a more valuable home, T_s is kept in liquid savings and T_m is used to lower the initial mortgage principal. The residual, T_c, is immediately consumed. Let Ω_t denote transfer wealth, that is, the additional net wealth at time t that recipients derive for every euro of parental transfer received at $t = 0$, when a home purchase is made. Transfer wealth at $t = 0$ is equal to:

$$\Omega_0 = T_h + T_s + T_m. \tag{3.4}$$

Every euro that is not immediately used for consumption raises net wealth one-to-one, either by increasing assets or lowering mortgage debt. Over time, funds that are initially invested into a more valuable home (T_h) or liquid savings (T_s) yield a constant return of r_h and r_s, respectively.[6] In contrast, how the initial reduction in mortgage debt (T_m) shapes future wealth depends on two additional decisions that recipients must make. First, they must decide whether the smaller principal should be amortized over a shorter period. This in turn determines the reduction in recurring debt service cost, compared to a counterfactual mortgage without reduced principal. Second, they must decide whether reductions in debt service cost are offset by investments into liquid savings, or whether they finance higher consumption.

Formally, let D_0 denote the initial principal on a fully amortizing mortgage absent a transfer, and $D_0{}' = D_0 - T_m$ denote the reduced principal for transfer recipients. To simplify the exposition, mortgage interest rates (r_m) are as-

[6] To simplify the exposition, I abstract from uncertainty in future returns.

sumed constant and therefore unaffected by the reduction in principal.[7] The initial term of the mortgage is n, but transfer recipients may opt for a faster amortization schedule, such that $n' \leq n$. Crucially, in choosing n', transfer recipients effectively also pin down their debt service payments $P' \leq P$.[8] Debt service cost and outstanding mortgage debt at time t are given by the standard annuity formulas:

$$P = \frac{r_m D_0}{1-(1+r_m)^{-n}} \qquad D_t = (1+r_m)^t D_0 - P\,\frac{(1+r_m)^t - 1}{r_m} \tag{3.5}$$

The first way in which mortgage reductions shape future wealth is determined by how the outstanding mortgage debt of transfer recipients evolves mechanically, relative to the counterfactual debt level without a transfer:

$$\Delta D_t(\Delta n') = D_t - {D_t}' = (1+r_m)^t\, T_m - \Delta P\,\frac{(1+r_m)^t - 1}{r_m}, \tag{3.6}$$

where $\Delta D_t = D_t - D_t'$, $\Delta P = P - P'$ and $\Delta n = n - n'$. The first term captures the future value of the initial principal reduction, and the second term captures the future value of reduced debt service cost. Note that $\Delta D_0(n') = T_m$, i.e., the initial gap in mortgage debt is equal to the transfer amount allocated to debt reduction.

Additionally, transfer wealth from mortgage reductions is shaped by the savings decisions that recipients make in every period. Concretely, if recipients pay lower debt service cost, they can use the periodic savings in two ways: to invest in liquid assets, or to finance additional consumption. Assuming that of each euro that is saved in debt service, a constant fraction $\varphi \in [0,1]$ is invested

[7] In practice, reducing the principal will lower mortgage interest rates. Incorporating this would imply that—for a constant term n—ΔP would increase, as debt service costs decrease even more. Therefore, in reality, the reductions in debt service cost that transfer recipients enjoy are even larger than the ones derived here.

[8] If transfer recipients decide to keep debt service payments unchanged ($P' = P$), they will amortize more quickly ($n' < n$). Concretely, the implied term is given by $n' = -\ln(1 - \frac{r_m {D_0}'}{P'})/\ln(1+r_m)$. In contrast, if they decide to keep the same term ($n' = n$), their debt service cost will be lower ($\Delta P > 0 \iff P' < P$).

every period, additional liquid savings at time t will amount to:

$$\varphi \, \Delta P \, \frac{(1+r_s)^t - 1}{r_s}. \tag{3.7}$$

Combining all terms, total transfer wealth at time t equals:

$$\begin{aligned} \Omega_t = {} & (1+r_h)^t \, T_h + (1+r_s)^t \, T_s + (1+r_m)^t \, T_m \\ & - \Delta P \left[\frac{(1+r_m)^t - 1}{r_m} - \varphi \, \frac{(1+r_s)^t - 1}{r_s} \right]. \end{aligned} \tag{3.8}$$

As this decomposition shows, how transfer wealth continues to grow depends on the initial allocation of funds, as well as on the behavioral response of the recipient. The first three terms highlight how parental transfers mechanically promote further wealth accumulation, regardless of how they are used: additional real estate or liquid assets yield increased capital gains, and deleveraging is profitable if mortgage debt is costly. However, the fourth term captures the decline in transfer wealth if recipients lower their monthly debt service cost. Effectively, if first-time buyers without parental support make higher periodic mortgage payments, they close the wealth gap over time by accumulating home equity more quickly. This does not occur in a linear way, but slowly at first and faster later on, in line with the way mortgage amortization schedules are structured. Transfer recipients can prevent this in two ways: either by not reducing their debt service cost too much and opting for faster amortization, or by offsetting their lower debt service with liquid savings. If $r_s > r_m$ and φ is sufficiently large, transfer wealth may actually increase, rather than decrease, over time, if recipients lower their mortgage payments.

To highlight the different possibilities, I simulate the mortgage-related components of transfer wealth based on the standard product in the Dutch market: a 30-year, fully amortizing mortgage with monthly payments. Consider two extreme scenarios: on the one hand, transfer recipients may leave the term of their mortgage unchanged—thereby lowering their monthly cost—and fully consume their monthly savings. On the other hand, they may lower the term of their mortgage to the extent that their monthly debt service costs stay unchanged ($\Delta P = 0$). While the mortgage gap is eventually closed in both

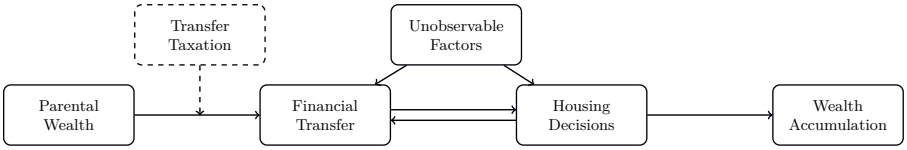


Figure 3.7: The Causal Relationships Around Transfers, Housing, and Wealth

Notes: This figure visualizes the causal relationships between parental wealth, parental transfers, housing outcomes, and wealth accumulation in a directed acyclic graph (DAG). Arrows indicate the direction of causality. Transfer taxation, visualized in a dashed box, is a moderating factor for the relationship between parental wealth and transfers.

scenarios, they imply very different wealth accumulation patterns, as visualized in Figure 3.16. In the same-term scenario without offsetting investments, transfer wealth declines steadily over the course of the mortgage. In the fast amortization scenario, transfer wealth continues to grow, as transfer recipients build up equity at a faster rate. Only after recipients have fully amortized (dashed line) can non-recipients begin to catch up, as they continue to pay off their debt.

3.4.2 Identifying the Effect of Home Purchase Transfers

In Figure 3.7, I present a simple directed acyclic graph (DAG) illustrating the causal relationships between parental transfers, housing market decisions, and future wealth accumulation—the main focus of this section. Identifying the effect of parental transfers on housing decisions and subsequent wealth accumulation is challenging for at least two reasons. First, within the extended family, transfer and home purchase decisions are likely to be jointly determined, making it difficult to disentangle the direction of causality. While parents control the supply of financial transfers, children determine their demand for them based on their housing preferences. For example, children who identify a dream home that exceeds their budget may request a (larger) transfer from their parents. Second, even after conditioning on observable characteristics, transfer recipients may differ from other buyers in unobserved ways that also influence their housing market behavior and economic outcomes. For instance, those receiving transfers might be better informed about the housing market, leading them to make systematically different purchase decisions. For these reasons, simple regressions of housing market outcomes on received transfer amounts are unlikely to capture the underlying causal mechanisms.

The supply of financial transfers is ultimately driven by the availability of parental assets that can be converted into cash. Isolating exogenous variation in parental resources would therefore help identify the causal effect of a financial transfer. For instance, Daysal et al. (2023) study variation in parental wealth due to differences in house price appreciation using a long administrative panel. However, this approach is less suitable for the setting I study. In particular, there is likely insufficient variation in parental housing returns to generate meaningful changes in transfer behavior. Moreover, converting housing wealth into financial transfers requires extracting home equity, which is less common in the Netherlands than in other countries (Benetton et al., 2024).

In contrast, the Dutch setting offers variation not in parental resources per se, but in how financial transfers from parents to children are taxed. As shown in Figure 3.7, the taxation of transfers moderates the relationship between parental wealth and financial transfers. In late 2013, the Dutch government implemented a policy reform that altered the taxation of transfers for housing-related expenditures (e.g., home purchases, mortgage prepayments, and renovations). Prior to the reform, the special exemption limit for housing-related transfers was set at €51,407. In October 2013, as part of a broader effort to stimulate the housing market, the government temporarily increased this one-time exemption limit to €100,000 for transfers earmarked specifically for housing. In practical terms, this implied potential tax savings of approximately €5,000 for parents who fully utilized the exemption, relative to the pre-reform regime. From 2015 onward, the special exemption limit was reduced to its pre-reform level.

During the initial phase of the increased exemption, the measure received wide media attention and triggered a surge in tax-free gifts, with around 150,000 recorded transfers. While many of these transfers were used to finance first-time home purchases—the focus of my analysis—the primary policy goal was to help existing homeowners with underwater mortgages following the prolonged housing market downturn after the global financial crisis. Between 2017 and 2022, the policy was reintroduced in a slightly modified form. However, according to Statistics Netherlands, take-up rates during this second policy window remained well below those observed in the initial rollout phase (Statistics Netherlands, 2021).

To study the effect of the initial policy rollout on financial transfers—and, in turn, on housing market outcomes—I make several adjustments to the dataset.

First, I restrict the sample to the years 2011–2016, surrounding the policy reform. Next, I collapse the original person-by-year panel into a repeated cross-section of potential first-time buyers, retaining one observation per individual. Individuals who make a first-time home purchase between 2011 and 2016 are assigned to the year of their purchase. For all others—those who remain renters throughout this period—I randomly assign a year between 2011 and 2016. This procedure allows me to compare, within each year, housing market outcomes of transfer recipients and non-recipients. Crucially, because this dataset originates from the underlying panel data, I can include information on outcomes measured up to 7 years after the observation year. For example, for an individual observed in 2016, I track household wealth information from 2016 (year t) through 2023 (year $t+7$) in separate columns. Similarly, for a person observed in 2014, the corresponding entries for year t and $t+7$ track her wealth in 2014 and 2021, respectively.

Empirically, I attempt to investigate the causal effect of parental transfers targeted for first-time home purchases—which I will refer to as *home purchase transfers* for brevity—on housing decisions and future wealth accumulation.[9] A straightforward way to estimate the effect of home purchase transfers on housing and wealth outcomes would be the following:

$$Y_{it}^{t+k} = \delta_t + \beta \, \mathit{Transfer}_{it} + \mathbf{X}_{it} + e_{it}, \tag{3.9}$$

where $\mathit{Transfer}_{it}$ indicates the amount of parental transfer that person i has received in year t, Y_{it}^{t+k} denotes a housing or wealth outcome from year $t+k$, $\mathbf{X}_{it}$ contains a vector of control variables and δ_t is a time fixed effect. However, for the reasons outlined above, this approach suffers from endogeneity and cannot reveal a causal effect. Therefore, my empirical strategy for identifying the effect of transfers on housing market outcomes is based on the relationship between parental wealth and the taxation of parental transfers. The policy change announced in late 2013 creates exogenous variation in this relationship, which I will capture in a two-stage least squares (2SLS) framework. Before introducing the empirical strategy formally, I visualize the underlying intuition

[9] I focus on targeted transfers, since the policy change related specifically to transfers made for housing-related expenditures. Conceptually, this is distinct from studying the effect of untargeted transfers, i.e., pure cash gifts. In a targeted transfer, the causal link between the transfer and the decision to make a home purchase is established mechanically. Therefore, this setting is not suitable to estimate how the probability of making a purchase is increased upon receiving a transfer.

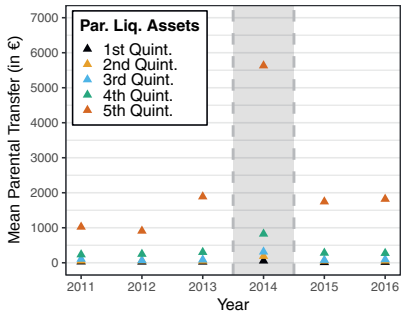
in Figure 3.8.

In 2014, when the exemption limit was raised from around €50,000 to €100,000, there is a substantial increase in parental transfers. However, that is the case only for people with wealthy parents (measured via parental liquid assets in 2011). For people with parents in the top quintile of liquid assets (orange triangles), the unconditional net transfer increases to €6,000 in that year. In 2015, when the exemption limit was reduced again, this number drops back to less than €2,000. In contrast, for those with less affluent parents, transfers respond substantially less to the special tax-exemption policy, and for those in the bottom quintile, they remain virtually flat around zero. This unequal increase in transfers in 2014 captures the first stage in the 2SLS framework. Similarly, the share of individuals making a first-time purchase increases in 2014, relative to other years, but this increase is driven by those with affluent parents.

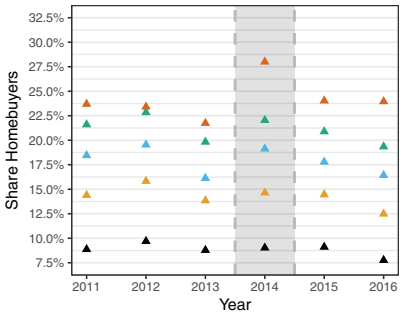
The last two panels are based on the buyer-only sample and depict the intensive-margin response to the policy. Panel (c) suggests that, among first-time buyers, the policy led to a significant increase in down payments: buyers with wealthy parents enter the housing market with lower loan-to-value (LTV) ratios in 2014, compared to other years.[10] In panel (d), average property values across all first-time buyers are shown. Here, no visible pattern stands out with respect to the impact of the policy. At first glance, this suggests that, at the intensive margin, parental transfers mainly lead to a deleveraging of mortgage debt.

Of course, these visual patterns merely represent unconditional averages and should not be interpreted as causal effects per se. To formally identify the effect of parental transfers on housing outcomes, I estimate a two-stage least squares (2SLS) model. In this framework, the receipt of a parental home purchase transfer is instrumented using an interaction between parental liquid assets (measured in 2011) and an indicator for the year 2014, when the exemption policy was in effect. This interaction captures plausibly exogenous variation in transfer receipts stemming from the policy's differential impact across the parental wealth distribution. I estimate the following model:

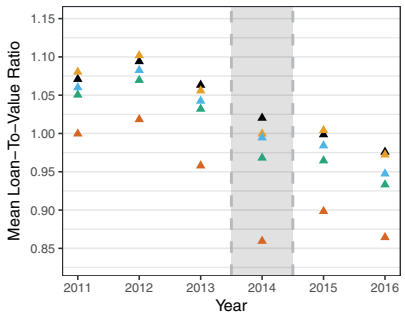
[10] Strikingly, in the early years of the sample, LTV ratios reach values of up to 110%. To some extent, this reflects imprecise measurement, as I calculate LTV ratios based on tax-assessed values, rather than the unobserved market values that mortgage lenders use. However, the statutory LTV limit in 2012 was as high as 106%. In the following years until 2018, that limit was reduced by 1 pp per year.



(a) Net Transfers



(b) Share First-Time Buyers



(c) LTV Ratio At Purchase



(d) Property Value At Purchase

Figure 3.8: Housing Market Outcomes Around the Policy Reform

Notes: This figure shows transfers and housing outcomes between 2011 and 2016, depending on parental liquid assets (measured in 2011). Panels (a) and (b) are based on the full sample, panels (c) and (d) consider first-time homebuyers only. Non-buyers are randomly assigned to a year, based on a uniform draw between 2011 and 2016. The year 2014, highlighted in grey, marks the period in which the special tax exemption for housing-related transfers was raised to €100,000.

$$Transfer_{it} = \delta_t + \pi \left(LiquidAssets^{Par}_{i,2011} \times \mathbf{1}\{t = 2014\} \right) + \mathbf{X}_{it} + u_{it} \quad (3.10)$$

$$Y^{t+k}_{it} = \delta_t + \beta \, \widehat{Transfer}_{it} + \mathbf{X}_{it} + \varepsilon_{it}, \quad (3.11)$$

where I instrument for the net financial transfer that person i receives in year t using an interaction between her parents' liquid assets, measured in 2011, and an indicator for the year 2014. In the second stage, predicted transfers are linked to various housing decisions and, in the next step, to downstream wealth outcomes. In both models, I include a granular set of controls and fixed effects, all measured in 2011: year-by-residential-location (NUTS3 region) fixed effects; a single-household dummy; gender; highest degree; birth year. Moreover, I control for the following wealth and income variables, measured in 2011 and captured in percentile bins: household income, household liquid assets, parental household liquid assets and household net wealth. Since I control for parent household liquid assets, the variation is driven by how people with wealthy parents behave differently during the policy period, relative to other years.

The relevance of the constructed instrument is supported by the visual evidence in Figure 3.8 and will be formally confirmed in the first-stage regressions. In contrast, the exclusion restriction cannot be tested directly. In this context, it requires that the housing decisions—and subsequent wealth outcomes—of children with wealthy parents observed in 2014 differ from those of others only because they received larger parental transfers. Given the salience of the policy change at the time, it is plausible to assume that any increase in parental support is captured by a higher propensity to make cash transfers, rather than by other forms of assistance such as co-signing mortgage applications.

Nonetheless, the exclusion restriction could be violated if other concurrent policy changes or market trends affected wealthy parents' children differently from others observed in the same or nearby years. For example, in 2013, the Dutch mortgage interest deduction rules were reformed (Bernstein and Koudijs, 2024), and throughout the sample period, lending limits were gradually tightened (Rouwendal and Petrat, 2022). However, there is little reason to believe that these changes uniquely affected children of wealthy parents in 2014. A key strength of the design is that the tax exemption reform was temporary and centered in the middle of the sample period. As a result, only those with wealthy parents in 2014 are "treated," while those with equally wealthy parents observed before or after serve as a control group. This feature mitigates concerns about linear time trends that might otherwise affect the relative outcomes of children from wealthier versus less wealthy families.

However, a remaining concern is that due to the policy, buyers may select into the year 2014, in order to make use of the higher exemption limit. This may

introduce selection bias for the intensive margin analysis, as it involves conditioning on a variable (making a home purchase) that is itself affected by the treatment (receiving a parental transfer). Concretely, selection bias arises if the marginal homebuyer that is induced to buy during the policy period differs from other homebuyers in ways that are related to the outcomes of interest. In that case, the estimated effects do not only capture the local average treatment effect (LATE) of parental transfers on the behavior of homebuyers, but also reflect compositional changes in the sample of buyers. To address this, I implement inverse probability weighting, whereby the probability of making a purchase during the policy period (rather than in another year) is estimated on all observable characteristics (Wooldridge, 2002; Carry et al., 2021).[11]

3.4.3 Transfers and Housing Market Outcomes

To examine how parental home purchase transfers are used, I estimate the 2SLS model from Equation (3.11). The resulting coefficients are reported in Table 3.2. These estimates are compared with results from an OLS model including similar control variables, in which each outcome variable is regressed on the net parental transfer received by the individual in a given year.

First, I estimate the extensive-margin effect of parental transfers by examining how the probability of making a first-time home purchase changes with each additional euro of transfer received. This estimation is based on the full sample, including individuals who do not make a purchase between 2011 and 2016. The dependent variable is an indicator equal to one if a first-time purchase occurs in a given year. Among non-recipients, the mean annual probability of making a first-time purchase is 16.6%.

According to the IV estimates, this probability increases by about 0.9 percentage points for every €1,000 in parental transfers received. However, as discussed above, this relationship is mechanical, as the policy supported only transfers that were earmarked for housing. The extensive-margin effect is therefore reported primarily for completeness and to illustrate the empirical relevance of the policy change.

[11] Concretely, I estimate probabilities $\hat{p}$ from a logistic regression of the probability of making a purchase in 2014, based on all covariates and fixed effects listed above, excluding the year FE. Next, weights are obtained by calculating $w = \Pr(year = 2014)/\hat{p}$ for all those who buy in 2014 and $w = \Pr(year \neq 2014)/(1 - \hat{p})$ for those who buy in another year, whereby $\Pr(year = 2014)$ is the actual share of buyers making their purchase in 2014.

Table 3.2: Transfers and Housing Outcomes

	Ext. Margin	*Int. Margin*			*Mortgage Outcomes*	
	(1) Purchase (in %)	(2) Home Value	(3) Mortgage Size	(4) Δ Liq. Assets	(5) Interest Rate	(6) Repay. Rate
Panel A: OLS						
Net Transfer (€)	0.0008***	0.050	−0.700***	0.035***		
	(0.0001)	(0.040)	(0.039)	(0.002)		
Transfer (Y/N)					0.000	0.005*
					(0.000)	(0.003)
Observations	288,084	26,504	28,087	28,087	27,265	27,658
Adjusted R^2	0.200	0.467	0.444	0.350	0.376	0.082
Panel B: IV						
Net Transfer (€)	0.0009***	−0.100	−1.004***	−0.044		
	(0.0001)	(0.194)	(0.172)	(0.070)		
Transfer (Y/N)					−0.006**	0.031
					(0.003)	(0.025)
Observations	288,084	26,504	28,087	28,087	27,265	27,658
Adjusted R^2	0.200	0.467	0.442	0.349	0.367	0.076
First Stage F-Stat	6,804	1,339	1,393	1,393	345	344

Notes: This table shows the OLS and IV estimates of the effect of receiving a parental transfer on various housing market outcomes. IV estimates are based on the 2SLS model specified in Equation (3.11). Model (1) is estimated on the full sample, including non-buyers. Models (2)–(6) are estimated on the buyer-only sample. All models include the following controls as categorical variables: residential location × transaction year; gender; single household dummy; birth year; highest degree; household income percentile; household liquid assets percentile; household net wealth and parental household liquid assets percentile (all measured in 2011). Heteroskedasticity-robust standard errors are reported in parentheses (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

Next, I turn to the intensive-margin response to a transfer. All outcome variables are measured in euros, allowing a straightforward interpretation of treatment effects: *for each euro received in transfers, how many cents are allocated to purchasing a more valuable home, reducing the mortgage balance, and increasing liquid savings?* Thus, the estimated effects can be directly mapped to the channels introduced in the previous section. Investments in additional real estate (T_h) are captured by the tax-assessed value of a purchased property, observed in the year prior to the transaction. Investments in mortgage reduction (T_m) are measured at the household level as the outstanding balance at the beginning of the following year. Finally, liquid savings (T_s) are measured

as the change in liquid assets (bank deposits plus stocks and bonds) from the purchase year to the subsequent year. The resulting coefficients are reported in columns (2) to (4) of Table 3.2.

The results indicate that the intensive-margin response to parental transfers is driven entirely by the mortgage channel. The IV estimates suggest that every euro of parental transfer leads to a one-for-one reduction in mortgage principal. This compares to a reduction of 70 cents per euro based on the OLS model. In other words, transfer recipients appear to use the entirety of the funds to deleverage. In contrast, I find no evidence that transfer receipts are spent on buying more valuable homes or increasing liquid savings: the effect of transfers on assessed values is insignificant in both models. For liquid assets, only the OLS model suggests a significant but small increase of 3.5 cents per euro of transfer.

In columns (5) and (6), I report two additional outcomes that provide further evidence on the mortgage channel. In both cases, the parental transfer is measured as an indicator variable rather than by amount, to ease interpretation. First, I examine mortgage interest rates. Since rates are not directly observed, I impute them by dividing total mortgage interest payments in year $t + 1$ by the mortgage principal outstanding at the beginning of that year. While the OLS estimates show no difference, the IV results indicate that transfer recipients pay approximately 0.6 percentage points lower interest rates on their mortgages. Hence, their debt service costs are reduced not only because they take on less debt, but also because the cost of that debt is lower.

Next, I study the amortization schedule chosen by transfer recipients. Because the data are right-censored, I cannot observe the full mortgage term. I therefore proxy for the repayment rate as the reduction in mortgage principal during the first year after purchase divided by the initial principal.[12] The OLS estimates suggest higher repayment rates among transfer recipients, but the 2SLS estimates are statistically insignificant. Thus, it cannot be rejected that recipients retain similar mortgage terms rather than opting for faster amortization.

[12] Most Dutch mortgages during this period were either fully amortizing or included an interest-only component, whereas linear mortgages were uncommon. Consequently, repayment rates are not constant over the life of the mortgage.

3.4.4 Transfers and Future Wealth Accumulation

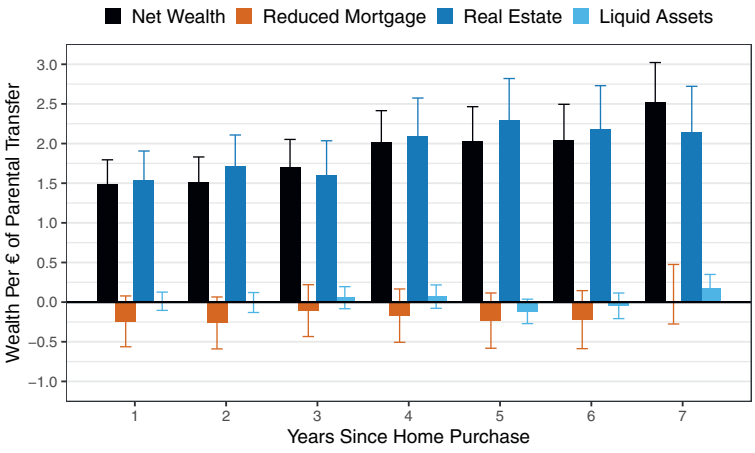
I now examine how this allocation of funds translates into subsequent patterns of wealth accumulation. To this end, I reestimate the 2SLS model from Equation (3.11), using household net wealth in year $t + k$ after the purchase as the outcome variable, where k ranges from 0 to 7. This allows me to track the evolution of transfer wealth over the seven years following the purchase. I also estimate future levels of real estate assets, liquid assets, and mortgage debt. This allows for an empirical decomposition of wealth changes into its different components. Since the (instrumented) treatment variable in each model is the size of the parental home purchase transfer, the coefficients are readily interpretable as the change in wealth in year $t + k$ per euro of transfer received in year t.

The analysis is done separately for two different counterfactual groups. The first is the full sample of not-yet-homeowners observed between 2011 and 2016, including those who remain renters. The second is the smaller sample of first-time buyers only. Making this distinction is important because, due to the nature of the quasi-experiment that I examine, one important counterfactual outcome is unobserved: whether or not transfer recipients would have bought a home absent the transfer.[13] If they would have bought a home anyway, the relevant counterfactual group consists of other first-time buyers in the same year. Only if they would have remained renters otherwise is the relevant comparison that with the full sample including never-buyers. I therefore present the empirical results for both counterfactual groups, before discussing which one is likely more appropriate.

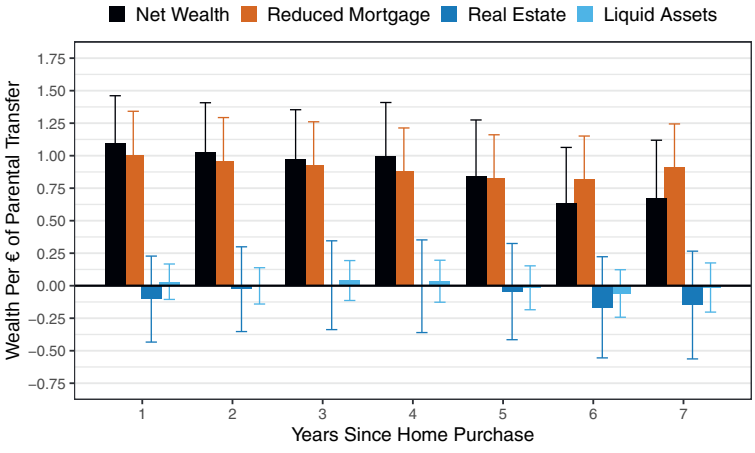
Figure 3.9 presents the second-stage results, using various forward-looking wealth measures as dependent variables. As expected, relative to the full sample that includes non-buyers, recipients of a home purchase transfer exhibit higher net wealth in the following year, driven entirely by an increase in real estate assets. Over time, this wealth advantage continues to grow as the gap in real estate holdings widens. By year seven, the estimated wealth difference between transfer recipients and non-recipients reaches approximately €2.50 for every euro of initial parental transfer.

While this represents a substantial increase, it should be noted that during the sample period the Dutch housing market recovered from its 2013 low to a

[13] This is due to the fact that the policy change incentivized parents to make targeted home purchase transfers, rather than unconditional cash gifts.



(a) Full Sample



(b) Buyer-Only Sample

Figure 3.9: Transfers and Wealth Accumulation (IV Estimates)

Notes: This figure shows the estimated evolution of *transfer wealth*, i.e., the gain in net wealth per euro of parental home purchase transfer, in the first seven years after the purchase. Total transfer wealth is decomposed into three components: reduced mortgage debt, increased real estate assets and increased liquid savings. Panel (a) is based on the full sample including non-buyers, whereas panel (b) is estimated on a buyer-only sample. All coefficients are based on 2SLS estimates from Equation (3.11). Error bars represent 95% confidence intervals, based on heteroskedasticity-robust standard errors.

record high in 2022. Consequently, it is unsurprising that, relative to a sample largely composed of non-buyers, individuals who purchase a home with parental support experience substantial housing gains. In contrast, there are no differences in mortgage debt, consistent with the earlier finding that transfers lead to a one-for-one reduction in mortgage principal.

Comparing these results to a simple OLS model, shown in Figure 3.17, highlights a key limitation of OLS in this context: because transfer recipients mechanically have higher mortgage debt than non-buyers, OLS would misleadingly suggest that transfers causally increase mortgage borrowing, when in reality they do not. Finally, estimated differences in liquid assets remain flat, indicating that transfer wealth is fully invested in real estate, with no evidence of subsequent investment in other assets.

Turning to the buyer-only analysis, shown in panel (b), a different picture emerges. In the first year after the purchase, the estimated wealth differences mirror the intensive-margin results from Table 3.2: compared to other homebuyers, parental transfers increase recipients' net wealth one-for-one, and this is exclusively driven by a reduction in mortgage borrowing. However, in the following years, the initial gap in net wealth narrows gradually. By year seven, the estimated gap in mortgage debt per euro of transfer received falls to approximately €0.85. This suggests that homebuyers without parental support close the initial mortgage gap by amortizing faster. In Figure 3.18, I illustrate this by estimating the reduction in total debt service cost (interest expenses plus principal repayments) per euro of transfer received, using the 2SLS model. Transfer recipients enjoy lower debt service cost, both due to lower interest expenses and smaller principal repayments.[14] In other words, by making larger principal repayments, first-time buyers without parental support slowly close the gap in outstanding mortgage debt, relative to transfer recipients.

Crucially, I find no evidence that the reduction in debt service costs is used to accumulate other savings. Differences in real estate assets remain near zero across all years, indicating that transfer recipients are not more likely to make a secondary purchase or upgrade to a more valuable home over time. Likewise, differences in liquid assets remain virtually flat, suggesting that the reduced debt service is not used to invest in securities or increase bank balances. In terms of the transfer wealth decomposition introduced above, this implies that

[14] While interest expenses are reported directly from the income data, principal repayments are imputed based on changes in mortgage debt between year t and $t+1$. Therefore, the estimates are less precise.

φ is close to zero. Conversely, the marginal propensity to consume (MPC) out of reductions in debt service costs is close to one. This finding aligns with Ganong and Noel (2020), who argue that a reduction in mortgage debt only affects consumption if it is accompanied by a reduction in monthly debt service payments. I therefore conclude that recipients use their increased residual income—net of debt service costs—to finance higher levels of consumption.

For comparison, results from the OLS model are reported in panel (b) of Figure 3.17. Unsurprisingly, the OLS estimates show a more muted pattern, reflecting the model's smaller predicted effect on initial mortgage balances. Qualitatively, however, the OLS results similarly suggest that the initial mortgage debt gap closes over time—albeit much more slowly—while no compensating investment in liquid savings occurs.

Taken together, the findings therefore imply that, when compared with other first-time buyers, parental transfers primarily frontload wealth. Initially, recipients use the additional funds to reduce their liabilities rather than to make additional investments or finance consumption. Over time, however, consumption increases as recipients draw on the additional disposable income made available through lower debt service costs.

3.4.5 Discussion of the Findings

A primary goal of this empirical exercise was to better understand how parental transfers contribute to wealth inequality in the longer run. If transfer wealth decreases over time, this would imply that the contribution of transfers to longer-run wealth inequality may be lower than previously anticipated. If, instead, transfer wealth increases because recipients accumulate additional assets in the future, transfers may amplify wealth inequality even more in the longer run.

As the empirical analysis has shown, which of these scenarios applies comes down to the question of what counterfactual is the relevant one. This in turn depends on when—if at all—transfer recipients would have bought their first home in the absence of parental support. Consider both extremes: if recipients would have otherwise never made a home purchase, the extensive margin effect of a transfer is dominant and recipients should be compared to a full sample including non-buyers. If, instead, transfers do not affect the timing of the purchase, the extensive margin effect is negligible. The truth likely lies between

these extremes. It seems plausible that most aspiring homeowners eventually purchase a home, particularly those with wealthy parents. More concretely, it is unlikely that parental transfers lead individuals to become homeowners many years earlier than they would have otherwise. For instance, Engelhardt and Mayer (1998) find that recipients of parental transfers purchase their first home only nine months earlier than comparable individuals without financial support. Consequently, I view the intensive-margin effect, reported in Figure 3.9, as the main empirical result of this analysis. However, in countries where starter homes are particularly unaffordable or loan-to-value limits are stricter, the extensive-margin effect may be relatively more important.

This implies that, at least in practice, cash transfers for home purchases are ineffective as a vehicle for dynastic wealth building. Instead, their main effect seems to be to frontload wealth, thereby allowing higher consumption earlier in the life cycle. These findings contrast with those by Wold et al. (2024), who argue that housing is a key channel for intergenerational wealth persistence. The key difference is related to leverage: Wold et al. (2024) find that homebuyers with wealthier parents take on more leverage, likely due to the fact that additional collateral or parental guarantees relax borrowing constraints. In contrast, I show that the direct effect of parental cash transfers is to reduce leverage.

One remaining question, beyond the scope of this empirical analysis, is whether the observed pattern of wealth frontloading is intentional. On the one hand, the decision to make a transfer and allocate it toward mortgage principal reduction may be part of a strategy to smooth consumption over the child's life cycle. By reducing debt service costs, investing a transfer in this way increases residual monthly income available for consumption at a time when liquidity constraints likely still bind for many young homeowners.

On the other hand, the observed wealth accumulation patterns could reflect the fact that most people accumulate wealth passively, with mortgage amortization serving as a "default" mechanism (Thaler and Sunstein, 2021). Bernstein and Koudijs (2024) document this mechanism explicitly in the Dutch housing market context: as mortgages amortize faster, individuals accumulate more wealth because the additional monthly investment into home equity is not offset by reductions in other forms of savings. The narrowing of wealth differences between transfer recipients and other homebuyers observed here may reflect the same phenomenon. In this case, transfer recipients unconsciously adapt their spending patterns to the lower monthly debt service costs, while

non-recipients reduce spending to meet their higher amortization targets.

3.5 Conclusion

Using rich administrative data from the Netherlands, I study parental inter vivos transfers from two perspectives, examining both their drivers and their consequences. Compared with other major life events, first-time home purchases clearly stand out as the dominant driver of parental cash transfers. In the year of the first-time purchase, the mean likelihood of receiving a registered parental transfer increases by six percentage points—a tenfold increase relative to individuals who experience no life event. On average, first-time homebuyers receive €3,200 in additional parental transfers in the year of the purchase, followed by annual transfers of around €300–400 over the subsequent eight years.

Transfers are strongly concentrated among parents with high levels of liquid assets. While first-time buyers receive around €1,000 in parental support at the median, those with parents at the 95th and 99th percentiles of liquid assets receive €10,000 and €34,000, respectively. Similarly, children who are wealthier to begin with receive higher levels of parental transfers, suggesting that these transfers increase absolute wealth inequality in the short run. In contrast, parents do not make transfers following an involuntary job loss, although some transfers occur around marriage. Regarding the birth of the first grandchild, there is evidence that parents provide informal monetary support by covering some of their child's expenditures.

Motivated by the dominance of transfers around first-time home purchases, I next examine how receiving a parental home purchase transfer shapes housing market decisions and, in turn, downstream wealth accumulation. To do so, I leverage quasi-random variation from a Dutch policy that exogenously increased the tax exemption limit for housing-related transfers. Using an IV design based on parental liquid wealth prior to the policy change, I find that, compared with other homebuyers without parental support, recipients use the windfall from a parental transfer in only one way: to reduce their initial mortgage principal.

Transfer recipients do not appear to purchase more valuable properties or increase their liquid savings. This generates wealth frontloading: in the years following the purchase, the net wealth gap between transfer recipients and

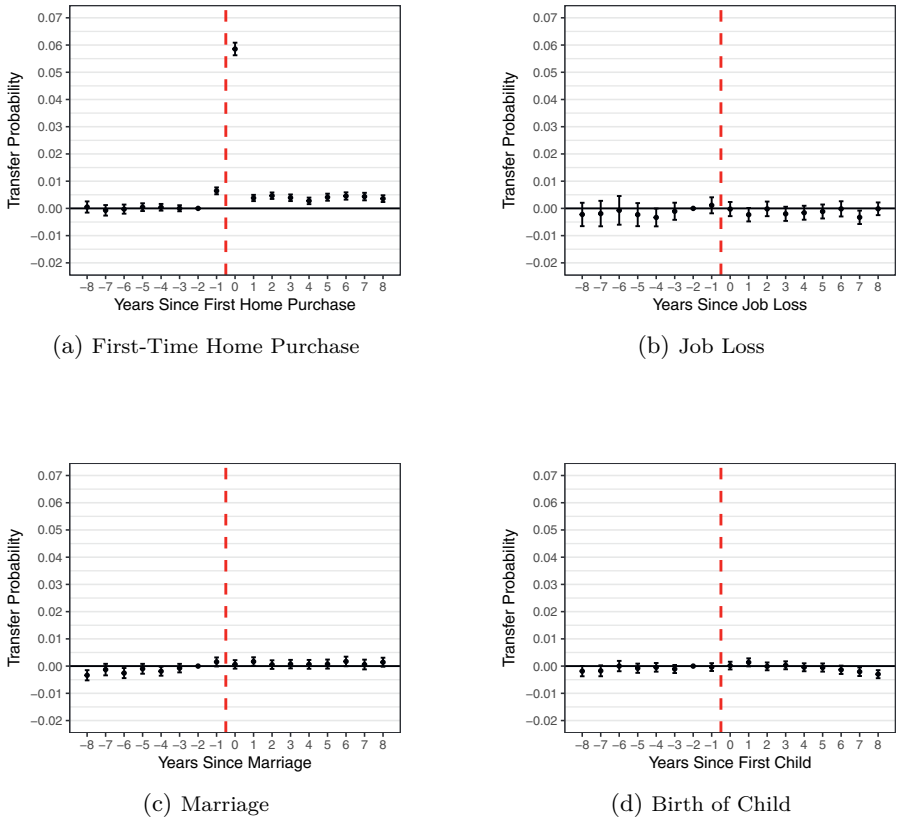
other homebuyers narrows. This occurs because transfer recipients, benefiting from lower debt service costs due to their reduced mortgage principal, do not invest more in liquid assets. In other words, while transfer recipients enter homeownership with lower liabilities, non-recipients gradually close the gap by amortizing their mortgages at a faster rate.

These findings have several important implications. First, the spike in parental transfers around first-time home purchases underscores the need to better understand their underlying drivers. Rather than providing insurance against negative income shocks, parents appear primarily motivated to help their children overcome liquidity constraints when investing in homeownership. At the same time, the fact that transfers are used entirely to reduce mortgage debt suggests that parents may aim to facilitate consumption smoothing, rather than dynastic wealth accumulation. An alternative explanation is that, due to bounded rationality, parents (and their children) fail to fully account for mortgage amortization as a wealth-building mechanism.

Overall, the findings of this study add nuance to the discussion of parental transfers and wealth inequality. On the one hand, because transfers are received primarily by children who are already relatively well-off, they clearly increase short-run inequality. On the other hand, wealth frontloading suggests that the long-run impact of inter vivos transfers on inequality within the recipient generation may be more muted.

A Appendix

A.1 Combined Event Studies



(a) First-Time Home Purchase

(b) Job Loss

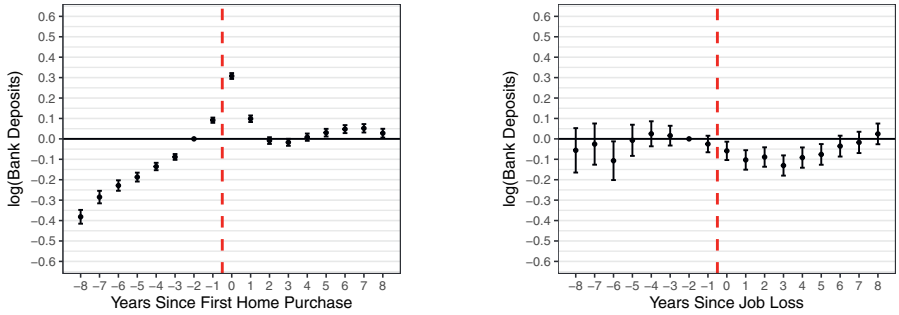
(c) Marriage

(d) Birth of Child

Figure 3.10: Parental Transfers Around Major Life Events (Combined Model)

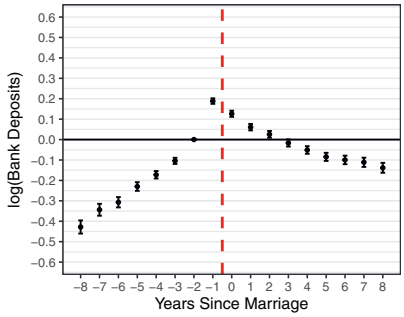
Notes: This figure presents event study coefficients estimated from Equation (3.2), where all four major life events are included in the same model. The dependent variable is a binary indicator equal to one if a parental transfer was received in a given year. The omitted period is $t-2$, corresponding to two years before the event occurs. Standard errors are clustered at the individual level, and the error bars represent 95% confidence intervals.

A.2 Own Bank Deposit Changes Around Life Events

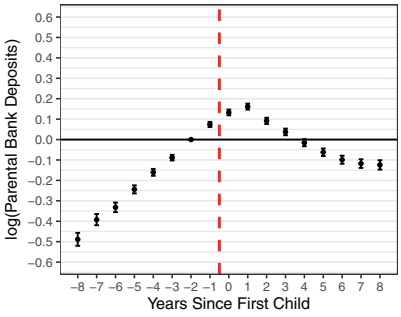


(a) First-Time Home Purchase

(b) Job Loss



(c) Marriage

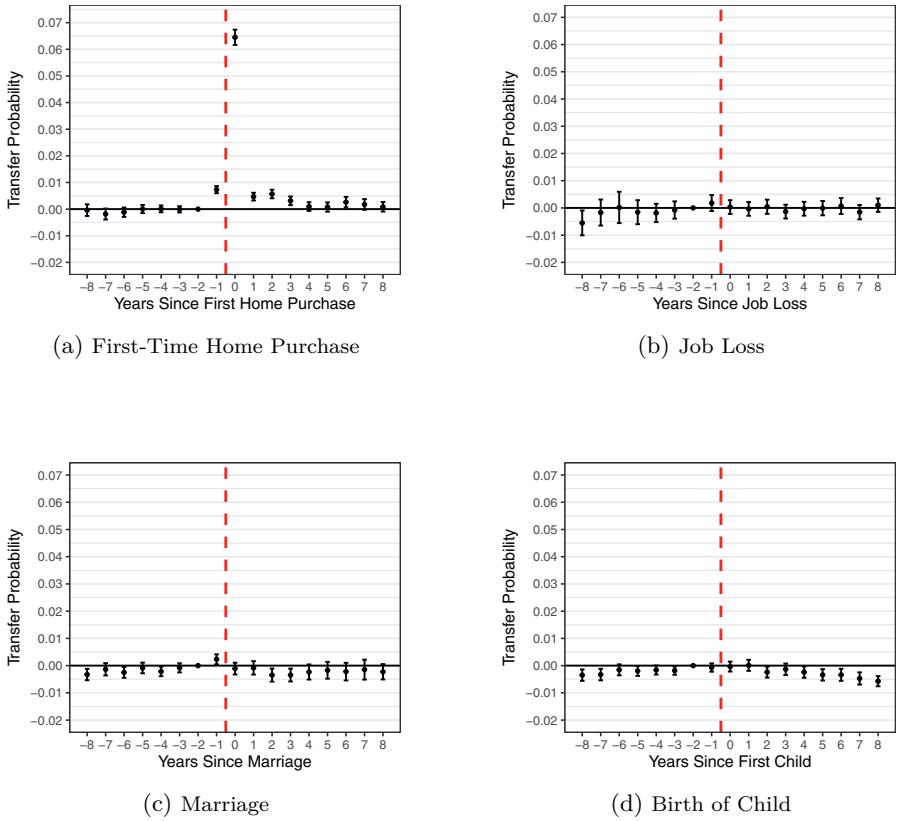


(d) Birth of Child

Figure 3.11: Household Bank Deposits Around Major Life Events

Notes: This figure presents event study coefficients estimated from Equation (3.1), where all four major life events are modeled separately. The dependent variable is the natural logarithm of household bank deposits in a given year. The omitted period is $t-2$, corresponding to two years before the event occurs. Standard errors are clustered at the individual level, and the error bars represent 95% confidence intervals.

A.3 Censored Event Studies



(a) First-Time Home Purchase

(b) Job Loss

(c) Marriage

(d) Birth of Child

Figure 3.12: Parental Transfers Around Major Life Events: Event Studies Removing Observations After Second Event

Notes: This figure presents event study coefficients, where all four major life events are included in a single model. Here, in comparison to Equation (3.2), observations are removed as soon as a second event occurs. This aims at isolating the effect of the first life event only. The dependent variable is a binary indicator equal to one if a parental transfer was received in a given year. The omitted period is $t-2$, corresponding to two years before the event occurs. Standard errors are clustered at the individual level, and the error bars represent 95% confidence intervals.

A.4 Transfers Across Years

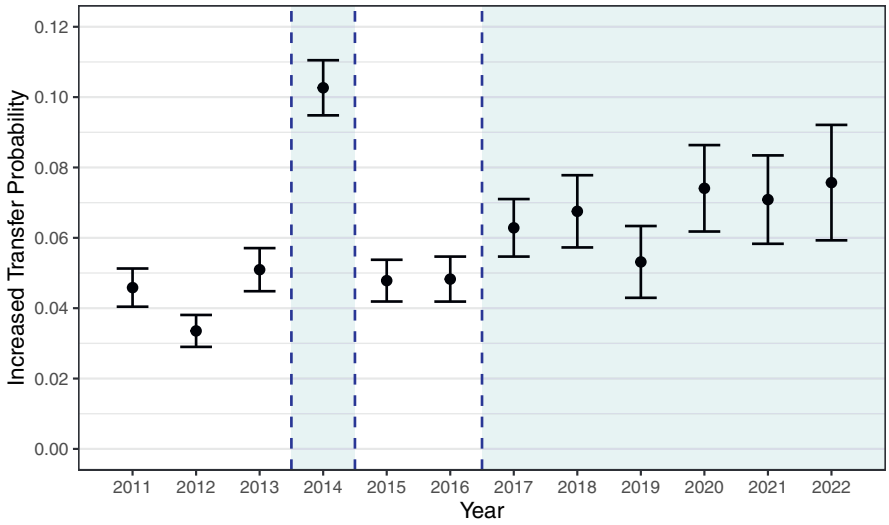
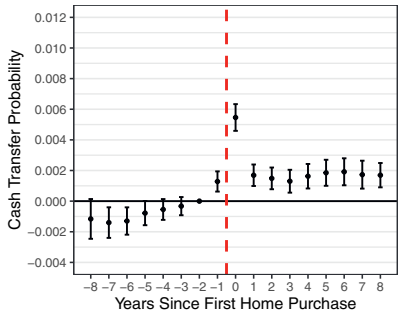


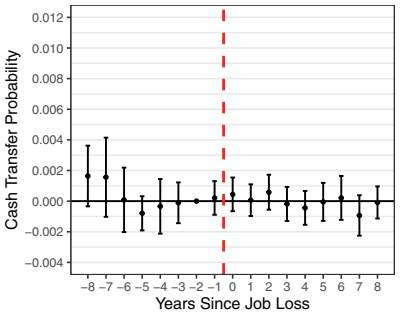
Figure 3.13: Home Purchase Transfers Across Years

Notes: This figure presents estimates from a model in which differences in parental transfers between homebuyers and others are compared across years. Years in which the special tax exemption threshold for housing-related transfers was raised are shaded in blue. Standard errors are clustered at the individual level, and the error bars represent 95% confidence intervals.

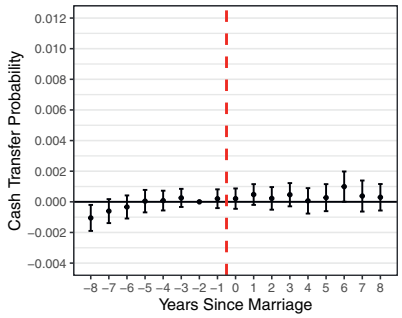
A.5 Excluding Tax-Exempt Transfers



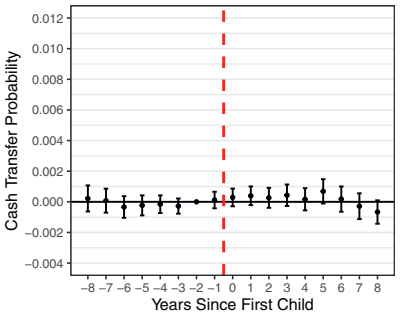
(a) First-Time Home Purchase



(b) Job Loss



(c) Marriage

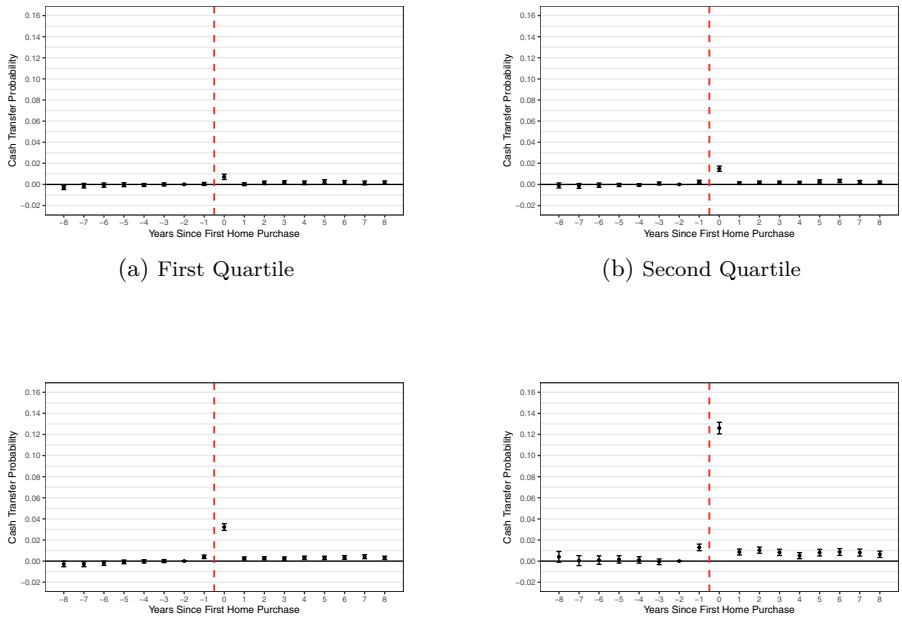


(d) Birth of Child

Figure 3.14: Regular (Non-Tax-Exempt) Transfers Around Major Life Events

Notes: This figure presents event study coefficients estimated from Equation (3.2), where all four major life events are included in the same model. The dependent variable is a binary indicator equal to one if a regular (taxed) parental transfer was received in a given year. The omitted period is $t-2$, corresponding to two years before the event occurs. Standard errors are clustered at the individual level, and the error bars represent 95% confidence intervals.

A.6 Transfers by Liquid Assets Quartile



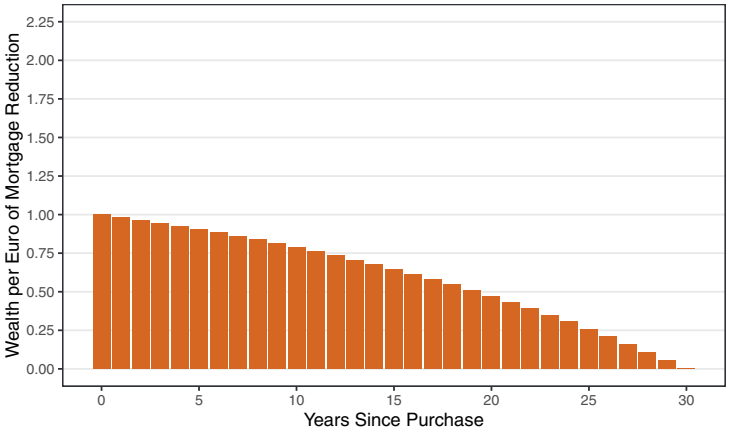
(a) First Quartile (b) Second Quartile

(c) Third Quartile (d) Fourth Quartile

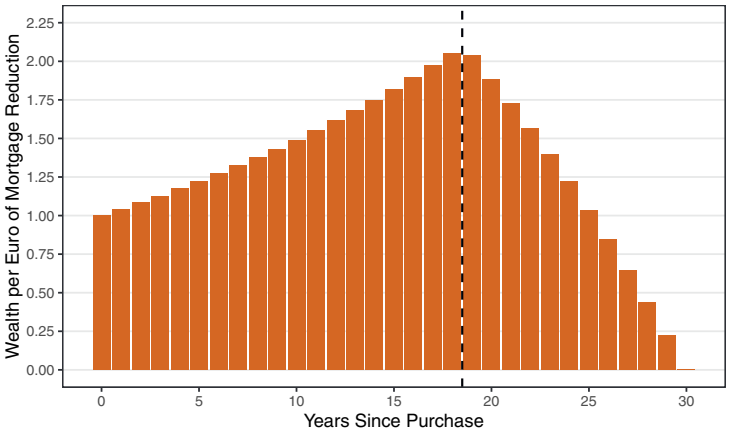
Figure 3.15: Transfers Around First-Time Home Purchase (by Parental Assets Quartile)

Notes: This figure presents event study coefficients estimated from Equation (3.1), using first-time home purchases as life event. The dependent variable is a binary indicator equal to one if a parental transfer was received in a given year. In the four different panels, the model is estimated separately in different quartiles of the distribution of parental liquid assets (measured in 2011). The omitted period is $t-2$, corresponding to two years before the event occurs. Standard errors are clustered at the individual level, and the error bars represent 95% confidence intervals.

A.7 Transfer Wealth Simulations



(a) Same-Term Scenario

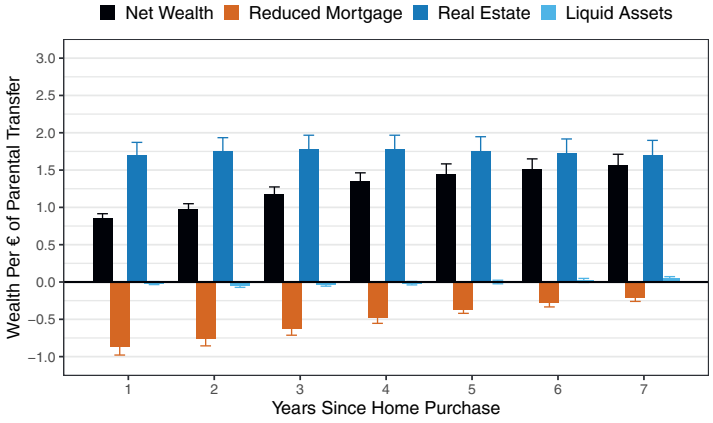


(b) Fast Amortization Scenario

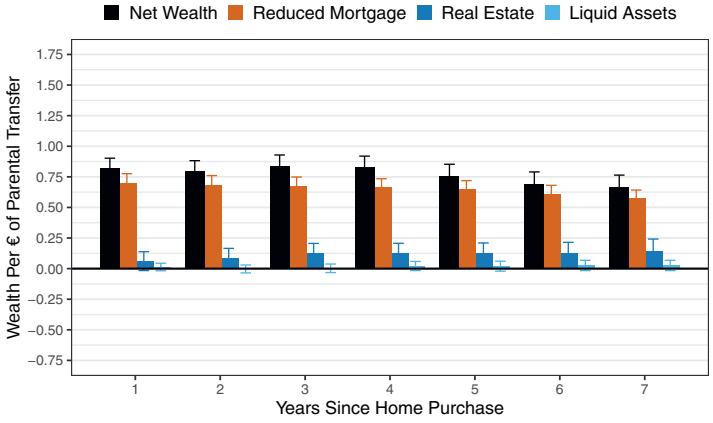
Figure 3.16: Simulated Transfer Wealth From Reduced Mortgage

Notes: This figure shows a simulation of the mortgage component of transfer wealth, based on two hypothetical scenarios: first, a same-term scenario, whereby transfer recipients do not reduce the term of their mortgage and do not reinvest reduced mortgage payments. Second, a fast amortization scenario, whereby transfer recipients reduce the term of their mortgage to the extent that their debt service costs remain unchanged.

A.8 Transfer Wealth Over Time: OLS Results



(a) Full Sample



(b) Buyer-Only Sample

Figure 3.17: Transfers and Wealth Accumulation (OLS Estimates)

Notes: This figure shows the estimated evolution of *transfer wealth*, i.e., the gain in net wealth per euro of parental home purchase transfer, in the first seven years after the purchase. Total transfer wealth is decomposed into three components: reduced mortgage debt, increased real estate assets and increased liquid savings. Panel (a) is based on the full sample including non-buyers, whereas panel (b) is estimated on a buyer-only sample. All coefficients are based on OLS estimates from Equation (3.9). Error bars represent 95% confidence intervals, based on heteroskedasticity-robust standard errors.

A.9 Debt Service Cost Over Time

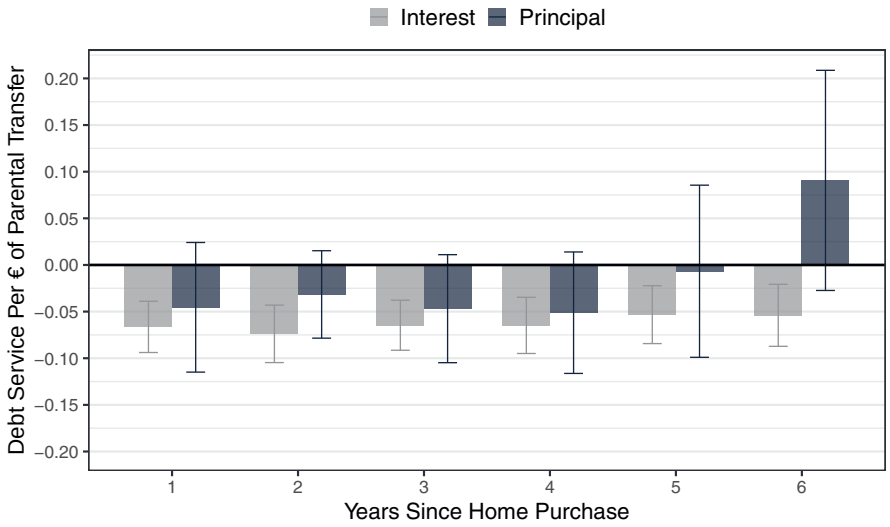


Figure 3.18: Debt Service Cost Over Time: Interest Expenses and Principal Repayment

Notes: This figure shows estimated differences in debt service cost (interest expenses and principal repayment) per euro of parental home purchase transfer, in the first six years after the purchase. While interest expenses are directly observed, principal repayments are imputed by taking the difference in outstanding mortgage debt between year t and $t + 1$. All coefficients are based on 2SLS estimates from Equation (3.11). Error bars represent 95% confidence intervals, based on heteroskedasticity-robust standard errors.

4

Bidder Beware: Intergenerational Wealth Transfers in the Residential Housing Market

This chapter is co-authored with Jaap Bos (Maastricht University) and Nils Kok (Maastricht University).

Abstract

Using an unanticipated tax-exemption policy in the Netherlands, we study the effect of wealth transfers on housing market outcomes. We find that buyers who receive a wealth transfer purchase homes that are 20–35% more expensive, and overpay for a given home by 0.5–2.0%, relative to other buyers. Overpayment is driven by the tightness of the local market in which a home is bought. At the local market level, the policy increases prices due to a spillover effect, whereby homebuyers who did not receive a transfer similarly raise their bids in response to an influx of wealth transfer recipients.

4.1 Introduction

After more than a decade of rapid house price growth in most developed economies, the issue of housing affordability has once again gained attention (Favilukis et al., 2023). In the owner-occupied sector, the situation of young households with limited savings is particularly worrisome. This has led policymakers across jurisdictions to design a variety of schemes supporting first-time homebuyers, ranging from shared-equity mortgages to outright cash grants.[1] Against this backdrop, the role of family wealth in housing markets has come under increased scrutiny. According to a recent report from Redfin, 36% of young homebuyers in the United States rely on family wealth transfers or inheritances to finance their down payment on a home (Redfin, 2024). This trend could exacerbate existing wealth inequalities by limiting access to homeownership for young households lacking familial wealth.

Following increased availability of data on family network links, economists have recently begun to study the role of family wealth in housing transactions. For instance, Benetton et al. (2024) find that parents extract home equity to help their children make a purchase. Brandsaas (2021) shows that parental transfers are important for relaxing borrowing constraints, and account for almost a third of the homeownership rate among young adults. What is missing, thus far, is evidence on how wealth transfers affect housing market outcomes. This study aims to fill the research gap, using granular administrative data from the Netherlands in conjunction with a quasi-experimental policy that facilitated housing-related wealth transfers. We begin by studying how the receipt of a financial transfer affects bidding behavior and transaction outcomes. Using a variety of ways to match buyers who receive a transfer with those who do not, we find that transfer recipients pay higher prices for a given property, holding the list price constant. To address endogeneity concerns, we devise an instrumental variable strategy based on the amount of liquid assets that a buyer's parents hold in their portfolio. Depending on the specification, we estimate that transfer recipients pay 0.5–2.0% higher prices for the same home, compared to other buyers. This effect does not seem to be driven by liquidity constraints of the buyer or by the type of home that is purchased. Instead, we document that transfer-induced overbidding is most pronounced in local markets that are relatively tight to begin with.

[1] For instance, during the 2024 campaign, US presidential candidate Kamala Harris proposed a $25,000 cash grant for first-time homebuyers as part of her economic agenda (Bloomberg, 2024).

Next, we investigate how an influx of wealth transfer recipients affects prices in local housing markets. We rely on variation in exposure to the tax-exemption policy across Dutch postcodes. In a difference-in-differences setting, we compare postcodes in which a larger fraction of buyers received transfers with those areas in which fewer buyers received a wealth transfer. We find that, while largely trending parallel prior to the policy, average bids made by non-recipients increase sharply shortly after the introduction of the policy. This suggests that untreated buyers adapt their bidding behavior in response to competition from wealth transfer recipients. Given that the treatment variable is assigned ex post and subject to endogeneity concerns, we run two additional tests, which support the interpretation of our findings. First, we show that the effect size is increasing monotonically: prices increase most in those postcodes where the share of transfer recipients is the largest. Second, we argue that our findings do not simply reflect the fact that transfer recipients select into gentrifying markets: using tax-assessed values as a placebo outcome, we do not find a significant effect of increased exposure to the policy.

Our work contributes to two strands of the literature. First, it adds directly to the literature on family wealth in housing markets. Charles and Hurst (2003), among the first to study intergenerational wealth patterns, document that the children of homeowning parents are much more likely to be homeowners themselves. Early research from Engelhardt and Mayer (1998) shows that wealth transfers allow households to make purchases earlier in their life cycle, by reducing the required time to save for down payments. More recent empirical studies show that wealth transfers also have a sizable impact on the probability of becoming a homeowner. Using Swiss data, Blickle and Brown (2019) find that a wealth transfer (of any size) increases the transition rate into homeownership (among renters) by 35%. Similarly, Brandsaas (2021) calibrates an overlapping-generations model to data from the United States and finds that wealth transfers account for 31% of the homeownership rate among young adults. Benetton et al. (2024) study the underlying mechanism driving intergenerational transfers and find that they are often preceded by home equity extraction from parents. They show, therefore, that financial transfers are an important driver behind the observed intergenerational persistence in housing wealth.

To our knowledge, no study has yet examined the repercussions that housing-related wealth transfers have on market outcomes. However, a number of closely related articles have identified a positive link between buyer wealth and house prices. For instance, Aiello et al. (2024) find that repeat buyers

who make a gain on selling their previous home tend to overbid when buying the next home. They also find that an inflow of wealthy repeat buyers can drive up prices at the neighborhood level. Cvijanović and Spaenjers (2021), using transaction data from Paris, find that wealthy "out-of-town" buyers make systematically higher bids than others, and that their inflow has a small, but positive effect on neighborhood prices.

Second, our finding of spillovers to untreated homebuyers is in line with a growing literature on spillover effects and social interaction in real estate markets. Burnside et al. (2016) borrow from the epidemiology literature to model social interactions between buyers and the transmission of different beliefs about housing fundamentals. Using social network data from Facebook, Bailey et al. (2018) find a strong effect of peers in residential housing markets: recent experiences of friends and relatives directly influence purchasing decisions, both at the extensive and intensive margin. More closely related to our work, a number of recent papers examine information spillovers across agents within local housing markets. Gupta (2019) shows that foreclosure probabilities increase due to other foreclosures in close proximity. In a similar spirit, McCartney and Shah (2022) establish that refinancing decisions are affected by the activities of close neighbors via the word-of-mouth channel, and Bayer et al. (2021) find that even real estate *investment* decisions are based on the activity of close neighbors. We add another dimension to this literature, by showing that bidding behavior in residential transactions can also be affected by the activity of other buyers in a local market.

The findings in this paper have some important policy implications. At a time when housing prices are high and getting on the "housing ladder" is hard, policymakers are inclined to resort to subsidies, especially for the young. However, such subsidies are at least partially absorbed into higher house prices for *all* market participants. While this benefits existing homeowners, it hurts those who are supported in the first place—first-time homebuyers. In addition, there are distributional effects in providing tax incentives for intergenerational wealth transfers, potentially further increasing inequality in society.

4.2 Institutional Setting and Policy Design

Ownership rates in the Netherlands are around 60%, comparable to the level in the United States. Another commonality with the US is the co-existence

of supply-constrained "superstar cities," such as Amsterdam and Utrecht, and rural areas where recent price growth has been more muted. The most distinctive feature of the Dutch market is a very strong dependence on mortgage finance, which is often explained by generous tax deduction rules for mortgage loans and the general absence of down-payment requirements (see Bernstein and Koudijs, 2024, for a discussion of the Dutch mortgage market).

4.2.1 Stabilizing Policies after the Financial Crisis

As in most other European countries, a period of strong real house price growth in the Dutch market since the 1990s ended abruptly in the wake of the global financial crisis in 2008. However, while house values across the European Union began to recover after 2011, the Netherlands experienced a second downtrend that lasted until 2013. Rouwendal and Petrat (2022) argue that this prolonged slump was due to a tightening of mortgage underwriting rules at the beginning of 2011. Real house prices fell by more than 30% in the period from 2008 to 2013, whereas the annual transaction volume dropped from more than 200,000 transactions in 2008 to just over 100,000 transactions in 2013 (Boelhouwer, 2017).

In the aftermath of the financial crisis, the Dutch government initiated a number of policy changes in the housing market. These were predominantly macroprudential measures to improve stability. For example, eligibility criteria for the national mortgage insurance scheme (NHG) were tightened. The maximum purchase price to qualify for the insurance was gradually lowered from €350,000 in 2011 to €290,000 by 2013, causing the number of new loans covered by the scheme to drop by more than 20%. Another important policy change involved an update to the tax-deduction criteria for mortgage payments. From 2013 onward, only interest payments on fully amortizing mortgages remained tax deductible, which largely eliminated the market for interest-only and savings mortgages (Bernstein and Koudijs, 2024). Finally, in 2012, a statutory limit on loan-to-value (LTV) ratios was introduced. This limit, which was initially set at 106% and subsequently lowered to 100% in 2018, reduced (but did not eliminate) the potential problem of negative home equity in downward markets. At the same time, however, it was another tightening measure that curbed demand and contributed to the prolonged fall of house prices, given that buyers commonly used their mortgages to finance not only the property, but also transaction costs, such as notary fees, transfer taxes, and brokerage fees (Biesenbeek et al., 2022).

4.2.2 The *Jubelton* Policy for Tax-Exempt Transfers

Following the restrictive macroprudential measures, the government eventually also took policy steps to stimulate housing demand. In 2011, the transfer tax was lowered from 6% to 2%, resulting in a significant reduction of effective purchase prices, given that this tax is levied on buyers. While initially planned as a temporary measure to be phased out in 2012, the reduced rates continued to stay in place for almost ten years.[2] Nevertheless, transaction volumes remained low and prices continued to fall after 2011, with the reduction of the transfer tax missing its intended purpose (Solinge and Fiedler, 2020).

Whereas the reduction in transfer tax directly lowered purchase prices, a second policy measure, introduced in 2013, was aimed at stimulating homeownership and reducing indebtedness among young and first-time homebuyers by facilitating *inter vivos* transfers. In particular, this policy offered a generous tax exemption for financial gifts related to housing expenditures. Figure 3.1 in Chapter 3 sketches the timeline of events. Before 2013, the annual tax-exemption threshold for financial gifts was €2,057, with an extended threshold of €5,141 for gifts from parents to children. However, an exemption threshold of €51,407 was in place for one-time transfers from parents to children, conditional upon usage for housing-related expenditures. On September 17, 2013, the government announced a change to the extended threshold along two dimensions. First, it was raised to €100,000, and second, gifts could be made by everyone, regardless of the relationship between donor and recipient.[3] Moreover, it was communicated that the policy would be in place for a limited time period: between October 31, 2013 and December 31, 2014.

The policy had a high take-up rate and received broad media coverage. Overall, more than 150,000 housing-related transfers were made during the time that the policy was in place. This translated into forgone tax revenue of about €1 billion, around 0.5% of the Dutch government's annual tax revenue at the time (Statistics Netherlands, 2019).[4] For our analysis, we identify all people who

[2] In 2021, two changes to the flat 2% tax rate were introduced. First, for young homebuyers, transfer taxes were fully eliminated. Second, for recreational homes, commercial real estate or investment properties, they were increased to 10.4%.

[3] This fiscal policy has become known as the *"Jubelton"*, Dutch for a "celebration of one hundred thousand". Given that gifts are taxed at 10–30%, depending on the relationship between donor and recipient, the policy offered substantial tax relief for not just intergenerational transfers, but also transfers between unrelated individuals.

[4] The total sum of transfers related to the *Jubelton* policy was €9.16 billion in 2013 and 2014 combined.

received a financial transfer under the extended tax-exemption threshold for housing expenditures in 2013 and 2014. This group includes a total of 87,000 gift recipients. Based on our available data, we estimate that around 65% of the recipients used the funds to pay off existing mortgages. In this paper, we focus instead on the smaller group of recipients using the gift to finance the purchase of a home. A third group of beneficiaries invested the money in upkeep or improvement of their existing home, or to acquire ground lease rights. The vast support for the *Jubelton* prompted the government to re-introduce the policy in 2017 in a slightly adapted form, whereby gift recipients had to be between 18 and 40 years old. However, take-up rates for the new scheme, which was ultimately phased out at the end of 2022, never reached the original levels. For instance, in 2017 and 2018, just 24,000 and 10,000 transfers were made, respectively (Statistics Netherlands, 2021).

4.3 Identification Strategy

Since the *Jubelton* policy triggered a large number of financial transfers earmarked for housing expenditures, it offers the ideal laboratory to study the effect of wealth transfers on transaction outcomes. Of course, the group of transfer recipients represents a highly selected sample of the total population of buyers, an issue that we will address through a combination of matching and instrumental variables. In order to analyze market-level effects, we will rely on variation in exposure to the policy across postcodes.

4.3.1 Data Sources

The analysis in this paper draws on data from two different sources. We use a comprehensive dataset on owner-occupied, residential real estate transactions from the largest Dutch real estate broker association (NVM). This dataset covers more than 70% of all sales in the Netherlands and has been used extensively for research purposes (e.g., Aydin et al., 2020; Dröes and Koster, 2016; Leib et al., 2021). For each property that is transacted, it contains a large set of structural characteristics, reported by the listing agent. More importantly for this study, the dataset reports both the price at which a house was initially listed, as well as the final sale price. The ratio between these two, which we refer to as "spread", is the main outcome of interest in our empirical analysis.

We link the transaction data with a number of administrative datasets from Statistics Netherlands (CBS). We begin by identifying, for each transaction from January 2012 until December 2014, the main buyer, based on ownership records that are reported annually for property taxation purposes. Then, we identify each buyer's parents, based on a network dataset that includes all links between family members. Finally, we add buyer demographics, household-level income, and wealth information for buyers and their parents, measured in 2012. The wealth data are particularly detailed, as they report not only total net wealth, but also a breakdown into its different components, including financial assets, bank balances, home equity and mortgage debt. This allows us to compare, through a matching exercise, transfer recipients to other buyers based on their financial situation prior to purchase. Finally, we add information on all financial transfers that a household received in a given year. Importantly, this includes both gross and net gifted amounts, as well as information on the type of transfer. This detailed information allows us to cleanly identify transfers that were made under the *Jubelton* policy. Table 4.1 reports key summary statistics for all buyers, depending on whether they received a tax-exempt transfer or not. While recipients and non-recipients are relatively equal in terms of disposable household income and age at the time of purchase, they differ greatly in terms of wealth. The median net worth among treated buyers (prior to receiving the transfer) is more than three times larger than that of control buyers. This is predominantly driven by differences in liquid assets (bank balances, stocks, and bonds) and lower levels of debt.

Table 4.1: Summary Statistics: Treated vs. Untreated Buyers

	Regular Buyers			Transfer Recipients		
	Mean	Med.	SD	Mean	Med.	SD
Buyer						
Disp. Income (€1,000)	52.14	46.67	34.29	56.42	45.74	57.54
Net Wealth (€1,000)	147.26	26.85	365.94	393.50	98.52	704.39
Total Assets (€1,000)	289.40	199.94	450.14	545.92	260.96	803.02
Liquid Assets (€1,000)	52.29	16.53	135.04	157.26	44.24	335.01
Real Estate (€1,000)	191.16	171.04	240.98	256.91	190.79	351.80
Mortgage Debt (€1,000)	199.94	176.36	127.73	149.87	119.35	150.21
Age (years)	35.63	33.00	10.39	35.39	34.00	10.28
Property						
List Price (€1,000)	237.66	199.50	144.20	291.90	245.00	175.51
Sale Price (€1,000)	224.54	190.00	132.65	275.62	232.50	160.31
Floor Space (m^2)	115.04	110.00	43.54	119.23	115.00	51.86
Number of Buyers		99,183			5,026	

Notes: All income and wealth variables are reported at the household level and measured before *Jubelton* transfers were received. Total assets are the sum of all assets a household owns. Liquid assets comprise bank balances, stocks and bonds. Real estate assets are assessed at their taxation value and include both owner-occupied and other (investment) properties.

4.3.2 Matching Process

Our empirical analysis is based on a comparison of homebuyers who received a tax-exempt financial transfer (*treated buyers*) with those who did not (*untreated buyers*). The challenge to identification in this context arises from the fact that treated buyers are not selected at random from the homebuyer population. To address this concern, we implement two different matching strategies, based on buyer characteristics. In addition, we deal with endogeneity concerns by implementing an instrumental variable strategy based on the share of liquid assets in the buyer's parents' wealth portfolio. Details of the

IV approach are described in the next subsection. We begin here by laying out the matching procedure.

The purpose of matching is twofold. First, it ensures comparability between treated and untreated buyers in observable characteristics, without assuming a functional form for the conditional expectation function. Second, it restricts the analysis to an area of common support. That is, it retains only those treated buyers for which comparable buyers exist in the untreated group, and vice versa. This entails a bias-variance trade-off. On the one hand, more granular matching reduces model dependence, because it does not require extrapolation for observations outside the area of common support. On the other hand, it reduces the sample size and hence the precision of the estimate. What works in our favor is that our dataset includes a large number of households that can potentially serve as control observations, since they did not receive a transfer under the *Jubelton* policy. That means we can be very selective when matching treated households with untreated households, in order to alleviate concerns about selection bias as much as possible.

The starting point for the matching procedure is the full dataset of all 104,209 transactions taking place between November 2013 and December 2014, including 5,026 treated and 99,183 untreated buyers. We start by removing all buyers with recorded net wealth above €5 million or net debt above €1 million, and those with disposable household income above €1 million or below 0. In addition, observations are removed if the buyer's parents are missing in the dataset, either because they are deceased or because they do not live in the Netherlands. This is a first step in improving comparability among buyers, given that the vast majority of transfers were made from parents to their children.

We apply two separate matching strategies to the remaining sample. First, we perform nearest neighbor matching based on propensity scores. To that end, we model the probability of receiving treatment as a linear function of buyer age, net wealth, total assets, disposable income and prior homeownership. Then, a greedy matching algorithm pairs each treated buyer with exactly one untreated buyer that is most comparable. The resulting sample contains 10,052 observations and is perfectly balanced in treatment and control observations.

In addition, we perform coarsened exact matching, as proposed by Iacus et al. (2012). This method is more flexible than exact matching, since it allows the researcher to define strata for each matching variable. In our case, we consider

two observations matched if they are in the same decile of the distribution of net wealth, total assets, disposable income and age. Additionally, they are matched on their previous homeownership status. Conversely, observations are removed if no match exists within the other group. In the estimation, treated and untreated buyers within the same stratum are compared by including matching-cell fixed effects. The resulting matched sample comprises 63,777 observations. Most treated buyers are retained, whereas almost two-thirds of untreated buyers are excluded due to lack of an adequate match.

4.3.3 Empirical Strategy

Based on our matched samples, we estimate the effect of wealth transfers on transaction outcomes via the following linear regression:[5]

$$Y_i = \beta \cdot \textit{Transfer}_i + \phi \cdot X_i + \lambda_j + \epsilon_i \tag{4.1}$$

where $\textit{Transfer}_i$ equals 1 if the buyer in transaction i has received a tax-exempt transfer and 0 otherwise. We study two separate outcomes, denoted by placeholder variable Y_i. First, the natural logarithm of property i's list price. In this model, β reveals to what extent buyers use a financial transfer to select more valuable homes. Second, the natural logarithm of the ratio between sale price and list price, which we denote as spread. Here, β captures the degree to which transfer recipients increase their bid for a given home. Given that price spreads might differ across price segments, we control for a home's list price in this model. Finally, λ_j captures matching-cell fixed effects in the CEM sample.

The spread between sale price and list price serves as an intuitive proxy for relative over- versus underpayment. It is based on the assumption that, conditional upon a property's unobserved true value, list prices are independent of the buyer's treatment status. Given that list prices are determined before potential buyers arrive, this assumption seems plausible.

[5]Since we estimate Equation (4.1) as a pooled cross-section, we omit time subscripts.

4.3.4 Addressing Endogeneity in Financial Transfers

By virtue of matching, we only compare buyers who are similar in a number of observable characteristics. However, matching cannot help overcome endogeneity concerns due to unobserved confounders. The estimated treatment effect from Equation (4.1) could therefore be subject to bias. If, for instance, transfer recipients have more housing market experience or hire better real estate agents, which in turn helps them to better negotiate the price of a property, the estimated least-squares coefficient would underestimate the true causal effect of the transfer. Assuming that donors (typically parents) who make housing-related wealth transfers are familiar with the housing market, it is plausible that transfer recipients are more sophisticated than the average buyer in our sample. Addressing this concern requires exogenous variation in financial transfers that is unrelated to any confounding factor. To isolate such variation, we rely on the high quality of Dutch administrative data and propose an instrument based on the financial situation of buyers' parents.

The policy we study was announced surprisingly and only in place for a relatively short time. Therefore, we argue that it is not parental wealth per se, but their ownership of liquid assets, that drives the supply of wealth transfers. Two potential donor families with equal net wealth could have had a very different propensity to make a transfer, depending on the type of assets that they owned as the policy was announced. Parents who owned a sizable amount of liquid assets were well-positioned to make a transfer, whereas those who derived their wealth from illiquid assets, such as real estate, may have been unable to do so. We therefore use the amount of liquid assets (bank deposits plus stocks and bonds) that a buyer's parents own as an instrument. Given that we also control for the parents' net wealth, this instrument captures the share of liquid assets in the parental wealth portfolio.

In order for our first instrument to qualify, it needs to meet two conditions. The first is instrument relevance. It requires the existence of a strong first stage, with nonzero correlation between the instrument and the endogenous treatment variable. We provide a test of this condition in the next section. The second condition, the exclusion restriction, cannot be formally tested. It requires that the instrument is uncorrelated with the error term in the reduced form regression. Intuitively, conditional on treatment assignment and control variables, a buyer's bid for a house must not be related to the share of liquid assets that her parents hold.

There are a number of reasons why this assumption could be violated. On the one hand, parental liquid wealth may be related to a buyer's socioeconomic status, which in turn affects the bidding behavior. However, we argue that differences in unobserved characteristics would rather be related to a family's total net wealth, and not to the composition of their asset portfolio. A second threat to the exclusion restriction would arise if parents used their liquid assets to support a home purchase in ways other than by making a wealth transfer. For instance, mortgage underwriters could offer more favorable terms to applicants whose parents can give financial support or even act as a guarantor. However, while we cannot rule this out theoretically, conversations with mortgage practitioners have suggested that it is very uncommon in practice. Finally, parents with a low share of liquid assets typically have much of their wealth tied up in real estate. While this restricts them in their capacity to make financial transfers, it could imply a higher level of sophistication in dealing with housing transactions, due to more experience. To address this, we include parental homeownership as a control variable in the IV estimation.

We implement the IV strategy by estimating the following system of equations:[6]

$$Y_i = \beta \cdot \widehat{Transfer}_i + \phi \cdot X_i + \lambda_j + \epsilon_i \tag{4.2}$$

$$Transfer_i = \boldsymbol{\pi} \cdot \mathbf{Z}_i + \phi \cdot X_i + \lambda_j + \eta_i \tag{4.3}$$

Equation (4.2) is the second stage in the 2SLS system and Equation (4.3) is the first stage. In the first stage, we estimate a linear probability model to predict the receipt of a wealth transfer based on our instrument (Z_i): the natural logarithm of liquid assets that a buyer's parents owned in 2012. To account for potential non-linearity in this relationship, we also include the square of this instrument. Additional control variables include: parental net wealth, financial assets and homeownership status. The predicted transfers from Equation (4.3) are then used to study transaction outcomes in Equation (4.2). We apply this model to the CEM sample and control for matching-cell fixed effects, in order to maintain comparability in observable buyer characteristics.

[6] As was the case for Equation (4.1), we omit time subscripts.

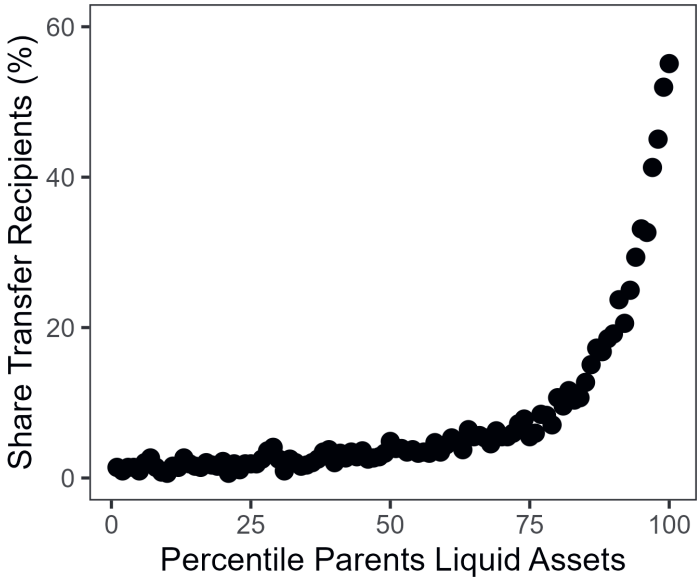


Figure 4.1: Parental Liquid Assets and Wealth Transfers

Notes: This figure plots the share of transfer recipients for homebuyers whose parents are in different percentiles of the liquid asset distribution.

4.4 Empirical Results

4.4.1 First Stage Results

We begin by providing graphical evidence for the instrument's relevance. Figure 4.1 displays a binned scatterplot, in which the share of transfer recipients is depicted for each percentile in the distribution of parental liquid assets. Two observations stand out: first, transfers increase largely monotonically as a function of parental liquid assets. Second, the relationship is strongly non-linear: while 4.9% of children of parents at the median of the parental liquid assets distribution (€37,000) receive a tax-free transfer, 19.2% of those with parents at the 90th percentile (€282,000) do.

First stage results, reported in Table 4.5, confirm the graphical impression. The instruments, log(Liquid Assets Parents) and log(Liquid Assets Parents)^2, are highly significant after controlling for other financial characteristics of buy-

ers and their parents. The Kleibergen-Paap F-Statistic in the full model, including a home's list price, is equal to 738.3, far exceeding any conventional rule-of-thumb threshold.

4.4.2 Financial Transfers and Transaction Outcomes

The outcomes of the baseline analysis are provided in Table 4.2. Column (1) reports unmatched OLS results, while nearest neighbor and CEM results are reported in columns (2) and (3), respectively. Finally, column (4) contains the second-stage results of our IV estimation, applied to the CEM sample. As discussed above, OLS and matching results could be downward-biased if treated buyers are more sophisticated in real estate negotiations. However, our first-stage visualization in Figure 4.1 suggests that most transfers are made by parents far in the right tail of the distribution of liquid assets. If children of these wealthy parents are more prone to using their financial windfall for overbidding and for buying pricier homes, the local average treatment effect in our IV estimation would be larger than the average treatment effect. We therefore believe it is reasonable to view matching and IV estimates as lower and upper bounds for the true causal effect of a financial transfer.

The estimated effect of wealth transfers on house selection is reported in panel A. Matching results suggest that transfer recipients buy homes that are approximately 20% more expensive than others. To be precise, the nearest-neighbor coefficient of 0.205 translates into an increase of 22.75%, or €47,500 above the average-priced home. The IV estimates suggest an even larger effect size, of roughly 35%. Taken together, these results indicate that financial transfers have substantial intensive-margin effects. Of the €100,000 that the median treated buyer was given in our sample, at least half was spent on purchasing more expensive homes (reflecting, for example, higher quality, larger living area etc.).

The analysis of bidding outcomes is reported in panel B. Here, the estimated effect of a transfer ranges from 0.005, for the OLS and CEM analysis, to 0.020 for the IV model. The results imply that for a home of given quality, transfer recipients pay a premium of approximately 0.5% to 2.0%, relative to non-treated buyers. Compared to the average transfer size and the list price effect reported above, this is relatively modest. However, it is economically non-trivial and especially noteworthy in view of the overall market condition at the time. Transfer recipients were willing to raise their bids in an economic

environment where most homes sold in bilateral negotiations and at a significant discount relative to the list price. Below, we will discuss the motive of wealth-transfer recipients for doing so and the ripple effects that this behavior might have on other buyers.

Table 4.2: Effect of the Wealth Transfer on Bidding Outcomes

	(1) OLS	(2) PSM	(3) CEM	(4) IV
Panel A: Housing Consumption	*Dep. Variable: log(List Price)*			
Transfer	0.200^{***}	0.205^{***}	0.163^{***}	0.300^{***}
	(0.007)	(0.010)	(0.006)	(0.032)
Buyer Controls	Yes	No	No	No
Parents Controls	No	No	No	Yes
Mean Outcome Control	12.26	12.25	12.28	12.28
Observations	104,209	10,052	63,777	63,777
Adjusted R^2	0.152	0.041	0.278	0.282
Panel B: Bidding Outcomes	*Dep. Variable: log(Spread)*			
Transfer	0.005^{***}	0.006^{***}	0.005^{***}	0.020^{***}
	(0.001)	(0.001)	(0.001)	(0.004)
log(List Price)	-0.016^{***}	-0.018^{***}	-0.016^{***}	-0.017^{***}
	(0.001)	(0.001)	(0.001)	(0.001)
Buyer Controls	Yes	No	No	No
Parents Controls	No	No	No	Yes
Mean Outcome Control	−0.053	−0.055	−0.054	−0.054
Observations	104,209	10,052	63,777	63,777
Adjusted R^2	0.027	0.025	0.035	0.030

Notes: This table reports estimates of the effect of wealth transfers on bidding outcomes. The treatment variable is a dummy, indicating whether a buyer has received a tax-exempt wealth transfer around the time of the purchase. Models in columns (1)–(3) either control for observable buyer characteristics or match buyers based on these characteristics. Column (4) employs an instrumental variables model based on the amount of liquid assets that a buyer's parents own. Heteroskedasticity-robust standard errors are reported in parentheses, and stars indicate 10% (*), 5% (**) and 1% (***) significance levels.

4.4.3 Robustness Tests

For the remainder of the empirical analysis, we focus on our second outcome, overpayment relative to untreated buyers. We first subject the baseline findings to a series of additional robustness tests. These are constructed to address different empirical threats to our design, related to the measure of overpayment, the use of a binary treatment dummy, and the specifics of the policy itself. The results of these tests, each of which we conduct both on the unmatched and on the CEM sample, are reported in Table 4.3.

First, we relax the assumption that the value of all structural and location characteristics is subsumed in the list price, such that buyers do not overbid systematically for homes of a certain type. This considers the possibility that transfer recipients could have preferences for certain amenities or locations, and that these carry a premium over the list price.[7] Columns (1) and (2) show that, while point estimates decrease slightly to 0.4% in the OLS and CEM model, the baseline effect appears to be robust to the inclusion of these additional hedonic characteristics.

Next, we address the concern that overbidding by transfer recipients could be driven by the particular design of the policy we study. One feature of the *Jubelton* policy was that after receiving a transfer, recipients were obliged to use the money during the policy period, in order not to forgo the tax exemption. It is possible, therefore, that transfer recipients only started overpaying towards the end of the policy period, due to time pressure. This would threaten the external validity of our findings, since wealth transfers typically do not involve a time constraint. However, these concerns appear to be unjustified, as an estimation of the baseline model excluding the last three months of the policy window shows. The results, reported in columns (3) and (4), are in line with the baseline model, suggesting that the tendency of transfer recipients to overbid is not tied to the particular policy under study, but inherently linked to the financial windfall they receive.

Finally, we replace the binary treatment dummy with a series of bins, indicating the size of the (net) transfer a buyer has received. If recipients are actually leveraging their extended budget to overbid, the estimated premium should increase with the size of the transfer. The results, reported in columns (5) and (6), confirm that this is indeed the case. Interestingly, for buyers

[7] Specifically, we control for the following variables: apartment dummy, construction period, living area, building type, number of rooms, interior and exterior condition.

with transfers below €50,000, we observe no significant effect. This suggests that bidding behavior is not the initial margin of adjustment for homebuyers as their budget grows. Beyond €50,000, the estimated effect is positive and increasing in the size of the transfer. Based on the OLS and CEM results, we find that buyers with transfers larger than €100,000—those exceeding the exemption limit—pay a premium of about 1%, relative to comparable buyers without any transfer.

Table 4.3: Robustness Analysis

	(1) OLS	(2) CEM	(3) OLS	(4) CEM	(5) OLS	(6) CEM
	Dep. Variable: log(Spread)					
Transfer	0.004***	0.004	0.005***	0.005***		
	(0.001)	(0.004)	(0.001)	(0.001)		
log(List Price)	−0.007***	−0.007***	−0.017***	−0.017***	−0.016***	−0.016***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Transfer ≤ 50k					0.000	−0.002
					(0.002)	(0.002)
50k < Transfer < 100k					0.004***	0.005***
					(0.001)	(0.001)
Transfer = 100k					0.006***	0.006***
					(0.001)	(0.001)
Transfer > 100k					0.011***	0.009***
					(0.003)	(0.003)
Matching Cell FE	No	Yes	No	Yes	No	Yes
Hedonic Controls	Yes	Yes	No	No	No	No
Sample Period	Full	Full	Short	Short	Full	Full
Observations	104,209	63,777	76,728	47,122	104,209	63,777
Adjusted R^2	0.081	0.084	0.030	0.037	0.028	0.035

Notes: This table reports additional robustness tests of the baseline analysis of bidding outcomes. In columns (1) and (2), hedonic characteristics are added to the model. In columns (3) and (4), the final months of the policy are excluded from the sample. In columns (5) and (6), the treatment variable is grouped into a series of bins. Heteroskedasticity-robust standard errors are reported in parentheses (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

4.4.4 Underlying Mechanisms

The baseline analysis shows that buyers increase their bid for a home after receiving a wealth transfer. This result can be driven by different underlying mechanisms, including characteristics of the buyer, the purchased property or the local market in which this property is located. In this section, we include a series of interaction terms in the baseline model, to assess the merit of these different explanations.

The Dutch mortgage market is characterized by quite generous limits on loan-to-value ratios. In 2013, the maximum mortgage size was equal to 105% of the assessed house value. This allowed buyers to finance not just the full purchase, but additional cost as well, including transaction taxes or broker fees. Nonetheless, home purchases entail a range of additional expenses, including those on home-related durables and improvements (Benmelech et al., 2023). Therefore, liquidity-constrained buyers may have an incentive to negotiate aggressively in order to secure a low price. We hypothesize that, among previously liquidity-constrained buyers, the receipt of a transfer will have a more pronounced effect on price spreads. We test this in columns (1) and (2) of Table 4.4, where we interact the treatment with an indicator that takes the value 1 for buyers with below-median liquid assets. However, the results show that the relevance of this channel seems to be limited, at best. While the main coefficient remains largely unchanged, at 0.4%, we do not find a significant difference between transfer recipients with large and small holdings of liquid assets.

Alternatively, financial transfers may help buyers find homes with a higher match value. Given that a larger budget allows buyers to broaden their search into higher price segments, it increases their chances of finding a "dream home", for which they are willing to pay a premium. We proxy for this effect by distinguishing between non-detached and detached units. Non-detached units—including apartments, row houses and corner houses—are relatively standardized in most Dutch markets. Therefore, we argue that they are less dispersed with respect to the idiosyncratic match value they offer to buyers. In contrast, a buyer's chance of finding a unique property that is a perfect match is larger in the market for detached homes, which offers much greater variety. In order to test this mechanism, we interact the treatment with a dummy for detached units. Additionally, we include postcode-level fixed effects in this model, to control for the fact that detached and non-detached units tend to be located in different areas. The results, reported in columns (3) and (4), do not

support the hypothesis that overpayment is driven by higher match values. In contrast, the estimated interaction effect—Transfer × Detached Home—is actually negative, albeit only significant at the 10% level. We therefore conclude that transfer recipients do not overpay to optimize idiosyncratic match value, as measured by the bidding behavior on detached homes.

Finally, bidding behavior could be driven by the local market environment. In competitive markets, buyers with the necessary means might raise their bid in order to reduce the risk of being outbid and to avoid a lengthy search process. To test this, we group postcodes into "tight" and "relaxed" markets, based on whether the average time on market was below or above the median (weighted by the average transaction volume) before the introduction of the policy. The results of this model, reported in columns (5) and (6), suggest that the local market environment does indeed determine the bidding behavior of transfer recipients. In fact, we find that transfer recipients only overpay, relative to other buyers, in relatively tight housing markets. While the uninteracted treatment does not have a significant effect in these models, the estimated coefficient on the interaction terms is 0.7% and 0.4% in the OLS and CEM models, respectively. In the context of our analysis, overpayment therefore seems to be driven by the extent of competition in a local market. Recalling that our analysis is set against the backdrop of a sluggish housing market, we expect this effect to be even stronger in the context of a heated housing market.

4.5 Spillovers to Local Housing Markets

In the previous section, we established that recipients of wealth transfers make higher bids for a given home and that this effect is driven by the degree of competition in a market. Motivated by the growing evidence on spillovers in residential housing markets (Bayer et al., 2021), we now examine whether home purchases by wealth transfer recipients affect the bidding behavior of other buyers in a local market. Within a local market, potential buyers are often competing for the same units, and information about neighboring transactions is communicated indirectly via real estate agent networks. Therefore, it is plausible that during their search, buyers are informed about transaction outcomes in a local market. To the extent that buyers take this into account during their bidding decision, the influx of wealth transfer recipients into a local market could lead other buyers to increase their bids.

Table 4.4: Underlying Mechanisms

	Buyer		**Building**		**Market**	
	(1) OLS	(2) CEM	(3) OLS	(4) CEM	(5) OLS	(6) CEM
	Dep. Variable: log(Spread)					
Transfer	0.004***	0.004***	0.003***	0.003***	0.000	0.002
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Liquidity Const.	−0.001***	−0.002**				
	(0.000)	(0.001)				
× Transfer	0.002	0.001				
	(0.002)	(0.002)				
Detached Home			−0.019***	−0.018*		
			(0.001)	(0.001)		
× Transfer			−0.004*	−0.005*		
			(0.003)	(0.003)		
Tight Market					0.012***	0.013***
					(0.000)	(0.000)
× Transfer					0.007***	0.004**
					(0.002)	(0.002)
log(List Price)	−0.016***	−0.016***	−0.015***	−0.014***	−0.015***	−0.015***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Matching Cell FE	No	Yes	No	Yes	No	Yes
Postcode FE	No	No	Yes	Yes	No	No
Observations	104,209	63,777	104,209	63,777	103,969	63,618
Adjusted R^2	0.028	0.035	0.131	0.145	0.040	0.048

Notes: This table reports results from three heterogeneity analyses aimed at studying the mechanism driving the bidding response. In columns (1) and (2), the treatment is interacted with a dummy indicating liquidity-constrained buyers. In columns (3) and (4), it is interacted with a dummy for detached homes. In columns (5) and (6), it is interacted with a dummy for tight postcodes. Heteroskedasticity-robust standard errors are reported in parentheses (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

To investigate this, we rely on the fact that there is substantial regional variation in where transfer recipients buy homes. We define exposure as the share of transfer recipients among homebuyers in a postcode during the policy window, ranging from 1.54% at the 25th percentile to 6.93% at the 75th percentile, with

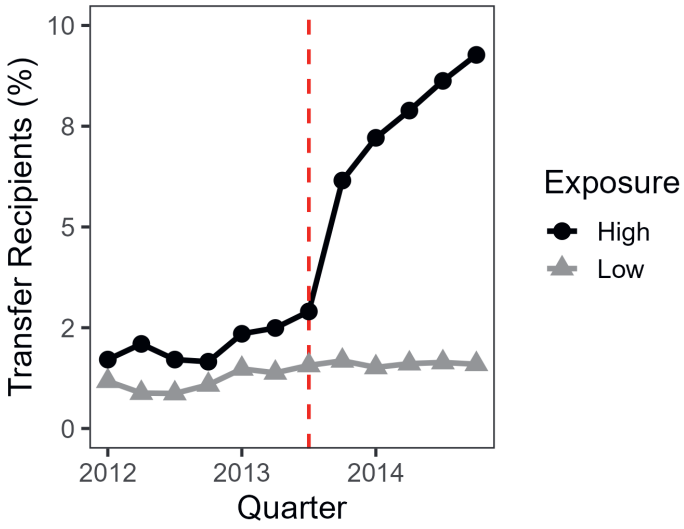


Figure 4.2: Share of Transfer Recipients over Time

Notes: This figure plots, on a quarterly level, the share of transfer recipients among all buyers in postcodes classified as high exposure versus low exposure. The dashed red line indicates the announcement of the policy.

a median of 3.98%. We base the empirical analysis on postcodes, not entire municipalities, in order to control more precisely for unobserved confounders that may correlate with treatment exposure. Dutch postcodes are somewhat smaller than their US equivalent, with a mean population size of 4,000.

In Figure 4.2, we split postcodes into those with below and above median exposure, and plot the average share of transfer recipients among all homebuyers in a given quarter. Even before the introduction of the policy, some purchases are made by buyers who received a tax-exempt transfer. As the introduction of the *Jubelton* policy raises the exemption threshold to €100,000, the share of transfer recipients rises quickly in high-exposure postcodes, from 2.9% in 2013Q3 to 7.2% in 2014Q1. Interestingly, transfers become increasingly popular over time, and they peak just before the policy is phased out in 2014Q4. In contrast, in below-median exposed postcodes, the share of transfer recipients stays virtually flat after the introduction of the policy.

4.5.1 Price Spreads in Exposed Markets

We examine whether competition from wealthy transfer recipients drives up average bids in a local market. To this end, we compare postcodes with varying exposure to the tax-exemption policy. The main challenge to identification stems from the fact that regional exposure to the treatment, much like the individual-level treatment, is not randomly assigned. To address this issue, we include a pre-period into the analysis and control for postcode- and quarter-fixed effects. Any remaining bias would therefore have to come from time-varying factors that correlate both with postcode-level bidding patterns and the location choices of transfer recipients. If, for instance, transfer recipients would systematically select into postcodes that "gentrify" around the time of the policy introduction, our estimates include the effect of other, omitted, factors. Given that we lack (quasi-) random variation in the allocation of treated buyers across space, we cannot rule out these concerns entirely. However, we will provide additional evidence below to support the interpretation that changes in local prices are indeed driven by the influx of transfer recipients.

Empirically, our approach resembles a dynamic difference-in-differences model with a continuous treatment variable. The following model is estimated, based on the extended sample, covering all transactions from January 2012 until December 2014:

$$Y_{ipt} = \sum_{k} \beta_k \cdot 1\{t = k\} \cdot Exposure_p + \mu_p + \delta_t + \varepsilon_{ipt} \tag{4.4}$$

As in the individual level analysis, Y_{ipt} captures the log sale-to-list ratio for transaction i. The additional subscripts t and p indicate the quarter in which a transaction took place and the postcode in which the home is located. $Exposure_p$ denotes the share of transfer recipients among all homebuyers in postcode p during the policy window. While the baseline differences across postcodes are captured by the unit fixed effects, μ_p, the variables of interest are the interaction terms between postcode exposure and the quarter-dummies, $\sum_k 1\{t = k\}$, ranging from 2012Q1 until 2014Q4. We omit 2013Q2, the quarter just before the announcement of the *Jubelton* policy. Therefore, the coefficient vector β_k captures differences in price spreads across postcodes with varying exposure, relative to the baseline differences in 2013Q2. Finally, quarter fixed effects δ_t absorb aggregate time trends.

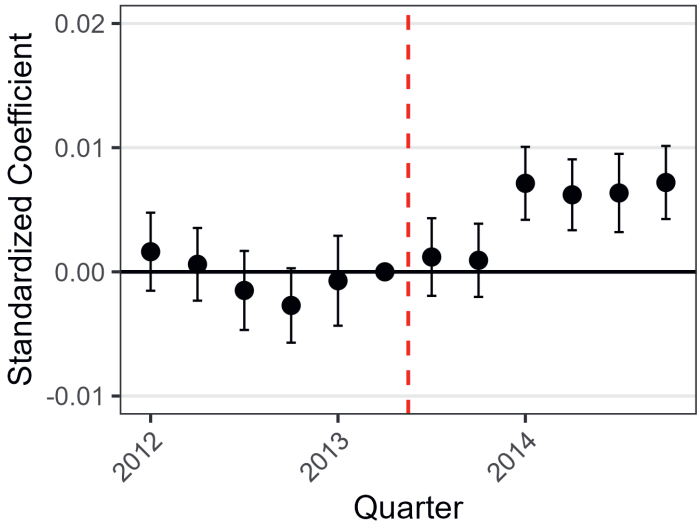


Figure 4.3: Price Spreads in Exposed Markets

Notes: This figure visualizes the coefficients from a dynamic difference-in-differences model, studying the effect of a higher share of transfer recipients among buyers on price spreads. Error bars represent the 95% confidence interval. The dashed red line indicates the announcement of the policy.

We plot the standardized model coefficients in Figure 4.3. Prior to the introduction of the policy, when hardly any housing-related wealth transfers were made, average price spreads evolved similarly across postcodes. Importantly, we do not find any evidence for significant pre-trends. After the introduction of the policy (at the end of 2013Q3), price trends stay the same for an additional quarter, before increasing significantly in those postcodes where many transfers are made. This is consistent with the actual share of transfer recipients in high-exposure postcodes documented in Figure 4.2, rising after the policy introduction, but only peaking over the course of the following year. Throughout 2014, we estimate that in postcodes with a one standard deviation increase in exposure, average sale prices rose by approximately 0.7%, relative to the list price. Importantly, this effect is not driven by transfer recipients alone. In Figure 4.6, we repeat the same analysis, but omit all transfer recipients from the estimation. The results are largely unchanged, suggesting that regular buyers updated their bidding behavior in response to an influx of transfer recipients.

4.5.2 Robustness Tests

Even though we can control for all time-invariant differences across postcodes, our findings need not reflect a true causal effect of competition from transfer recipients on the bidding behavior of other buyers. Particularly, the fact that postcode-level exposure is defined ex post opens the door for alternative explanations. For instance, wealth transfer recipients could self-select into gentrifying neighborhoods where prices would have risen anyway. Alternatively, our findings could be driven by reverse causality if, due to rising prices, some postcodes become relatively less attractive to prospective buyers who did not receive financial support. Fully overcoming these concerns would require (quasi-) random variation in the allocation of transfer recipients across local housing markets, which is not available in our setting. Instead, we implement two additional tests, in order to rule out alternative explanations.

We begin by estimating an alternative model where, instead of using the continuous exposure variable from Equation (4.4), we group postcodes into exposure deciles, weighted by the average transaction volume in the pre-policy period. Then, we reestimate the baseline model in its static form, comparing outcomes before and after the introduction of the policy. Postcodes in the bottom two exposure deciles, which did not attract any transfer recipients, are the omitted group. The intuition behind this test is straightforward: if price increases truly are driven by competition from transfer recipients, we expect the effect size to grow monotonically as a function of postcode exposure. Moreover, the relationship could be non-linear, if transfer recipients are only "noticed" in a local market once they are sufficiently numerous. The results, depicted in Figure 4.4, are in line with this interpretation. The figure plots coefficients from the interaction terms between the post-policy dummy and each exposure decile. In addition, it shows, on the secondary y-axis, the average share of transfer recipients in each decile. Up until the sixth exposure decile, there is no measurable price effect, relative to postcodes which attracted no transfer recipient at all. Beyond that, however, the effect size is growing quickly. Going from the 9th to the 10th exposure decile, the estimated coefficient almost doubles, from 0.55% to 1%.

Next, we address the concern that transfer recipients are drawn into gentrifying neighborhoods, where prices would have also risen in the absence of the policy. To investigate this, we implement a placebo test, relying on another piece of information about properties that is available in the data from Statistics Netherlands. In the Dutch market, properties are assessed for taxation

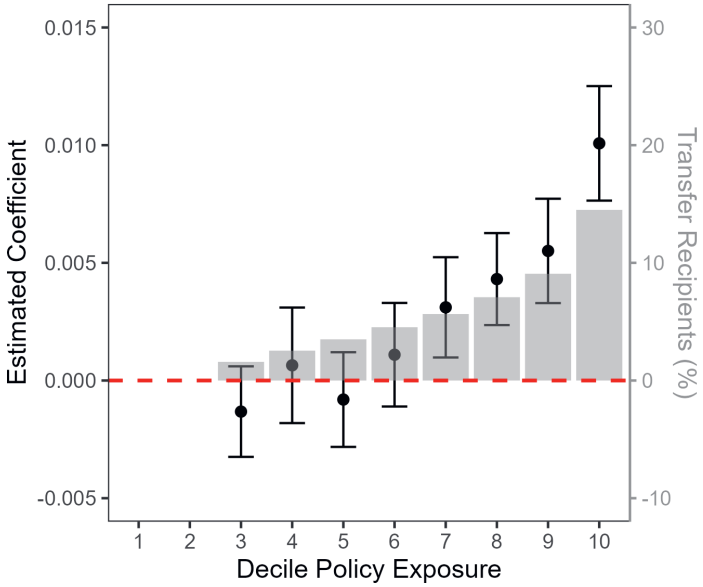


Figure 4.4: Price Spreads in Exposure Deciles

Notes: This figure visualizes the coefficients from a static difference-in-differences model, studying the effect of a higher share of transfer recipients among buyers on price spreads. Here, the continuous treatment variable is replaced with dummies indicating the respective decile of treatment exposure. Error bars represent the 95% confidence interval.

purposes on an annual basis, and all tax-assessed values are public information. Municipalities are responsible for the assessment, which is supposed to provide an objective measure of a property's current market value. The intuition underlying our placebo test is the following: in case the share of transfer recipients in a postcode is indeed unrelated to confounding factors, it should not predict changes in tax-assessed values over time. If, instead, assessed values would be on a differential trend in exposed and unexposed postcodes, this would rule out spillovers from transfer recipients as the (sole) driver behind our findings.

We implement this placebo test by replacing, in Equation (4.4), the dependent variable with the natural logarithm of the current assessed value of property i. The resulting coefficients are plotted in Figure 4.5. While estimates are noisier than those from the baseline model, we do not find any evidence for a rise in

tax-assessed values in postcodes that are more exposed to transfer recipients. Taken together, our findings therefore suggest that in exposed markets, average bids increase relative to list prices, and that this increase cannot be explained by rising property values in these markets.

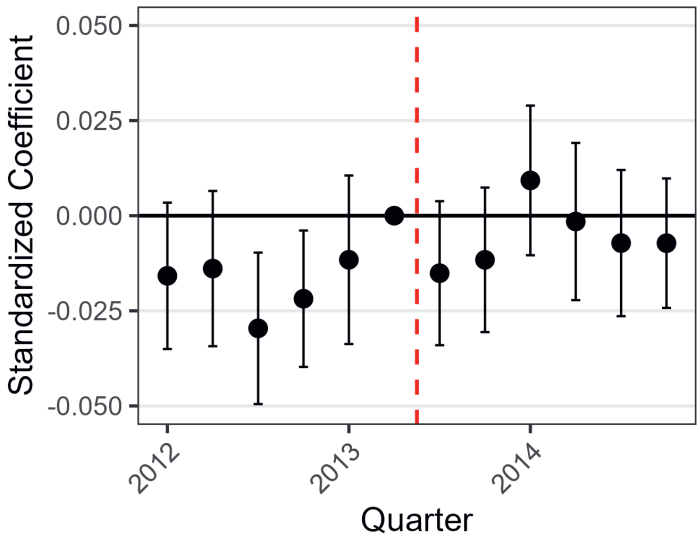


Figure 4.5: Tax-Assessed Values in Exposed Markets

Notes: This figure visualizes the coefficients from a dynamic difference-in-differences model, studying the effect of a higher share of transfer recipients among buyers on average tax-assessed values. Error bars represent the 95% confidence interval. The dashed red line indicates the announcement of the policy.

4.6 Conclusion

A substantial number of homebuyers are relying on wealth transfers as a means of financing their down payment (Brandsaas, 2021). In this paper, we take a first step towards understanding the effect of wealth transfers on housing market outcomes. To do so, we rely on an exogenous tax-exemption policy for wealth transfers in the Dutch housing market. For a limited period of time, transfers up to €100,000 were exempt from taxation, as long as the funds were used for housing-related expenditures. We are able to identify homebuyers who received a transfer and compare them to others, buying a house at the

same time but without financial support. To overcome endogeneity concerns, we match treated and untreated buyers and instrument for the receipt of a wealth transfer using the share of liquid assets in a buyer's parents' wealth portfolio.

At the individual level, we document two main findings. First, we estimate that at least half of the total net transfer a buyer receives is spent on buying more expensive homes. Across the different models, we estimate that transfer recipients choose homes which are 20–35% more expensive than other buyers, all else equal. Since there are no direct positive externalities associated with larger homes, we conclude that exempting housing-related transfers from taxation is mainly a subsidy for those who are relatively well-off to begin with. Second, we find that transfer recipients pay a premium of 0.5–2% for a given home, relative to other buyers. This tendency to overpay is driven by the tightness of the local market in which the purchase takes place, suggesting that transfer recipients use their increased budget to avoid a lengthy search process. Therefore, we hypothesize that transfer recipients could be overpaying more aggressively as housing markets "heat up" and become more competitive.

In addition, we establish a local spillover effect in bidding behavior: as transfer recipients move into a local market, other buyers tend to increase their bids as well. To show this effect, we calculate each postcode's exposure to the policy, based on the share of transfer recipients among all buyers during the policy period. We compare local markets that vary in their exposure to the policy before and after it has been introduced, controlling for postcode- and quarter-fixed effects. We find that, for a one standard deviation increase in policy exposure at the postcode level, sale prices increase by 0.7%, relative to list prices. This effect increases in a non-linear way with exposure, suggesting that spillovers require a "critical mass" of transfer recipients among homebuyers.

Taken together, our findings suggest that intra-family wealth transfers could have important, and potentially distortionary, effects on housing markets. As housing affordability continues to decline, the share of first-time buyers relying on family resources will likely grow even further. More research needs to be done in order to understand the implications of this along a number of dimensions, including wealth inequality, housing market outcomes and, ultimately, optimal policy design.

A Appendix

A.1 IV First Stage Results

Table 4.5: IV First Stage: Predicting Wealth Transfers

	Dep. Variable: Transfer		
	(1)	(2)	(3)
log(Liquid Assets Parents)	0.022***	−0.163***	−0.160***
	(0.001)	(0.005)	(0.005)
log(Liquid Assets Parents)^2		0.010***	0.009***
		(0.000)	(0.000)
log(List Price)			0.034***
			(0.003)
log(Total Assets Parents)	−0.001	0.020***	0.019***
	(0.001)	(0.002)	(0.001)
Net Wealth Perc. Parents	0.002***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)
Homeowners Parents	−0.055***	0.072***	0.070***
	(0.005)	(0.005)	(0.005)
Matching Cell FE	Yes	Yes	Yes
Kleibergen-Paap F-statistic	587.0	747.3	738.3
Observations	63,777	63,777	63,777
Adjusted R^2	0.123	0.165	0.167

Notes: This table reports estimates of the first stage regression. The endogenous variable is a dummy indicating whether the buyer has received a tax-exempt transfer. The instrument is the natural log of parental liquid assets (linear and squared). Additional exogenous variables include the log of parental (total) assets, parental net wealth (percentile rank) and a dummy indicating parental homeownership. All models include matching-cell fixed effects, to compare buyers within the same stratum for wealth, income, age and parental wealth. Heteroskedasticity-robust standard errors are reported in parentheses (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

A.2 Regional Diff-in-Diff Excluding Transfer Recipients

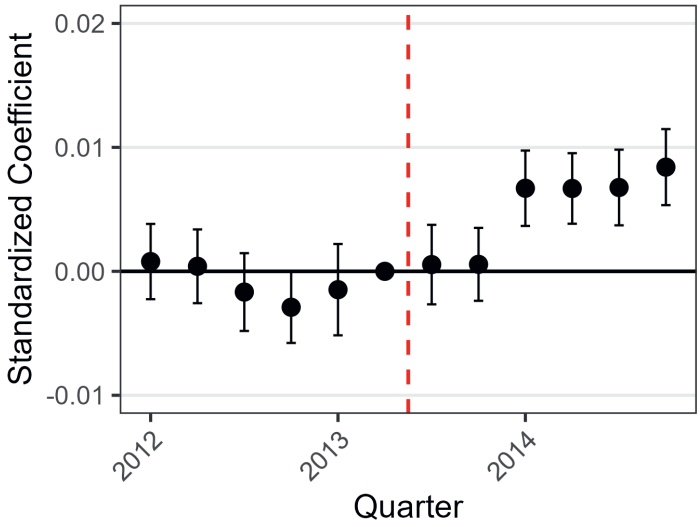


Figure 4.6: Price Spreads in Exposed Markets: Excluding Transfer Recipients

Notes: This figure repeats the baseline difference-in-differences analysis from Figure 4.3, excluding transfer recipients. Error bars represent the 95% confidence interval. The dashed red line indicates the announcement of the policy.

5

Ethnic Matching in the Housing Market

This chapter is co-authored with Jaap Bos (Maastricht University) and Olivier Marie (Erasmus University Rotterdam).

Abstract

Using Dutch administrative data and the near universe of residential property transactions, we document that homebuyers and homesellers with the same migration background match with each other at rates that substantially exceed random search predictions. Transactions between same-ethnicity buyers and sellers sell at a price discount of around 4%, relative to otherwise comparable transactions. This discount is highly robust and holds across most migrant groups we examine. While around half of the discount can be explained by non-arm's-length transactions, the remainder likely reflects ethnic in-group favoritism. Our findings therefore reveal a novel mechanism through which social preferences affect housing market outcomes.

5.1 Introduction

Residential segregation has long-run consequences for wealth accumulation and economic opportunity. Discrimination by sellers, brokers, and lenders makes it difficult for ethnic minorities to access desirable neighborhoods (Ihlanfeldt and Mayock, 2009; Christensen and Timmins, 2022; Ambrose et al., 2021), while preferences for co-ethnic proximity and shared amenities may further concentrate minority households in ethnic neighborhoods (Aliprantis et al., 2024; Waldfogel, 2008). Despite extensive evidence on both supply-side discrimination and demand-side sorting, another mechanism has received little attention: the preference of buyers and sellers of the same ethnic background to transact with each other, independently of neighborhood composition. Whether such preferences lead to price premiums or discounts is theoretically ambiguous. Sellers may earn a premium from co-ethnic buyers if shared preferences for neighborhood or property amenities raise willingness to pay; alternatively, they may offer a discount out of in-group favoritism or trust. If such preferences generate systematic price differentials, they constitute a novel channel through which ethnic identity shapes housing market outcomes. Unlike discrimination or neighborhood sorting, this mechanism operates within neighborhoods and has direct implications for how housing wealth accumulates among migrant households.

In this paper, we study ethnic matching in the Dutch housing market, a setting that is uniquely suited to this question. The Netherlands is home to one of the most ethnically diverse migrant populations in Europe, with precise administrative records of origin country for each individual. We combine these records with the near universe of residential property transactions, allowing us to identify the exact migration background of every buyer and seller across more than a decade of transactions. Crucially, the administrative data extend well beyond demographics: they cover individual income, detailed household balance sheets, and family network links across the entire population, allowing us to control for financial differences between buyers and sellers and to identify pre-existing connections between transaction parties. A further advantage is the availability of annual tax assessments[1] for every property, which provide an unusually precise and exogenous benchmark for market value.

Using this setting, we document two main findings. First, buyers and sellers sharing the same migration background transact with each other far more

[1] In Dutch, this value is known as a property's *WOZ-waarde*, where "WOZ" stands for "waarde onroerende zaken" (i.e., value of immovable property).

often than random search would predict, even within narrowly defined local sub-markets and across migrant groups that differ widely in migration history and socioeconomic status. Second, these matched transactions sell at a price discount of around 4%, approximately €10,000 at the sample mean, relative to otherwise identical properties. Around half of this discount reflects non-arm's-length transactions—sales between parties with a pre-existing relationship, such as family members or neighbors—which are more common among co-ethnics and typically involve below-market prices. The remainder points to ethnic in-group favoritism: migrant sellers appear to voluntarily forgo financial surplus when transacting with co-ethnics, a phenomenon documented in experimental settings but not previously identified in the context of housing, the single largest financial transaction in most people's lives and their primary vehicle for wealth accumulation.

A key empirical challenge is separating preferences for same-ethnicity transactions from the tendency of ethnic groups to cluster in the same neighborhoods. If buyers and sellers of the same background simply concentrate in the same areas, matched transactions will be more frequent mechanically, without reflecting any active preference for co-ethnic counterparties. To address this, we follow Badarinza et al. (2022) and compare, within each neighborhood, the share of sellers from a given group among all transactions to the corresponding share among buyers from that same group. Under random search, these shares should coincide. Instead, we document substantial deviations: buyers systematically transact with sellers from the same migration background at rates that far exceed the group's baseline presence in the local market. In a neighborhood in The Hague, for instance, 34% of all sellers have a Turkish migration background, whereas this share rises to 71% among transactions involving Turkish buyers.

Studying price differentials in co-ethnic transactions is difficult, as it requires isolating the effect of buyer-seller identity from differences in the properties themselves. The key institutional feature that enables this is the availability of annual tax assessments for every Dutch property—administratively determined valuations set on January 1 of the year prior to the sale, and therefore exogenous to the transaction itself. Our main outcome is the gap between a property's assessed value and its realized transaction price, which captures the premium or discount negotiated between buyer and seller net of underlying property quality. The near-universe coverage of our transaction data means that many properties transact multiple times over our sample period, enabling property fixed effects that compare the same property sold under

different buyer-seller combinations at different points in time. As a first validation, we confirm that matched transactions do not involve systematically different properties: while same-ethnicity transactions involve properties of slightly lower assessed value, these differences disappear once property fixed effects are included.

Tightly controlling for local market conditions, buyer and seller income, and their ethnicity per se, we find that same-ethnicity transactions sell at a discount of around 4% relative to other buyer-seller combinations. This discount is robust to the inclusion of property fixed effects, suggesting that it is not driven by time-invariant property characteristics. Estimating separate models for each of the largest migrant groups active in the Dutch market yields similar discounts throughout, ruling out that the pooled result is driven by a small number of outlier groups. A remaining concern is that discounts may reflect time-varying property characteristics or maintenance levels, rather than buyer-seller interaction. We address this using information on subsequent sales: if a discount reflects the identity of the parties rather than property characteristics, the same home should earn a higher return when later resold. Consistent with this, properties previously transacted between same-ethnicity migrants earn subsequent returns comparable in magnitude to the initial discount, even after controlling for neighborhood trends and the timing of both transactions.

These findings point to buyer-seller identity as the driver of price discounts, but leave open the possibility that discrimination rather than co-ethnic preferences is the underlying force. If minority buyers systematically overpay when transacting with out-group sellers, same-ethnicity discounts may simply be the flipside of that premium rather than reflecting voluntary favoritism. We find little support for this interpretation. Migrant buyers do not systematically pay higher prices than Dutch buyers when transacting with out-group sellers, nor do they pay higher prices in neighborhoods with a high share of their own ethnic group—suggesting a limited role for taste-based discrimination or steering.

We then investigate the mechanisms behind same-ethnicity price discounts. A natural starting point is whether discounts simply reflect non-arm's-length transactions, which are more common among co-ethnics. The richness of our administrative data allows us to address this directly: using linked family network records covering the entire Dutch population, we identify transactions between family members, and using residential location data, we flag buyers purchasing in their own neighborhood. Together these proxies reveal that non-

arm's-length transactions account for around half of the observed discount, with the share especially large among smaller migrant groups. For larger and more established groups, however, a significant discount of around 2% persists even after excluding all transactions likely to involve pre-existing connections between buyers and sellers.

Disentangling the mechanisms behind the remaining discount is inherently challenging, as trust and in-group favoritism are conceptually related and both predict lower prices in same-ethnicity transactions. We cannot resolve this with a single clean test, but leverage the richness of our administrative data to construct a battery of proxies that shed light on which channel is most likely at work. If trust drives discounts, they should be largest when transaction uncertainty is highest—for instance, when buyers rely on mortgage financing, a common source of deal failure in the Dutch market (Reher and Valkanov, 2024), or when sellers are time-constrained having already committed to a new purchase. If instead in-group favoritism is the driver, discounts should be strongest among sellers who are most culturally embedded in their origin community. Such behavior is well documented across a range of economic settings, from donations (Fong and Luttmer, 2009) to lending decisions (Fisman et al., 2017), but has not previously been identified in housing transactions. Consistent with favoritism, discounts are unrelated to mortgage financing or seller time pressure, but are significantly larger for first-generation sellers, high-gain sellers, and sellers from high-kinship societies (Schulz et al., 2019).

This paper speaks to three strands of literature. First, we contribute to the literature on race and ethnicity in housing markets. Much of the existing work has focused on the supply side, examining whether minority buyers pay price premiums due to discrimination by sellers or brokers (Chambers, 1992; Kiel and Zabel, 1996; Ondrich et al., 2003; Myers, 2004; Ihlanfeldt and Mayock, 2009; Bayer et al., 2017; Christensen and Timmins, 2022; Box-Couillard and Christensen, 2024). Researchers have also studied how racial preferences shape residential location choices, documenting *White flight* following minority influx (Card et al., 2008; Boustan, 2010) and minority self-selection into segregated neighborhoods (Waldfogel, 2008; Aliprantis et al., 2024). Crucially, both strands study ethnicity on one side of the transaction only. We instead focus on the interaction between buyer and seller ethnicity, documenting that this shapes both matching patterns and transaction prices in ways that are orthogonal to discrimination and neighborhood sorting.

Most of this literature focuses on the United States, with its distinct history of

racial discrimination. One exception is Wong (2013), who uses data from Singapore to show that preferences for ethnic neighborhood compositions follow an inverted U-shape. One of the few studies that considers both buyer and seller ethnicity is Agarwal et al. (2019), who find that specialized brokers from the seller's own ethnic group facilitate same-ethnicity transactions, sometimes at the cost of lower sale prices for the seller. We extend this literature by documenting co-ethnic preferences in transactions as a broad-based phenomenon across a large number of distinct migrant groups, and by showing that ethnic matching has quantifiable financial consequences for housing wealth accumulation. The latter connects our paper to two recent contributions: Kermani and Wong (2026), who document racial gaps in housing returns driven by distressed sales, and Reher and Valkanov (2024), who show that seller uncertainty about buyer financing generates substantial price premia. We contribute to both by identifying a distinct mechanism—voluntary in-group favoritism between co-ethnic buyers and sellers—that shapes transaction prices and housing returns independently of financial distress or financing risk.

We also contribute to research on in-group bias in economic interactions. Following the introduction of *social identity theory* by Tajfel and Turner (1979), psychologists have long studied the importance individuals attach to group membership. In economics, group identity was first formalized by Akerlof and Kranton (2000), and a growing body of work has since examined how shared identity shapes economic behavior. One strand specifically investigates preferential treatment of in-group members (Ockenfels and Werner, 2014), often associated with shared race or ethnicity (Glaeser et al., 2000; Fong and Luttmer, 2009; Mantilla et al., 2021; Kudashvili and Lergetporer, 2022; Cetre et al., 2024), although this relationship is complex and highly context-dependent.[2] Most evidence comes from experimental studies (see Lane, 2016, for a meta-analysis). Exceptions include Fisman et al. (2017) and Chen et al. (2019), who document in-group favoritism in lending markets. We extend this literature by providing the first evidence of ethnic in-group favoritism in the housing market, where the financial stakes are at their highest.

The remainder of this paper is structured as follows. Section 5.2 describes the Dutch housing market and our data sources. Section 5.3 documents ethnic matching patterns and their variation across migrant groups. Section 5.4 presents our main findings on price differentials in same-ethnicity transactions

[2]For example, Fershtman and Gneezy (2001) find that in a trust game in Israel, players of Eastern origin are discriminated against, even by members of their own ethnic group, due to ethnic stereotypes.

and their robustness. Section 5.5 investigates the underlying mechanisms, distinguishing between non-arm's-length transactions and in-group favoritism. Section 5.6 concludes.

5.2 Data and Institutional Background

Homeownership rates in the Netherlands are around 60%, comparable to other Western European countries. However, due to rising house prices following the global financial crisis, access to homeownership has become increasingly difficult, especially in urban areas. One distinctive feature of the Dutch market is its very well-developed mortgage market. In theory, lending standards in this market are relatively generous, allowing for broad access. For instance, the stipulated upper limit on loan-to-value (LTV) ratios is 100%, and loans without down payment are commonly made in practice.[3] At the same time, relatively tight loan-to-income (LTI) limits make it difficult, particularly for young, low- to middle-income households, to access homeownership in their preferred locations. Moreover, lenders typically apply strict income verification requirements and require income stability, which poses a barrier to entry for self-employed individuals and those on temporary contracts.

5.2.1 Migrants in the Dutch Housing Market

The Netherlands is home to a large and diverse set of ethnic groups. In 2025, 16.8% and 11.8% of the population were first- and second-generation migrants, respectively (CBS, 2026). Broadly, migrants can be categorized into four groups. First, migrants from former Dutch colonies, including Indonesia, Suriname, and the former Netherlands Antilles in the Caribbean. Second, former guest workers and their descendants, particularly from Turkey and Morocco; these two groups are predominantly second-generation migrants and Dutch citizens. Third, cross-border migrants from Germany or Belgium. Fourth, highly skilled expats, often from China, India, or other European countries, who are typically first-generation migrants without Dutch citizenship.

Historically, ethnic minorities have been underrepresented among homeowners, relative to their population share, even after accounting for socioeconomic

[3] In the past, LTV limits exceeding 100% were common. After 2012, as a measure to reduce household debt, LTV limits were subsequently lowered from 106% to 100%.

characteristics (Zorlu et al., 2014). One potential explanation is that migrant buyers face additional barriers in the mortgage market compared with indigenous Dutch households (Aalbers, 2007). However, in line with positive net migration to the Netherlands for more than a decade (CBS, 2025), the share of migrants among homebuyers has risen from 16% in 2012 to 25% in 2024, driven largely by first-generation migrants. This is shown in Figure 5.1.

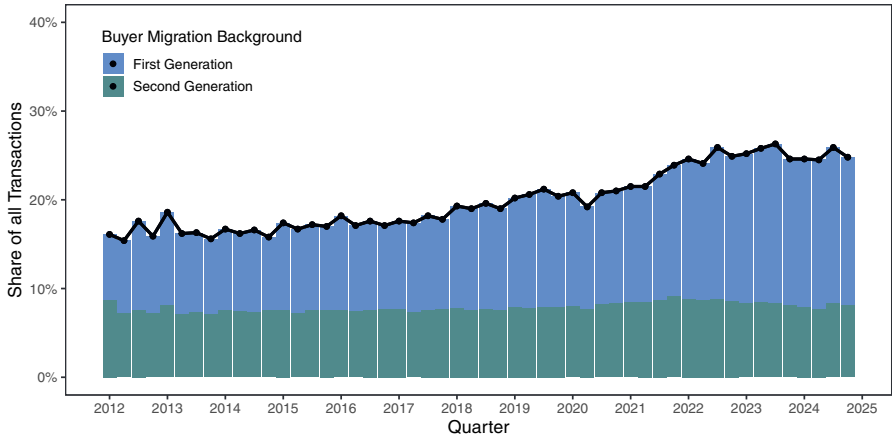


Figure 5.1: Share of Homebuyers with Migration Background

Notes: This figure shows the share of homebuyers with migration background in each quarter between 2012 and 2025. The solid line groups all buyers with migration background, while blue and green bars distinguish between first- and second-generation migrants.

5.2.2 Data

To study ethnic interactions in housing transactions, we combine data from two sources. Transaction records come from the Dutch Land Registry (Kadaster), which covers the universe of sales of existing properties. These records include the timing and sale price of each transaction, but they do not identify buyers and sellers, information that is essential for our analysis. To address this, we use microdata from Statistics Netherlands (CBS), which provides extensive administrative records at both the property and individual level. Property ownership records document the first owner of each residential unit for all years from 2011 to 2025. Using changes in these ownership records, we identify buyers and sellers for most transactions between 2012 and 2024. In total, buyers

and sellers can be identified for more than 87% of transactions. Transactions that are dropped include units with no recorded owner or cases where ownership changes are ambiguous, for example when they are recorded multiple years after the transaction date.

The richness of the CBS data allows us to infer a wide range of additional information by merging other datasets using unique identifiers for properties or for buyers and sellers. We determine each buyer's and seller's migration background using the CBS variable *origin country*. According to this variable, individuals are classified as Dutch only if they and both of their parents were born in the Netherlands. Otherwise, their origin is defined by the country in which they or one of their parents were born. If both parents were born in different countries, origin is determined by the mother's birth country. We define as migrant buyers or sellers all those with a non-Dutch origin according to this variable. This classification is independent of legal status and may differ from a person's nationality. Indeed, many individuals classified as migrants are Dutch citizens. The variable is therefore a proxy for ethnicity, which we use interchangeably with migration background in this paper.[4] Ethnicity is arguably more salient than nationality in capturing latent factors that may influence transaction outcomes, such as homophily or discrimination.

We add data on individual-level disposable income and household-level wealth for both buyers and sellers, measured in the transaction year. The wealth data are particularly detailed, as they decompose the household's portfolio into different asset types and reports current debt levels. Using residential location data from the year prior to the transaction, we identify local buyers who make a purchase in the neighborhood that they currently reside in. In addition, we use family network information, which identifies all family members and spouses across the Dutch population. This allows us to flag non-arm's-length "intra-family transactions," involving buyers and sellers who are related or are partners. At the property level, we include the property's tax value, measured on January 1 in the year prior to the transaction year.

Finally, we use an alternative transaction-level dataset from the largest Dutch association of real estate agents (NVM). Approximately 50% of the transactions in the main data involve an NVM-affiliated listing agent. By merging NVM records with the full Kadaster data, we identify whether a transaction

[4]We acknowledge that ethnicity and migration background are not synonymous. For example, an African American migrant and a White American migrant would both be classified as belonging to the same group in our analysis, despite clear ethnic differences.

was brokered. This classification is imperfect, however, as it misses transactions handled by non-NVM agents, and therefore includes false negatives.

The data are summarized in Table 5.1, where we distinguish between different types of buyer-seller interactions. Panel A reports transactions involving indigenous Dutch buyers, and Panel B focuses on migrant buyers. We further distinguish between "matched" and "unmatched" transactions, depending on whether buyer and seller share the same ethnicity. Among Dutch buyers, those purchasing from non-Dutch sellers have slightly higher income and asset holdings, on average, and tend to buy homes of somewhat higher value. Overall, migrant buyers are substantially less wealthy than Dutch buyers, though this may partly reflect measurement error if migrants hold assets abroad. Within the group of migrant buyers, there are large differences between transactions involving same-ethnicity or other-ethnicity sellers. Buyers purchasing from same-ethnicity sellers have lower income and liquid assets but more real estate. This may largely reflect the higher share of second-generation migrants among buyers in matched transactions. Transactions between migrants with shared ethnicity involve properties of much lower value and are significantly more likely to be unbrokered.

5.3 Ethnic Matching in Housing Transactions

The location choices of migrant homebuyers in the Netherlands are highly clustered. For example, although ethnic Turkish buyers account for less than 1.9% of all homebuyers nationwide, they make up 6.3% of buyers in The Hague and 48.1% in one neighborhood within the city. Previous research has typically analyzed the segregated home purchase patterns of minority buyers in terms of preferences or barriers to neighborhood access. In contrast, we are interested in preferences for ethnic matching—the tendency of buyers and sellers from the same group to transact with one another—independent of location. Disentangling these two phenomena is empirically challenging. If members of a given group share common preferences for neighborhoods, either due to ethnic homophily or due to shared amenity preferences, the share of "matched" transactions will mechanically be higher.

This pattern is illustrated in the cross-tabulation of buyer and seller ethnicity shown in Table 5.2. The first number in each cell reports the number of transactions between buyers and sellers of each group. Same-ethnicity trans-

Table 5.1: Summary Statistics

	Other-Ethnicity Seller			Same-Ethnicity Seller		
	Mean	Median	SD	Mean	Median	SD
Panel A: Dutch Buyers						
Income (€1,000)	50.89	37.67	47.50	46.33	35.72	40.99
Household Net Wealth (€1,000)	402.64	122.01	878.01	356.92	126.26	769.69
Household Liquid Assets (€1,000)	95.68	33.77	182.24	87.15	33.47	162.68
Household Real Estate (€1,000)	329.30	231.43	423.01	295.30	229.13	353.16
Sale Price (€1,000)	334.83	276.00	232.73	323.04	275.00	197.38
Share Brokered (%)		55.35			55.85	
Transactions		*190,622*			*1,481,981*	
Panel B: Migrant Buyers						
Income (€1,000)	48.63	37.69	43.76	42.42	32.92	39.35
Household Net Wealth (€1,000)	211.71	41.54	617.55	231.33	45.85	642.26
Household Liquid Assets (€1,000)	64.17	21.58	142.24	55.25	16.45	130.93
Household Real Estate (€1,000)	218.83	0.00	366.94	255.89	146.30	408.76
Sale Price (€1,000)	345.97	295.00	224.42	279.40	237.50	186.34
Share First Generation (%)		61.52			53.79	
Share Brokered (%)		56.75			27.44	
Transactions		*407,865*			*20,288*	

Notes: This table summarizes observable characteristics for transactions involving different types of buyer-seller interaction. Panel A summarizes all transactions featuring indigenous Dutch buyers, Panel B summarizes those involving buyers with migration background. Within each group, a distinction is made depending on whether the seller shares the buyer's ethnic background or not.

actions, shown along the main diagonal, consistently occur more frequently than predicted by the group's overall market share. This does not necessarily imply that buyers actively seek sellers from the same ethnic group. To separate ethnic matching preferences from location choice, we therefore analyze matching patterns at the sub-neighborhood level. Dutch neighborhoods are the smallest official geographic units defined by Statistics Netherlands, with an average population of around 1,500, roughly comparable to US census blocks. Within these small areas, variation in ethnic composition or amenities is limited, allowing us to more precisely isolate preferences for same-ethnicity transactions.

Table 5.2: Cross-Tabulation of Buyer and Seller Migration Background

Buyer \ Seller	Belgium	China	Dutch Caribbean	Germany	Great Britain	India	Indonesia	Morocco	Netherlands	Poland	Suriname	Turkey
Belgium	**N = 1519** **a = €237k** $\tilde{p}$ **= 0.111**	N = 35 a = €320k $\tilde{p}$ = 0.123	N = 73 a = €316k $\tilde{p}$ = 0.129	N = 392 a = €319k $\tilde{p}$ = 0.109	N = 136 a = €428k $\tilde{p}$ = 0.127	N = 17 a = €311k $\tilde{p}$ = 0.133	N = 396 a = €366k $\tilde{p}$ = 0.126	N = 48 a = €258k $\tilde{p}$ = 0.162	N = 13922 a = €308k $\tilde{p}$ = 0.121	N = 62 a = €262k $\tilde{p}$ = 0.118	N = 139 a = €297k $\tilde{p}$ = 0.149	N = 117 a = €224k $\tilde{p}$ = 0.125
China	N = 111 a = €336k $\tilde{p}$ = 0.124	**N = 681** **a = €266k** $\tilde{p}$ **= 0.082**	N = 123 a = €258k $\tilde{p}$ = 0.129	N = 274 a = €301k $\tilde{p}$ = 0.099	N = 148 a = €347k $\tilde{p}$ = 0.124	N = 62 a = €344k $\tilde{p}$ = 0.154	N = 535 a = €323k $\tilde{p}$ = 0.140	N = 210 a = €250k $\tilde{p}$ = 0.151	N = 11430 a = €303k $\tilde{p}$ = 0.141	N = 144 a = €249k $\tilde{p}$ = 0.207	N = 758 a = €262k $\tilde{p}$ = 0.148	N = 534 a = €244k $\tilde{p}$ = 0.158
Dutch Caribbean	N = 58 a = €290k $\tilde{p}$ = 0.144	N = 27 a = €334k $\tilde{p}$ = 0.124	**N = 198** **a = €262k** $\tilde{p}$ **= 0.082**	N = 171 a = €301k $\tilde{p}$ = 0.108	N = 52 a = €351k $\tilde{p}$ = 0.117	N = 13 a = €329k $\tilde{p}$ = 0.188	N = 309 a = €355k $\tilde{p}$ = 0.115	N = 63 a = €270k $\tilde{p}$ = 0.195	N = 7905 a = €313k $\tilde{p}$ = 0.124	N = 56 a = €270k $\tilde{p}$ = 0.160	N = 279 a = €281k $\tilde{p}$ = 0.103	N = 141 a = €241k $\tilde{p}$ = 0.166
Germany	N = 285 a = €336k $\tilde{p}$ = 0.103	N = 60 a = €313k $\tilde{p}$ = 0.144	N = 113 a = €323k $\tilde{p}$ = 0.128	**N = 1536** **a = €282k** $\tilde{p}$ **= 0.063**	N = 236 a = €440k $\tilde{p}$ = 0.130	N = 31 a = €338k $\tilde{p}$ = 0.132	N = 882 a = €368k $\tilde{p}$ = 0.119	N = 105 a = €274k $\tilde{p}$ = 0.131	N = 26957 a = €320k $\tilde{p}$ = 0.122	N = 111 a = €309k $\tilde{p}$ = 0.131	N = 308 a = €296k $\tilde{p}$ = 0.133	N = 292 a = €266k $\tilde{p}$ = 0.123
Great Britain	N = 131 a = €384k $\tilde{p}$ = 0.105	N = 46 a = €411k $\tilde{p}$ = 0.188	N = 75 a = €348k $\tilde{p}$ = 0.148	N = 377 a = €411k $\tilde{p}$ = 0.117	**N = 412** **a = €441k** $\tilde{p}$ **= 0.100**	N = 31 a = €598k $\tilde{p}$ = 0.082	N = 529 a = €436k $\tilde{p}$ = 0.139	N = 73 a = €326k $\tilde{p}$ = 0.199	N = 12346 a = €384k $\tilde{p}$ = 0.135	N = 70 a = €366k $\tilde{p}$ = 0.169	N = 247 a = €340k $\tilde{p}$ = 0.146	N = 129 a = €323k $\tilde{p}$ = 0.140
India	N = 104 a = €360k $\tilde{p}$ = 0.169	N = 209 a = €385k $\tilde{p}$ = 0.144	N = 147 a = €348k $\tilde{p}$ = 0.168	N = 265 a = €387k $\tilde{p}$ = 0.156	N = 191 a = €449k $\tilde{p}$ = 0.129	**N = 544** **a = €373k** $\tilde{p}$ **= 0.141**	N = 577 a = €388k $\tilde{p}$ = 0.163	N = 214 a = €350k $\tilde{p}$ = 0.206	N = 11254 a = €375k $\tilde{p}$ = 0.169	N = 176 a = €345k $\tilde{p}$ = 0.202	N = 820 a = €350k $\tilde{p}$ = 0.164	N = 499 a = €356k $\tilde{p}$ = 0.178
Indonesia	N = 329 a = €354k $\tilde{p}$ = 0.097	N = 99 a = €296k $\tilde{p}$ = 0.121	N = 195 a = €321k $\tilde{p}$ = 0.126	N = 974 a = €343k $\tilde{p}$ = 0.097	N = 364 a = €433k $\tilde{p}$ = 0.103	N = 41 a = €334k $\tilde{p}$ = 0.132	**N = 2143** **a = €366k** $\tilde{p}$ **= 0.090**	N = 170 a = €285k $\tilde{p}$ = 0.129	N = 36363 a = €345k $\tilde{p}$ = 0.118	N = 154 a = €293k $\tilde{p}$ = 0.154	N = 593 a = €290k $\tilde{p}$ = 0.124	N = 348 a = €266k $\tilde{p}$ = 0.121
Morocco	N = 67 a = €337k $\tilde{p}$ = 0.058	N = 32 a = €293k $\tilde{p}$ = 0.103	N = 72 a = €246k $\tilde{p}$ = 0.122	N = 187 a = €284k $\tilde{p}$ = 0.077	N = 64 a = €360k $\tilde{p}$ = 0.086	N = 11 a = €437k $\tilde{p}$ = 0.127	N = 257 a = €328k $\tilde{p}$ = 0.079	**N = 371** **a = €263k** $\tilde{p}$ **= 0.055**	N = 6069 a = €300k $\tilde{p}$ = 0.095	N = 67 a = €261k $\tilde{p}$ = 0.155	N = 335 a = €274k $\tilde{p}$ = 0.100	N = 322 a = €258k $\tilde{p}$ = 0.089
Netherlands	N = 11756 a = €315k $\tilde{p}$ = 0.100	N = 2556 a = €312k $\tilde{p}$ = 0.113	N = 5208 a = €296k $\tilde{p}$ = 0.122	N = 35049 a = €309k $\tilde{p}$ = 0.086	N = 8559 a = €380k $\tilde{p}$ = 0.110	N = 997 a = €355k $\tilde{p}$ = 0.126	N = 40020 a = €345k $\tilde{p}$ = 0.111	N = 4205 a = €272k $\tilde{p}$ = 0.132	**N = 1481981** **a = €309k** $\tilde{p}$ **= 0.108**	N = 5009 a = €272k $\tilde{p}$ = 0.145	N = 13397 a = €283k $\tilde{p}$ = 0.116	N = 10151 a = €262k $\tilde{p}$ = 0.121
Poland	N = 220 a = €237k $\tilde{p}$ = 0.116	N = 68 a = €249k $\tilde{p}$ = 0.126	N = 178 a = €218k $\tilde{p}$ = 0.111	N = 505 a = €239k $\tilde{p}$ = 0.098	N = 139 a = €267k $\tilde{p}$ = 0.118	N = 32 a = €281k $\tilde{p}$ = 0.140	N = 564 a = €249k $\tilde{p}$ = 0.113	N = 200 a = €208k $\tilde{p}$ = 0.130	N = 18333 a = €236k $\tilde{p}$ = 0.117	**N = 1192** **a = €221k** $\tilde{p}$ **= 0.159**	N = 561 a = €208k $\tilde{p}$ = 0.108	N = 601 a = €196k $\tilde{p}$ = 0.132
Suriname	N = 135 a = €300k $\tilde{p}$ = 0.117	N = 157 a = €274k $\tilde{p}$ = 0.147	N = 297 a = €258k $\tilde{p}$ = 0.098	N = 449 a = €281k $\tilde{p}$ = 0.100	N = 184 a = €319k $\tilde{p}$ = 0.119	N = 111 a = €303k $\tilde{p}$ = 0.139	N = 850 a = €309k $\tilde{p}$ = 0.118	N = 323 a = €258k $\tilde{p}$ = 0.154	N = 18970 a = €297k $\tilde{p}$ = 0.125	N = 205 a = €226k $\tilde{p}$ = 0.174	**N = 3400** **a = €245k** $\tilde{p}$ **= 0.068**	N = 746 a = €240k $\tilde{p}$ = 0.148
Turkey	N = 238 a = €299k $\tilde{p}$ = 0.099	N = 226 a = €315k $\tilde{p}$ = 0.127	N = 283 a = €286k $\tilde{p}$ = 0.122	N = 698 a = €288k $\tilde{p}$ = 0.096	N = 262 a = €372k $\tilde{p}$ = 0.120	N = 135 a = €384k $\tilde{p}$ = 0.118	N = 1083 a = €320k $\tilde{p}$ = 0.103	N = 714 a = €265k $\tilde{p}$ = 0.116	N = 25761 a = €304k $\tilde{p}$ = 0.118	N = 345 a = €248k $\tilde{p}$ = 0.164	N = 1711 a = €276k $\tilde{p}$ = 0.109	**N = 5641** **a = €239k** $\tilde{p}$ **= 0.102**

Notes: This table breaks down patterns of ethnic buyer-seller interaction in the Dutch housing market. Rows indicate the migration background of buyers, columns that of sellers. Each cell reports, for each type of buyer-seller interaction, the number of transactions (N), average tax values per January 1 of year $t-1$ (a) and average log point differences between sale prices and tax values ($\tilde{p}$). Diagonal entries (co-ethnic buyer and seller) are highlighted in bold.

5.3.1 Migrants Have Preferences for Ethnic Matching

To study matching rates within neighborhoods, we follow the methodology of Badarinza et al. (2022), which is based on a simple premise. If buyer-seller matching were random with respect to ethnicity, the share of sellers from country j among transactions involving buyers from the same country j would, in expectation, equal the share of sellers from country j across all transactions in that neighborhood. Deviations from this benchmark indicate preferences for (or against) ethnic matching. More formally, let $N_{ij,l}$ denote the number of transactions involving buyers from country i and sellers from country j in neighborhood l. The share of transactions involving sellers from country j among all transactions in a given neighborhood is then:

$$m_{j,l} = \frac{\sum_{i=1}^{I} N_{ij,l}}{\sum_{j=1}^{J}\sum_{i=1}^{I} N_{ij,l}}. \tag{5.1}$$

In contrast, the fraction of sellers from country j, only among transactions involving a buyer from country i is:

$$h_{ij,l} = \frac{N_{ij,l}}{\sum_{j=1}^{J} N_{ij,l}}. \tag{5.2}$$

In the absence of preferential matching, it should hold that $m_{j,l}$ is equal to $h_{ij,l}$ for $i = j$. For instance, if 17% of all sellers in a neighborhood are Turkish, the share of Turkish sellers among Turkish buyers alone should also be equal to 17%, if search is random and the sample is sufficiently large.

In Figure 5.2, we plot $m_{j,l}$ against $h_{ij,l}$ for the largest migrant groups in the Dutch housing market, and compare these patterns to indigenous Dutch buyers. For each group, we visualize the 100 neighborhoods most prominent among its buyers. The results show that preferences for ethnic matching exist in all groups, but to very different degrees. Among buyers from Indonesia, Germany, or Belgium, ethnic matching is relatively limited: $h_{ij,l}$ exceeds $m_{j,l}$ in most neighborhoods, but deviations from the neutral benchmark are generally small. In contrast, for buyers with Chinese, Indian, Moroccan, Polish,

Surinamese, or Turkish backgrounds, $h_{ij,l}$ exceeds $m_{j,l}$ in nearly all neighborhoods, often by a substantial margin. In other words, for most ethnic groups in our analysis, transactions between buyers and sellers of the same ethnicity occur far more frequently than standard search theory would predict, even when focusing on highly granular sub-markets.

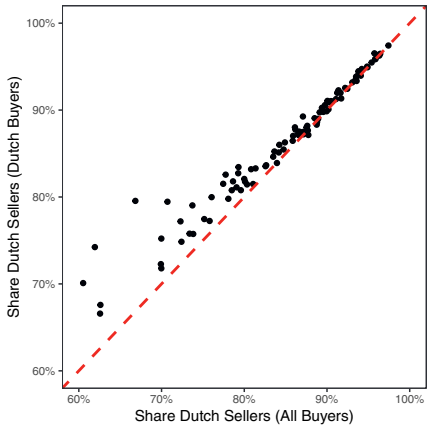
5.3.2 Preferences for Ethnic Matching Are Strongest in Small Groups

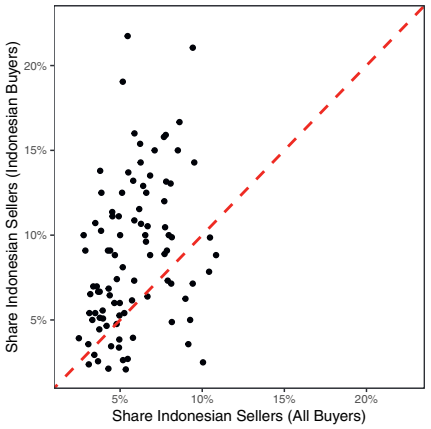
The plots in Figure 5.2 suggest that ethnic matching is widespread: most migrant groups exhibit at least some preference for same-ethnicity transactions. However, the strength of this preference appears inversely related to group size in the Dutch housing market. Among Indonesian migrants, the largest group of buyers with migration background (2.1% of all purchases), matching preferences are relatively weak. By contrast, Moroccan migrants, who account for only 0.4% of purchases, show much stronger preferences for same-ethnicity transactions.

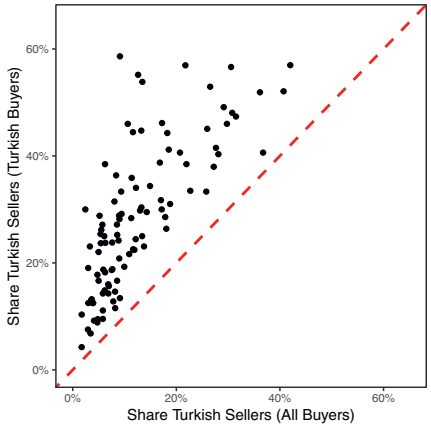
To investigate whether these patterns hold systematically, we extend the analysis to the 25 largest ethnic groups in the Dutch housing market. Rather than focusing on individual neighborhoods, we now aggregate transaction patterns across all neighborhoods. For each group, we first compute $h_{ij,l}/m_{j,l} - 1$ in each neighborhood, which measures preferential matching based on ethnicity (Badarinza et al., 2022). We then calculate the weighted median of this measure across all neighborhoods. A value greater than zero indicates evidence of preferential matching based on ethnicity.

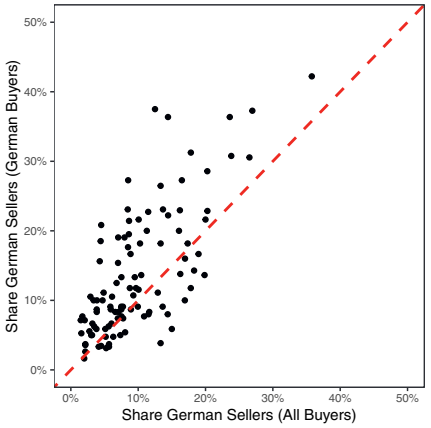
In Figure 5.3, aggregated matching preferences are plotted against the total number of purchases by buyers from each migration background. Panel (a) is based on the full transaction sample. Two main patterns emerge. First, all migrant groups exhibit some preference for ethnic matching. In contrast, indigenous Dutch buyers and sellers, who are excluded from the figure to preserve scaling, have a median value of $h_{ij,l}/m_{j,l} - 1$ equal to zero. Second, matching preferences are inversely related to group size: the largest migrant groups show only small deviations from random search, whereas smaller groups display increasingly non-random matching patterns.

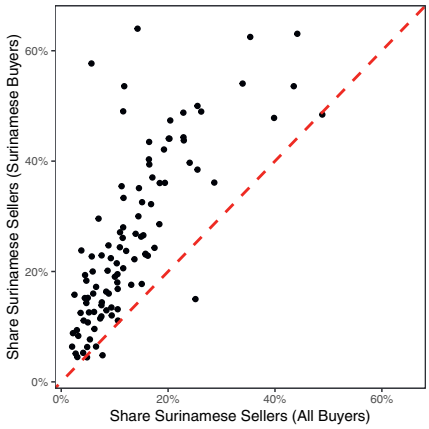
One concern is that same-ethnicity transactions in small groups may reflect non-arm's-length sales, such as transactions between family members. Be-

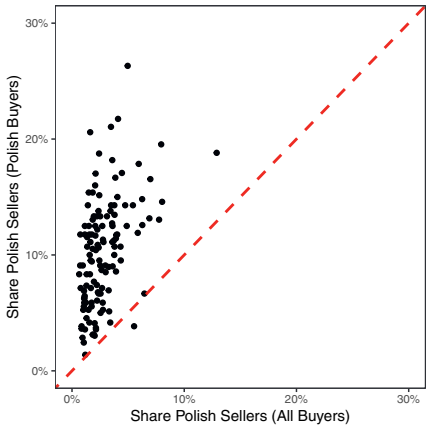

Share Dutch Sellers (Dutch Buyers)
100%
90%
80%
70%
60%
60%
70%
80%
90%
100%
Share Dutch Sellers (All Buyers)

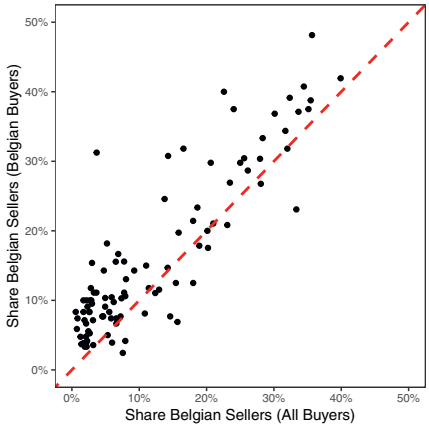

Share Indonesian Sellers (Indonesian Buyers)
20%
15%
10%
5%
5%
10%
15%
20%
Share Indonesian Sellers (All Buyers)

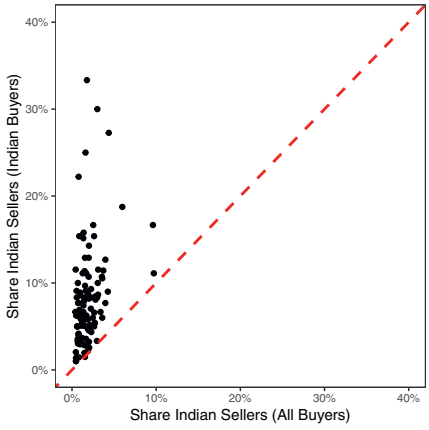

Share Turkish Sellers (Turkish Buyers)
60%
40%
20%
0%
0%
20%
40%
60%
Share Turkish Sellers (All Buyers)


Share German Sellers (German Buyers)
50%
40%
30%
20%
10%
0%
0%
10%
20%
30%
40%
50%
Share German Sellers (All Buyers)


Share Surinamese Sellers (Surinamese Buyers)
0%
20%
40%
60%
Share Surinamese Sellers (All Buyers)


Share Polish Sellers (Polish Buyers)
0%
10%
20%
30%
Share Polish Sellers (All Buyers)


Share Belgian Sellers (Belgian Buyers)
0%
10%
20%
30%
40%
50%
Share Belgian Sellers (All Buyers)


Share Indian Sellers (Indian Buyers)
0%
10%
20%
30%
40%
Share Indian Sellers (All Buyers)

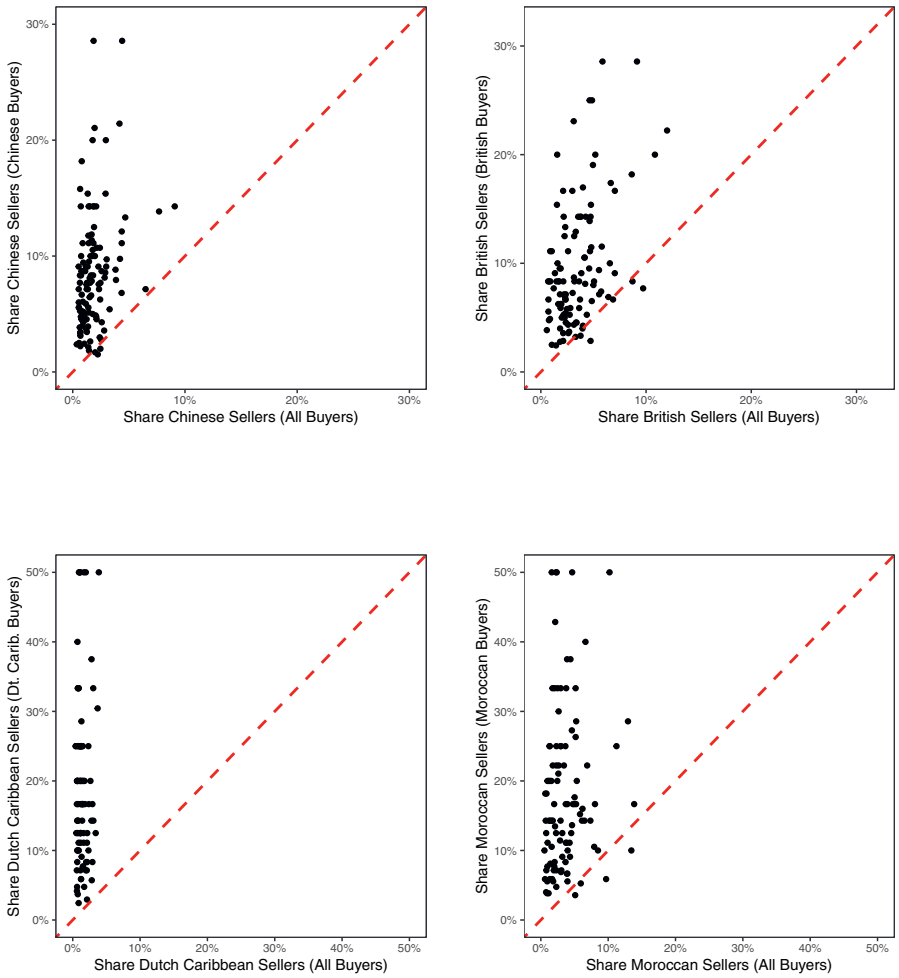


Figure 5.2: Matching Preferences by Ethnic Group

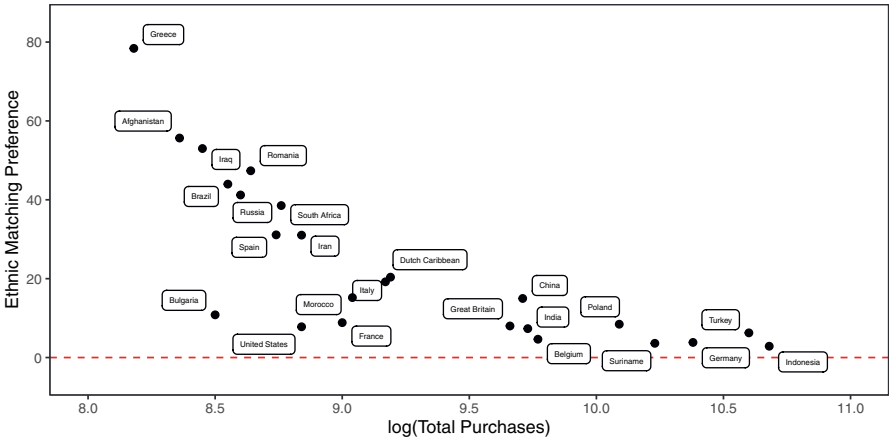
Notes: This figure visualizes preferences for ethnic matching across the largest migrant groups in the Dutch housing market. For each group, we select the 100 neighborhoods in which most purchases by group-members are made. Each dot represents a neighborhood. For each neighborhood, we compare the share of sellers from a given group (x-axis) to the share of sellers from that group, only among transactions involving buyers from the same group. Details are found in the main text.

cause we aim to isolate preferences for matching based solely on migration background, it is important to account for this confounding factor. In Panel (b), we therefore recalculate matching preferences after excluding all intra-family sales. This adjustment has little impact on the largest migrant groups but reduces the estimated preferences for smaller groups. Taken together, this suggests that non-arm's-length transactions are an important confounder, but leave the qualitative relationship between group size and ethnic matching preferences unchanged.

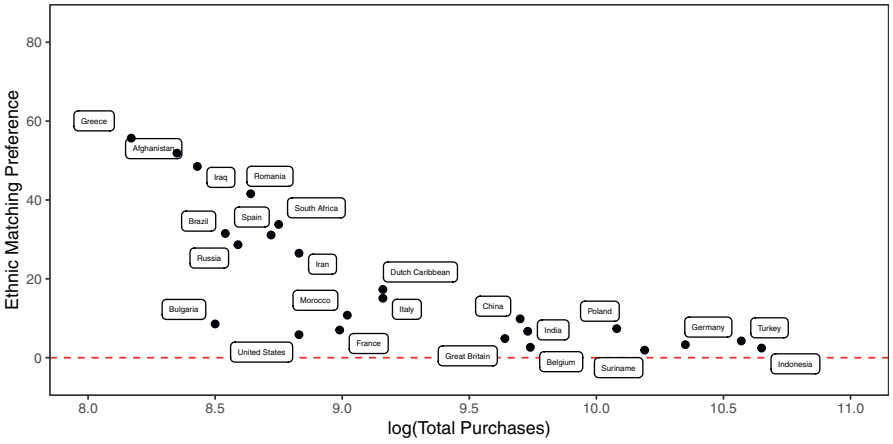
5.4 Price Differentials in Within-Ethnicity Transactions

We now turn to the question of whether within-group transactions among migrants entail a price differential, relative to other types of buyer-seller matches. Theoretically, the direction of such differentials is ambiguous. On the one hand, two channels suggest that migrant buyers may pay a premium to purchase from sellers of the same group. First, migrant buyers may face discriminatory constraints in their housing search, with the effect of steering them toward "ethnic neighborhoods", where same-ethnicity sellers are overrepresented. By concentrating the demand of migrant buyers in these neighborhoods, they may be forced to pay higher prices, relative to other buyers who do not have a constrained choice set. Evidence for discriminatory steering by real estate agents is well-established, particularly in the US context (e.g., Christensen and Timmins, 2022). Second, migrant buyers may exhibit homophily and therefore be willing to pay a premium for residing in close proximity to their co-ethnics (Aliprantis et al., 2024). Relatedly, migrants from the same group may share preferences for certain neighborhood amenities that other buyers do not value (Waldfogel, 2008).

On the other hand, within-ethnicity transactions may involve discounts if migrant sellers are willing to forgo some financial return in order to transact with buyers from their own group. Such discounts may arise from different underlying motivations. For instance, migrant sellers may have increased trust toward their co-ethnics and trade off a lower sale price in exchange for a higher (perceived) likelihood of successfully carrying out the transaction. The fact that trust levels are higher between people of shared ethnicity has been shown in survey data (Guiso et al., 2009). Alternatively, in-group favoritism may lead sellers to be altruistic toward their co-ethnics. In that case, sellers may be will-



(a) Full Sample



(b) Excluding Intra-Family Transactions

Figure 5.3: Preferences for Ethnic Matching Across Groups

Notes: This figure visualizes preferences for ethnic matching across the largest migrant groups in the Dutch housing market. For each group, the (weighted) median matching preference $(h_{ij,l}/m_{j,l} - 1)$ is calculated across all neighborhoods, and plotted against the total number of purchases made by buyers from that group. The red dashed line indicates the non-preferential benchmark. In panel (a), all transactions are included, whereas panel (b) excludes intra-family transactions. In both panels, indigenous Dutch buyers are excluded to avoid scaling issues on the x-axis. In either case, the median ethnic matching preference among indigenous Dutch buyers is equal to 0.

ing to forgo some financial surplus when negotiating with a group member, a behavior that has been shown in various economic interactions in experimental settings (Chen and Li, 2009).

5.4.1 Same-Ethnicity Transactions Involve Comparable Properties

In the canonical hedonic framework, pioneered by Rosen (1974), a home is viewed as a composite good, made up of all its structural and neighborhood attributes. Each attribute's shadow price is determined in a competitive market. The price of a complete home is then the linear combination of all its attributes, weighted by their respective shadow prices. In this model, buyers and sellers take the value for each attribute as given, such that there is no role for bargaining to influence transaction prices. However, Harding et al. (2003) show that if markets for individual home characteristics are thin, shadow prices are not well-defined, such that the interaction between buyers and sellers can influence final transaction outcomes. This insight is crucial to explain why different types of buyers end up paying differential prices for observationally equivalent homes. For instance, Bayer et al. (2017) find that African American and Hispanic homebuyers pay a price premium of around 2% and Goldsmith-Pinkham and Shue (2023) find a similar premium for female buyers. Bernstein et al. (2019) study the price effect of climate risk and discover that more sophisticated buyers apply a larger discount to properties exposed to sea level rise.

Our empirical analysis investigates whether same-ethnicity transactions among migrants involve price discounts. To test this, assume that the transaction price of a home is a function of its attributes, both observable (λ) and unobservable (ψ), of characteristics of the buyer (B) and the seller (S), and an indicator for "matched" transactions (M), involving buyers and sellers of the same ethnicity:

$$p = p(\lambda, \psi, B, S, M) \tag{5.3}$$

We are interested in the causal effect of M on prices. However, a reduced-form estimation of this relationship suffers from an obvious endogeneity problem:

homes sold in same-ethnicity transactions are different from others in ways that are at least partially unobservable to the researcher. For instance, matched transactions occur more frequently in ethnic (sub-)neighborhoods, which may offer a distinct set of amenities that is absent in predominantly Dutch neighborhoods. A threat to identification therefore arises because of a correlation between the matching dummy and unobserved property or neighborhood characteristics: $\text{Cov}(M, \psi) \neq 0$.

The identifying assumption we will make in order to recover the causal effect of ethnic matching on transaction prices is that same-ethnicity transactions do not differ from other transactions, after controlling for buyer and seller characteristics: $\psi \perp\!\!\!\perp M \mid (B, S)$. While this assumption cannot be tested directly, we evaluate its credibility by examining a related question—whether matched transactions are systematically different from others in *observed* characteristics. Our data allow for a compelling way to test this, by studying tax values of transacted homes. In the Netherlands, each residential property is assessed on an annual basis, with the objective of estimating its market value per January 1 of the previous year. Assessments are administered by municipalities and based on automated valuation models, comparing a property to others, which sold recently in the same neighborhood, and adjusting for observable characteristics. In some cases, tax value estimations are verified by human assessors. The natural logarithm of the most recent tax value for transaction i, which we denote by a_i, should therefore be a reasonable proxy for observable quality, mapping all property and neighborhood attributes into a single value. Tax values are generally considered to be good proxies for market values (Ahlfeldt et al., 2025) and in the Dutch context, they can explain over 90% of the observable variation in sales prices (Francke et al., 2024).[5] In Table 5.2, we show that across the country, average tax values vary substantially for various types of buyer-seller interaction.

To test for systematic differences in tax values depending on the ethnic interaction between buyers and sellers, we estimate the following model:

$$\begin{aligned} a_i = \mu_{b(i)} + \gamma_{s(i)} + \delta_{c(i)} + \beta\, migrant_i + \theta\, match_i \\ + \psi\,(migrant_i \times match_i) + \varepsilon_i, \end{aligned} \tag{5.4}$$

[5]Given that tax values are based on comparable sales, they should even capture unobserved neighborhood characteristics to some extent. However, given that they do not involve human inspection in most cases, it is plausible that they overlook some hard-to-observe property characteristics.

where $migrant_i$ indicates a buyer with migration background and $match_i$ indicates transactions involving same-ethnicity buyers and sellers. We include neighborhood-by-year-quarter-cell fixed effects, denoted by $\delta_{c(i)}$, to account for differences in tax values due to location choice and timing. In addition, we control for an increasingly restrictive set of fixed effects at the buyer and seller level, denoted by $\mu_{b(i)}$ and $\gamma_{s(i)}$, respectively. In the most restrictive specification, we further include property-level fixed effects, effectively estimating a repeat-sales model. While this reduces the sample to properties that transact multiple times, it allows us to control for all time-invariant unobserved characteristics at the property level.

Because both independent variables are binary, the model is fully interacted. Transactions between indigenous Dutch buyers and migrant sellers ($migrant_i = 0$ and $match_i = 0$) form the omitted baseline group. All coefficients are therefore interpreted relative to the average (log) tax value of these transactions within a given fixed-effect cell. This specification allows us to cleanly isolate ethnically matched transactions among migrants from all other types of buyer-seller interactions, including matches between Dutch buyers and sellers, which account for 71% of all transactions in the data.

The estimated results are reported in Table 5.3. The baseline model, reported in column (1), suggests that within a neighborhood-by-year-quarter cell, transactions between same-ethnicity migrants involve properties of 2.3% lower tax value, compared to transactions between migrant buyers and other-ethnicity sellers (including Dutch sellers). In comparison to Dutch-to-Dutch transactions, same-ethnicity transactions among migrants involve properties of 1.8% lower value ($0.005 - 0.023$). Including additional fixed effects reduces this difference. In column (3), we include fixed effects for the deciles of buyer income, seller income, buyer financial assets and seller financial assets. In addition, we include fixed effects for the buyer's and seller's ethnic group, thereby isolating the role of their interaction. This reduces the difference in tax values between same-ethnicity and other-ethnicity transactions among migrant buyers to 1.6%. Importantly, once property fixed effects are included (column (4)), all differences by buyer and seller ethnicity disappear. Similarly, we find no differences in tax values when restricting the sample to transactions involving only migrant buyers and sellers. Taken together, these findings suggest that, while same-ethnicity transactions between migrants of the same group do involve properties of slightly lower value, compared to other buyer-seller combinations, these differences can entirely be explained by time-invariant property or location characteristics. In other words, differences in

observable property values in same-ethnicity transactions are small and can fully be accounted for. Under the identifying assumption, this suggests that any remaining unobserved differences in property quality are unlikely to bias our estimates.

5.4.2 Same-Ethnicity Transactions Entail a Price Discount

We now return to the question of whether there is a price differential in matched transactions. To answer this, we adapt the model in Equation (5.4) slightly, by replacing the dependent variable with the difference between the natural logarithm of the sale price and the corresponding tax value:

$$\tilde{p_i} = p_i - a_i \tag{5.5}$$

$\tilde{p_i}$ approximates the percentage premium or discount paid in a given transaction, relative to the property's most recent tax value.[6]

Table 5.2 shows that average values of $\tilde{p_i}$ are lower in transactions involving migrant buyers and sellers from the same ethnic group. However, $\tilde{p_i}$ captures price premiums or discounts due to buyer-seller interaction only if differences in unobservable property or location quality are fully controlled for. Since our model accounts for virtually all observable property differences, we assume that unobserved differences are also controlled for. Based on this approach, we estimate the following regression model:

$$\begin{aligned}\tilde{p_i} = \mu_{b(i)} + \gamma_{s(i)} + \delta_{c(i)} + \beta\, migrant_i + \theta\, match_i \\ + \psi\,(migrant_i \times match_i) + \varepsilon_i,\end{aligned} \tag{5.6}$$

where, again, transactions between Dutch buyers and migrant sellers are the omitted group in the main specification.[7] Accordingly, all coefficients are in-

[6] This is similar to Avenancio-León and Howard (2022), who analyze the difference between market prices and tax values to study discrimination in property assessments among US homeowners. Here, we treat tax values as exogenous because they are set as of January 1 of the year prior to the transaction.

[7] In Appendix A.1, we estimate an alternative specification where we create a categorical

terpreted as relative to the average price premium or discount in these transactions, within a fixed-effect cell. The coefficient of interest, ψ, therefore does not measure whether same-ethnicity transactions among migrants differ from the property's tax value per se, but rather whether the difference between transaction price and tax value is systematically higher or lower than in other types of buyer-seller interactions.

We estimate price differentials analogously to tax values, using the same set of fixed effects across models. The resulting coefficients, reported in Panel B of Table 5.3, suggest that same-ethnicity transactions among migrants involve a substantial price discount. Within a neighborhood-by-year-quarter, migrant buyers who purchase a home from an in-group seller pay a 5.5% lower price, compared to those buying from another seller. Relative to indigenous Dutch buyers purchasing from Dutch sellers, they pay a 4.7% lower price (0.008 − 0.055). Controlling for buyer and seller ethnicity and financials reduces the estimated price discount to 4.2%, whereas including property fixed effects further reduces it to 3.8%. The stability of the observed price discount is notable, especially given that by including property fixed effects, which lowers the sample by more than three-quarters, all differences in property values are accounted for. Finally, in the migrant-only sample, the estimated discount in same-ethnicity transactions is equal to 4.5%.

Overall, our findings indicate that migrant homebuyers receive a substantial discount when purchasing from a seller who shares their ethnic background. Scaled to the average tax value in same-ethnicity transactions, this discount amounts to approximately €10,000, representing roughly 15% of the total liquid assets of the average buyer and seller in the sample. Before examining the mechanisms underlying this discount, we assess its robustness using additional tests.

5.4.3 Ethnic Matching Discounts Are Highly Robust

An obvious concern related to the estimated price discount is based on the fact that it pools all buyers and sellers with any migration background. Therefore, observed price discounts may be driven by a few ethnic groups. In order to understand whether ethnic matching discounts are a broad phenomenon, this needs to be addressed. Therefore, rather than combining all into one category

variable for all different forms of buyer-seller interaction.

Table 5.3: Ethnic Matching and Transaction Outcomes

	Dep. Variable: log(Tax Value)				
Panel A: Tax Values					
	(1)	(2)	(3)	(4)	(5)
Migrant Buyer	−0.003*	0.003			
	(0.001)	(0.002)			
Matched	0.005***	−0.002*	0.015***	0.002	
	(0.001)	(0.001)	(0.002)	(0.002)	
Migrant Buyer × Matched	−0.023***	−0.007	−0.016***	−0.003	−0.001
	(0.004)	(0.005)	(0.005)	(0.003)	(0.004)
Fixed Effects:					
Neighborhood × Year-Quarter	✓	✓	✓	✓	✓
Buyer & Seller Income Decile	✗	✓	✓	✓	✓
Buyer & Seller Fin. Assets Decile	✗	✓	✓	✓	✓
Buyer & Seller Ethnicity	✗	✗	✓	✓	✓
Property	✗	✗	✗	✓	✗
Observations	1,985,104	1,750,528	1,750,472	410,443	49,575
Adjusted R^2	0.657	0.710	0.707	0.992	0.838
	Dep. Variable: log(Price) − log(Tax Value)				
Panel B: Price Differentials					
	(1)	(2)	(3)	(4)	(5)
Migrant Buyer	0.010***	0.007***			
	(0.001)	(0.001)			
Matched	0.008***	0.004***	0.000	0.001	
	(0.001)	(0.001)	(0.001)	(0.002)	
Migrant Buyer × Matched	−0.055***	−0.050***	−0.042***	−0.038***	−0.045***
	(0.003)	(0.003)	(0.004)	(0.005)	(0.004)
Fixed Effects:					
Neighborhood × Year-Quarter	✓	✓	✓	✓	✓
Buyer & Seller Income Decile	✗	✓	✓	✓	✓
Buyer & Seller Fin. Assets Decile	✗	✓	✓	✓	✓
Buyer & Seller Ethnicity	✗	✗	✓	✓	✓
Property	✗	✗	✗	✓	✗
Observations	1,985,104	1,750,528	1,750,472	410,443	49,575
Adjusted R^2	0.473	0.488	0.489	0.834	0.580

Notes: This table shows differences in transaction outcomes between same-ethnicity transactions among migrants and other transactions. In Panel A, the dependent variable is the natural log of a property's last assessed tax value before the transaction. In Panel B, the dependent variable is the difference between the natural log of the sale price and the natural log of the property's last tax value. In columns (1)–(4), the model is estimated on the full transaction sample. In column (5), the sample includes only transactions involving both a migrant buyer and a migrant seller. Standard errors are clustered at the neighborhood and at the year-quarter level (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

($migrant_i$), we estimate separate models for each of the largest migrant groups in the Dutch market (e.g., $turkish_i$, $surinamese_i$, etc.).

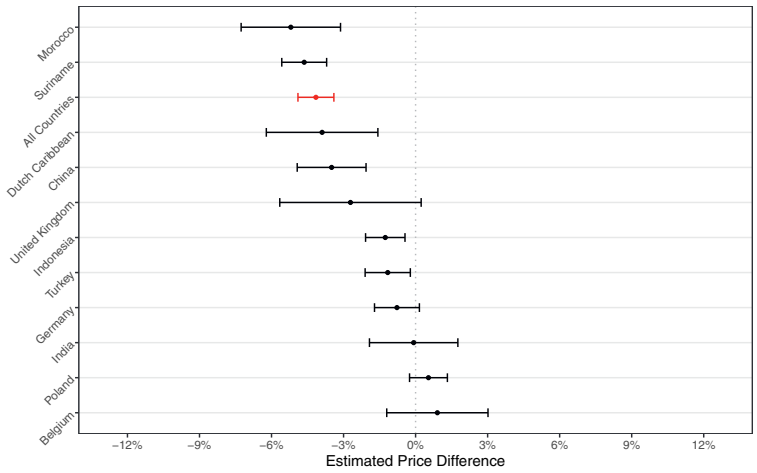
The resulting estimates are visualized in Figure 5.4. They yield two main insights: first, discounts in same-ethnicity transactions are widespread and not driven by a few outlier groups. Including property fixed effects substantially reduces the sample size and widens confidence intervals, but leaves the overall pattern and effect sizes unchanged. However, point estimates for individual ethnic groups are affected, changing the relative order of the coefficients. Second, price discounts in larger migrant groups are smaller than average. For most ethnic groups with substantial activity in the Dutch housing market, estimated discounts are smaller (in absolute terms) than the pooled effect, suggesting that the overall effect is partly driven by smaller migrant groups. We return to this issue below when investigating the underlying mechanism.

A more fundamental concern is that price discounts may not reflect ethnic preferences, but instead arise from unobservable property or location characteristics. Our empirical strategy addresses this concern in several ways. Since all models include neighborhood-by-year-quarter fixed effects, location effects would only matter if same-ethnicity transactions occur in housing blocks of lower unobservable quality within a neighborhood. Moreover, including property fixed effects ensures that all time-invariant property and location characteristics are controlled for. Any remaining bias would therefore arise from time-varying property characteristics. One potential source is differences in maintenance: if migrant homeowners who eventually sell to a co-ethnic invest less in property upkeep, this could lead to observed price discounts in same-ethnicity transactions.[8] Although we see no reason to expect this, it cannot be fully ruled out in the baseline framework.

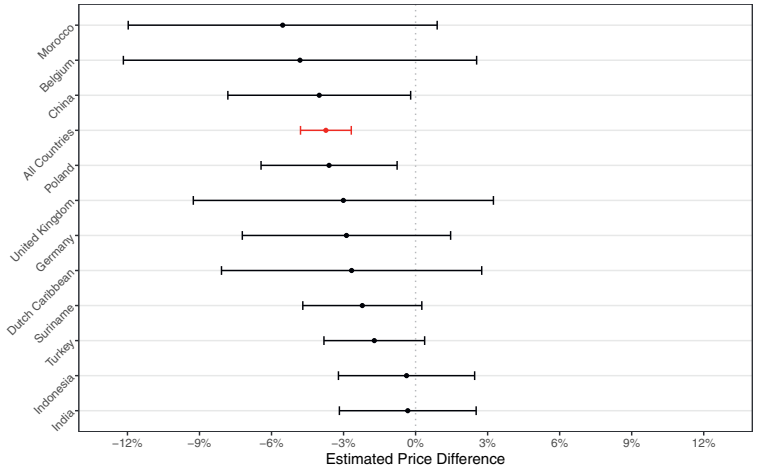
Therefore, we conduct an additional test, which allows us to distinguish genuine price discounts from confounding factors. This test is based on the simple premise that if a buyer gets a discount due to her interaction with the seller, she should earn a higher return when selling the same home in a subsequent transaction.[9] If, instead, the discount reflects any omitted factor associated

[8] This does not require migrant homeowners to invest less in maintenance on average, but specifically those who sell to a same-ethnicity buyer.

[9] The following simple example illustrates this: consider two identical properties, A and B, that are located in the same neighborhood and both sell twice, at identical dates, t and $t+1$. At time t, property A sells at a lower price than property B, while they both sell at the same price (i.e., their market value) at time $t+1$. Consequently, the initial buyer of property A has made a higher return than that of property B.



(a) Excl. Property Fixed Effects



(b) Incl. Property Fixed Effects

Figure 5.4: Price Discounts by Ethnic Group

Notes: This figure shows estimated price differentials in same-ethnicity transactions based on separate regressions for various migrant groups. Panels show results from models excluding (N = 1,750,472) and including (N = 410,443) property fixed effects, respectively. Estimates from the full model are shown in red. They correspond to the estimates from columns (3) and (4) in Panel B of Table 5.3. Error bands indicate 95% confidence intervals.

with the property or its location, this would also reduce the sale price in the future transaction. In that case, the initial buyer would earn no outsized return.

To implement this test, we restrict the sample to all transactions for which a subsequent transaction is recorded and calculate the annualized (unlevered) return between both transaction dates: $r_{is} = \frac{P_{is}}{P_{it}}^{\frac{1}{(s-t)}} - 1$, where P_{it} is a property's initial transaction price at time t, and P_{is} is the price in the subsequent transaction. Since we measure transaction dates at a quarterly level, we allow t and s to be non-integers. We reestimate the baseline model from Equation (5.6), using r_{is} as the dependent variable. Similar to Goldsmith-Pinkham and Shue (2023), we control for location choice and timing of both purchase and subsequent sale in an increasingly tight way.

The resulting coefficients are reported in Table 5.4. Across specifications, we find that those who make a same-ethnicity purchase earn a significantly higher return when reselling the property. In the baseline model with no controls, migrants buying from co-ethnic sellers make subsequent returns that are 4.2% higher ($0.083 - 0.041$) than those earned by indigenous Dutch buyers purchasing from Dutch sellers. When we include both neighborhood-by-initial-purchase-year-quarter and neighborhood-by-future-sale-year-quarter fixed effects as in Goldsmith-Pinkham and Shue (2023), estimated returns are 5.1% higher. In the full model including all buyer and seller financials as well as their ethnic background, the returns remain 3.6% higher, which is remarkably similar to the initial price discount estimated above. We conclude, therefore, that price discounts reflect a deliberate decision made by migrant sellers in favor of co-ethnic buyers, rather than any unobserved factors related to the property or its location.

5.5 Underlying Mechanism

Price discounts in housing transactions based on the shared ethnicity of buyers and sellers have not been documented before.[10] These observed discounts do not align with established mechanisms such as discriminatory steering or

[10] Box-Couillard and Christensen (2024) find lower prices in same-race transactions in the US context with only a few racial groups. However, they attribute this finding to discrimination by the majority group, rather than deliberate discounts made by same-race sellers. In contrast, Agarwal et al. (2019) find that same-ethnicity transactions can involve premiums or discounts, depending on the neighborhood composition.

Table 5.4: Ethnic Matching and Subsequent Returns

	Dep. Variable: Annualized Return				
	(1)	(2)	(3)	(4)	(5)
Intercept	0.558***				
	(0.002)				
Migrant Buyer	−0.032***	−0.030***	−0.028***	−0.031***	
	(0.003)	(0.003)	(0.002)	(0.003)	
Matched	−0.041***	−0.016***	−0.015***	−0.015***	0.000
	(0.007)	(0.002)	(0.002)	(0.003)	(0.006)
Migrant Buyer × Matched	0.083***	0.046***	0.057***	0.066***	0.036**
	(0.011)	(0.010)	(0.006)	(0.009)	(0.016)
Time Difference			0.134***		
			(0.002)		
Fixed Effects:					
Neighborhood × Year-Quarter Purchase	✗	✓	✓	✓	✓
Neighborhood × Year-Quarter Sale	✗	✗	✗	✓	✓
Year-Quarter Purchase × Year-Quarter Sale	✗	✗	✗	✓	✓
Buyer & Seller Income Decile	✗	✗	✗	✗	✓
Buyer & Seller Fin. Assets Decile	✗	✗	✗	✗	✓
Buyer & Seller Ethnicity	✗	✗	✗	✗	✓
Observations	405,939	305,592	305,592	214,236	177,454
Adjusted R^2	0.001	0.416	0.701	0.811	0.835

Notes: This table shows differences in annualized returns at subsequent sale, comparing same-ethnicity transactions among migrants and other transactions. The dependent variable is the annualized unlevered return between two sale dates. The independent variables Migrant Buyer and Matched are measured at the initial transaction. Standard errors are clustered at the neighborhood and at the year-quarter level (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

homophily (we investigate these alternative mechanisms in Appendix A.3). Instead, they suggest that sellers with migration background are willing to forgo financial surplus by giving a discount to their co-ethnics. It is natural to ask, therefore, which underlying mechanism leads to this behavior. We investigate this in the current section, by leveraging additional information from the administrative and transaction data. Before that, we consider the possibility that discounts are instead driven by non-arm's-length transactions, which may occur more frequently in same-ethnicity transactions.

5.5.1 Non-Arm's-Length Transactions Explain Half of the Observed Discount

Our transaction data include all sales of existing residential properties, including non-arm's-length transactions. This is a key confounding factor, as such transactions may bias estimates of price discounts. Non-arm's-length transactions—particularly intra-family sales—occur more frequently among buyers and sellers of the same ethnic group. Moreover, these transactions are likely to

involve price discounts relative to standard sales.[11] To investigate this, we use administrative records to construct three binary proxies for non-arm's-length transactions.

First, we identify intra-family transactions using linked family network data. This variable equals 1 whenever the buyer and seller share any family tie, regardless of the type, or if they are partners. This represents the clearest form of non-arm's-length transactions, as intra-family sales are typically not financially motivated. Intra-family sales are also more frequent among same-ethnicity transactions between migrants (22%) than among Dutch-to-Dutch sales (4%). Second, we define "local buyers" as those who previously resided in the same neighborhood where the transaction occurs. This is a less precise proxy for non-arm's-length transactions, as it may also reflect buyers' preferences or local market knowledge (Cvijanović and Spaenjers, 2021). Nonetheless, differences are substantial, as 26% of matched transactions among migrants involve local buyers, compared to 16% of Dutch-to-Dutch sales. Finally, we identify the presence of sell-side brokers and classify transactions without a broker as unbrokered. Here, again, large differences exist: 73% of transactions between same-ethnicity migrants are unbrokered, according to our proxy, compared to 44% of Dutch-to-Dutch transactions. As explained in Section 5.2.2, this variable is measured with error, as it only captures brokers affiliated with the largest industry association in the Netherlands. Moreover, the absence of a sell-side broker does not necessarily indicate that the transaction was not at arm's length. Sell-side brokers may also moderate price discounts arising due to ethnic in-group favoritism or trust—mechanisms that we will discuss below. To test how these three factors interact with observed price discounts in same-ethnicity transactions between migrants, we extend the baseline model as follows:

$$\begin{aligned}\tilde{p}_i = \mu_{b(i)} + \gamma_{s(i)} + \delta_{c(i)} + \beta\, migrant_i + \theta\, match_i + \omega\, interaction_i \\ + \psi\,(migrant_i \times match_i) + \phi\,(migrant_i \times interaction_i) \\ + \lambda\,(match_i \times interaction_i) + \pi\,(migrant_i \times match_i \times interaction_i) + \varepsilon_i,\end{aligned} \tag{5.7}$$

where $interaction_i$ is a placeholder for the three different variables introduced

[11] Specifically, if the sale price falls below the tax value ($\tilde{p}_i < 0$), the difference is subject to gift tax.

above: $intra_family_i$, $local_buyer_i$, and $unbrokered_i$. We are particularly interested in ω, which captures the baseline effect of each variable, as well as in ψ and π, which contrast same-ethnicity discounts in transactions where the additional factor is either present or not present.

The results of estimating the fully interacted model, separately for all three interaction terms, are reported in Table 5.5. First, intra-family sales—the clearest form of non-arm's-length transactions—account for roughly half of the observed price discount in same-ethnicity transactions among migrants. This is not because migrant sellers give larger discounts to family members, but rather reflects the higher frequency of intra-family sales in same-ethnicity transactions, combined with an average discount of about 10%. Local buyers pay lower prices on average, and those with a migration background receive an even larger discount from in-group sellers. Nonetheless, the baseline discount in same-ethnicity transactions remains significant at 3% for buyers from outside the neighborhood. Finally, the involvement of a sell-side broker is crucial for the existence of price discounts: when a broker is involved, no price differential is observed, whereas in unbrokered transactions, discounts exceed 5%. This suggests that price reductions reflect deliberate choices by sellers interacting directly with buyers, rather than transactions mediated by an intermediary.

To assess the remaining price discount after accounting for non-arm's-length transactions, we remove all transactions involving family members or local buyers and reestimate the baseline model separately for each major migrant group. The results are shown in Figure 5.5. Green bars replicate the baseline effect from panel (a) of Figure 5.4 (N = 1,750,472), black bars reflect the sample excluding intra-family transactions (N = 1,680,677), and blue bars further exclude transactions with local buyers (N = 1,444,043). Consistent with the pooled regression results, intra-family transactions contribute more to the observed price discount than local buyer transactions. Importantly, in the pooled sample, non-arm's-length transactions explain a larger share of the discount than in the largest migrant groups individually, suggesting that intra-family sales are especially prevalent among smaller migrant groups. Nonetheless, a highly significant price discount of approximately 2% persists even after excluding all transactions likely to involve pre-existing connections between buyers and sellers.

Table 5.5: Price Discounts: Underlying Mechanism

	Dep. Var: log(Price) − log(Tax Value)		
	(1)	(2)	(3)
Matched	0.002**	0.001	0.001
	(0.001)	(0.001)	(0.001)
Interaction	−0.099***	−0.018***	−0.010***
	(0.010)	(0.002)	(0.002)
Matched × Interaction	0.000	−0.005***	−0.003**
	(0.004)	(0.001)	(0.001)
Migrant Buyer × Interaction	−0.015***	−0.003*	0.003**
	(0.005)	(0.001)	(0.001)
Migrant Buyer × Matched	−0.019***	−0.030***	−0.002
	(0.003)	(0.003)	(0.003)
Migrant Buyer × Matched × Interaction	0.016***	−0.032***	−0.051***
	(0.006)	(0.004)	(0.004)
Fixed Effects:			
Neighborhood × Year-Quarter	✓	✓	✓
Buyer & Seller Income Decile	✓	✓	✓
Buyer & Seller Fin. Assets Decile	✓	✓	✓
Buyer & Seller Ethnicity	✓	✓	✓
Interaction Term	Intra-Family	Local Buyer	Unbrokered
Observations	1,750,472	1,750,278	1,598,731
Adjusted R^2	0.500	0.491	0.501

Notes: This table shows differences in price discounts between matched and unmatched transactions, and how these discounts vary based on additional transaction characteristics. It is based on the extended regression model from Equation (5.7). Each column reports results from a model with a different interaction term. Standard errors are clustered at the neighborhood and at the year-quarter level (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

5.5.2 Remaining Price Discounts Are Driven by In-Group Favoritism

We interpret the remaining 2% price discount, after excluding all potential non-arm's-length transactions, as reflecting a causal effect of shared ethnicity between buyers and sellers. The question remains, however, why sellers would be willing to give up around €5,000, on average, to sell their home to a co-ethnic, rather than to another buyer. In the final part of the analysis, we investigate two candidate explanations for this behavior, which are based on distinct but related concepts: in-group trust and in-group favoritism. While we cannot test the importance of these mechanisms in a causal way, our data offer a number of suitable proxies to assess whether observed discounts align with either of them.

Trust levels in economic exchanges tend to be higher between people who share

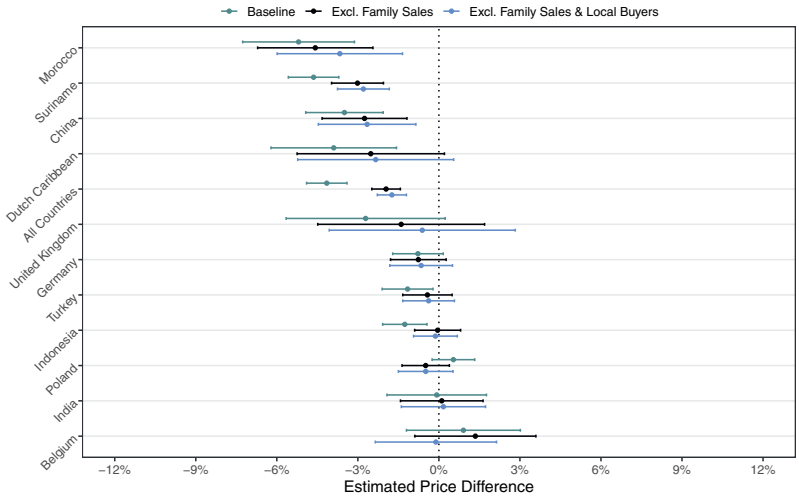


Figure 5.5: Price Discounts in Arm's-Length Transactions

Notes: This figure visualizes price discounts in transactions between migrants of the same ethnic background. Coefficient estimates (error bands show 95% confidence interval) are based on Equation (5.6), estimated either on the full sample (green), or a sample that excludes intra-family transactions (black) and local buyers (blue).

the same cultural background (Guiso et al., 2009). This could be of particular relevance in housing transactions, where buyers often make conditional offers that only become binding once certain conditions are met. In the Dutch market, it is very common that offers are conditional on mortgage approval. As Reher and Valkanov (2024) show, the uncertainty around buyers' mortgage approval can be so undesirable for sellers that they are willing to forgo substantial financial surplus in order to avoid it. In this setting, trust may act as a key moderating factor, leading sellers to offer price discounts when transacting with same-ethnicity buyers. Our first measure of trust therefore indicates whether buyers obtained a mortgage to finance the purchase.[12] If higher in-group trust is the driving factor behind price discounts, we would expect larger discounts in mortgage-financed transactions. For the second proxy, we identify sellers who have already made a subsequent home purchase just

[12]Since we do not have access to loan-level data, we construct this variable based on household-level balance sheet data, and label a transaction as mortgage-financed if the buyer's mortgage debt has increased between January 1 of the purchase year and January 1 of the following year.

before the transaction (i.e., in the transaction year of the previous one). Arguably, these sellers are more time-constrained, as they may need to access their home equity to finance their next purchase. In that case, a seller may be willing to offer discounts to buyers who they perceive as more trustworthy.[13] As a third proxy, we use a country-level measure of in-group versus out-group trust, which does not vary at the transaction level. This measure, developed by Schulz et al. (2019), combines items from the World Values Survey (Haerpfer et al., 2020) to capture the extent to which individuals place greater trust in in-group members (e.g., family, neighbors, acquaintances) relative to out-group members (e.g., foreigners or individuals of a different religion).[14] If higher trust of same-ethnicity buyers leads sellers to give a discount, we expect this behavior to be especially prevalent among sellers from countries with high levels of in-group trust.

Turning to measures of in-group favoritism, we start by comparing sellers who are first-generation migrants to those who are second-generation. This is motivated by the fact that, with increased length of stay, migrants tend to assimilate to the culture of the host country and relinquish that of their origin country (Constant and Zimmermann, 2008), even though this process is complex and heterogeneous. Therefore, we hypothesize that first-generation migrants exhibit stronger in-group bias toward their co-ethnics, compelling them to offer price discounts in same-ethnicity transactions. The second proxy relates to the concept of inequity aversion (Fehr and Schmidt, 1999), whereby individuals suffer a disutility if they perceive a distribution of outcomes as inequitable. While inequity aversion is most pronounced for disadvantageous outcomes, advantageous inequality can also entail a utility cost. In other words, decision makers are willing to give up some financial surplus in order to achieve a more equitable outcome. Chen and Li (2009) show that this behavior is more pronounced in interactions between in-group members. In contrast to experimental settings, housing transactions lack a clear reference point that defines a fair division of surplus between buyer and seller. To proxy for advantageous outcomes for the seller, we calculate the financial return that the seller has made since purchasing the unit, based on the difference in tax values.[15] We hypothesize that sellers who realize high returns on their property

[13] Relatedly, Koster and Rouwendal (2024) show that home sellers are willing to sell at lower prices in order to achieve a shorter time on market, especially when they have already moved to a new home.

[14] Concretely, the variable calculates the difference between out-group trust and in-group trust. Negative values indicate countries in which in-group trust is higher than out-group trust.

[15] Returns are calculated as: $(a_t - a_p)/a_p$, where a_t denotes the current tax value and a_p

are more likely to offer price discounts, particularly when interacting with co-ethnic buyers, consistent with in-group favoritism. For the final proxy, we rely on Schulz et al. (2019) once again, and incorporate the kinship intensity index (KII). This index, composed based on data from the Ethnographic Atlas, captures to what degree societies are based on kin-based institutions (e.g., co-residence of extended families, community endogamy). One aspect of kinship-based societies is a stronger emphasis on in-group loyalty, compared to more individualistic societies. We test, therefore, whether sellers with a migration background from a country with strong kinship norms are more prone to give discounts to in-group members.

Empirically, we incorporate these different proxies for in-group trust and favoritism as interaction terms in the baseline model in Equation (5.6). To ease interpretation by reducing the number of terms, we focus on the migrant-only sample, where all transactions involving indigenous Dutch buyers or sellers are removed.[16] This allows us to examine only variation in the average discount in transactions between same-ethnicity migrants. The estimation results are presented in Table 5.6. The estimates based on the trust proxies, reported in columns (1)–(3), provide little support for the hypothesis that price discounts are driven by elevated trust toward same-ethnicity buyers. Contrary to a trust-based explanation, discounts are smaller—rather than larger—when sellers have already made a subsequent purchase or when buyers rely on mortgage financing. Similarly, country-level differences in out-group-versus-in-group trust have no effect on observed price discounts.

Instead, the findings reported in columns (4)–(6) support the notion that observed discounts are related to in-group favoritism. First, price discounts are driven entirely by first-generation migrants, while no significant differences are observed among second-generation migrant sellers. Second, we find that discounts are related to the financial gain that a seller has realized: per 10 percentage point increase in housing returns, same-ethnicity transactions sell at a 0.2% lower price, compared to other transactions. Finally, sellers from countries characterized by higher kinship intensity give larger discounts to their in-group members. A one-unit increase in this standardized measure—approximately the difference between the United States and Hungary—is associated with a 1.4% lower price in same-ethnicity transactions relative to others.

denotes the tax value in the year that the current seller has purchased the property.

[16] Note that, since the model is based on the migrant-only sample, the price discount in same-ethnicity transactions among migrants is captured by $match_i$, rather than $migrant_i \times match_i$, as in the baseline model.

Taken together, while not causal, these findings suggest that price discounts in matched transactions reflect an altruistic inclination toward co-ethnic buyers.

Table 5.6: Ethnic Matching Discounts: Trust vs. In-Group Favoritism

	Dep. Variable: log(Price) − log(Tax Value)					
	In-Group Trust			In-Group Favoritism		
	(1)	(2)	(3)	(4)	(5)	(6)
Interaction	0.022^{***}	0.010^{***}		-0.017^{***}	-0.017^{***}	
	(0.002)	(0.003)		(0.002)	(0.006)	
Matched	-0.021^{***}	-0.033^{***}	-0.009^{***}	0.004	-0.010^{**}	-0.014^{***}
	(0.004)	(0.007)	(0.003)	(0.004)	(0.005)	(0.003)
Matched × Interaction	0.025^{**}	0.030^{***}	0.005	-0.026^{***}	-0.021^{***}	-0.014^{***}
	(0.011)	(0.008)	(0.004)	(0.005)	(0.005)	(0.004)
Fixed Effects:						
Neighborhood × Year-Quarter	✓	✓	✓	✓	✓	✓
Buyer & Seller Income Decile	✓	✓	✓	✓	✓	✓
Buyer & Seller Fin. Assets Decile	✓	✓	✓	✓	✓	✓
Buyer & Seller Ethnicity	✓	✓	✓	✓	✓	✓
Interaction Term	Subsequent Purchase	Mortgage-Financed	Out-vs-In Trust	First-Gen Seller	Seller Return	Kinship Intensity
Observations	38,938	22,340	20,335	38,938	24,305	35,341
Adjusted R^2	0.617	0.642	0.398	0.616	0.640	0.396

Notes: This table shows differences in price discounts between matched and unmatched transactions, and how these discounts vary based on additional transaction characteristics. It is based on the following model, estimated on the migrant-only sample: $\tilde{p}_i = \mu_{b(i)} + \gamma_{s(i)} + \delta_{c(i)} + \theta\, match_i + \psi\,(match_i \times interaction_i) + \varepsilon_i$. Each column reports results from a model with a different interaction term. Columns (1)–(3) include interaction terms that proxy for in-group trust as a mechanism behind price discounts in same-ethnicity transactions, whereas columns (4)–(6) test for in-group favoritism. Since the estimation is based on the migrant-only sample, there is no "Migrant Buyer" effect in this model. Therefore, in contrast to the baseline model from Equation (5.6), price differences in same-ethnicity transactions among migrants are captured by "Matched". Standard errors are clustered at the neighborhood and at the year-quarter level (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

5.6 Conclusion

This paper uses Dutch administrative data to study ethnic matching in the housing market and its implications for transaction prices. Combining the near universe of residential property transactions with detailed records on the migration background of buyers and sellers, we document that buyers and sellers who share the same migration background match with each other at rates that substantially exceed what random search would predict. This pattern is widespread: it holds across all migrant groups we examine, within narrowly defined local sub-markets, and is more pronounced among groups that represent a small share of the buyer population. Ethnic matching in housing transactions is therefore a broad-based phenomenon that is distinct from neighborhood sorting due to ethnic homophily.

Turning to transaction prices, we find that same-ethnicity transactions are associated with a significant price discount of around 4% relative to otherwise comparable transactions. This discount is highly robust. It holds across migrant groups, survives the inclusion of property fixed effects, and is confirmed by an analysis of returns on subsequent sales, which suggests that discounted properties earn higher resale returns of comparable magnitude. Taken together, these findings indicate that price discounts reflect intentional rebates by sellers toward buyers who share their migration background, rather than differences in unobserved property or neighborhood characteristics. Importantly, the discount does not appear to reflect discrimination: migrant buyers do not systematically overpay in transactions with out-group sellers, nor do they pay premiums in neighborhoods where their ethnic group is well-represented, which would be consistent with discriminatory steering by brokers.

Investigating the mechanisms behind same-ethnicity discounts, we find that around half of the discount is explained by non-arm's-length transactions, particularly sales to family members and buyers from the same neighborhood. This pattern is especially pronounced among small migrant groups. For larger and more established groups, however, a significant discount persists after removing likely non-arm's-length transactions. Further analysis of the remaining discount points to in-group favoritism, rather than elevated in-group trust, as the dominant mechanism. The discount is largest for first-generation migrant sellers, sellers who have accrued a larger financial gain from their home, and sellers with a migration background from countries where kin-based institutions are especially strong.

Our findings have implications for the understanding of residential segregation. Even in the absence of discrimination by sellers, brokers, or lenders, preferential matching between ethnically similar buyers and sellers can reinforce patterns of segregation. The financial incentives embedded in same-ethnicity transactions—a price discount that sellers give voluntarily—provide a novel mechanism through which in-group preferences shape housing market outcomes. This mechanism operates within local sub-markets and is therefore orthogonal to the neighborhood-level sorting that has received the bulk of attention in the literature.

More broadly, this paper contributes to the growing body of evidence on ethnic in-group favoritism in real economic transactions. Prior work has documented similar patterns in lending, experimental games, and donation behavior. Our results extend this to the housing market, where the stakes are considerably

higher: for most individuals, purchasing a home is the single largest financial transaction of their lifetime. That in-group preferences are strong enough to shape outcomes in a transaction of this magnitude underscores both their salience and their potential economic consequences.

A Appendix

A.1 Alternative Estimation of Price Differentials

A complication in the baseline model from Equation (5.6) is that transactions between migrant buyers and Dutch sellers may constitute a non-natural benchmark group. This complicates the interpretation of the estimated coefficients. We therefore implement an alternative model, where we interact the ethnic background of buyers and sellers to construct a categorical variable that can take on five distinct values for the different possible combinations of seller-to-buyer matches: Dutch-to-Dutch (omitted), Dutch-to-migrant, migrant-to-Dutch, migrant-to-migrant-matched, migrant-to-migrant-unmatched. The last two values distinguish between transactions involving two migrants of identical or different ethnicity. Using this variable, we reestimate the baseline model of price differentials. The resulting coefficients, reported in Table 5.7, align with the main findings. Compared to all other types of buyer-seller interactions, transactions between same-ethnicity migrants involve a substantial price discount. In the full specification, involving property fixed effects, the estimated discount is somewhat smaller than that in the main text, at 2.7%. One reason may be that, due to perfect collinearity, this specification does not allow for the inclusion of buyer and seller ethnicity fixed effects.

A.2 Placebo Test: Random Match Permutation

The baseline results suggest that price discounts arise exactly when buyers and sellers who share the same migrant background transact with each other. To investigate whether discounts require "exact" matching of buyer and seller ethnicity, or whether they arise in other cases as well, we implement a simple placebo test.

Instead of assigning $Matched = 1$ when buyers and sellers share the same ethnic background, we randomly permute the matching structure across buyer groups and reestimate the baseline model from Equation (5.6) 10,000 times using these counterfactual matches. For instance, in one simulated run, transactions by Turkish buyers may counterfactually be labeled as matched if they purchase from a Moroccan seller, whereas those by Chinese buyers may be treated as matched if they transact with Dutch sellers. Evidence of systematic discounts under these placebo assignments would indicate that factors other than ethnic matching drive the observed price differentials.

Table 5.7: Ethnic Matching and Price Differentials

	Dep. Var: log(Price) − log(Tax Value)		
	(1)	(2)	(3)
Dutch-to-Migrant	0.009^{***}	-0.004^{***}	-0.001
	(0.001)	(0.001)	(0.001)
Migrant-to-Dutch	0.004^{***}	0.004^{***}	0.008^{***}
	(0.001)	(0.001)	(0.001)
Migrants (Unmatched)	-0.005^{***}	0.000	0.010^{***}
	(0.001)	(0.001)	(0.002)
Migrants (Matched)	-0.048^{***}	-0.045^{***}	-0.027^{***}
	(0.003)	(0.003)	(0.005)
Fixed Effects:			
Neighborhood × Year-Quarter	✓	✓	✓
Buyer & Seller Income Decile	✗	✓	✓
Buyer & Seller Fin. Assets Decile	✗	✓	✓
Property	✗	✗	✓
Observations	1,985,104	1,750,528	410,665
Adjusted R^2	0.472	0.488	0.833

Notes: This table shows differences in transaction outcomes between same-ethnicity transactions among migrants and other transactions. The dependent variable is the difference between the natural log of the sale price and the natural log of the property's last tax value. The model is estimated on the full transaction sample. Standard errors are clustered at the neighborhood and at the year-quarter level (significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

However, our simulations provide no evidence for this. In Figure 5.6, we plot a histogram of estimated treatment effects (Migrant Buyer × Matched) based on 10,000 simulated runs. As expected, the estimated coefficients resemble a normal distribution. In fact, the average price differential across simulations is slightly positive. In contrast, the "true" coefficient of −4.2%, estimated from the baseline model with buyer and seller financial and ethnicity fixed effects, lies far in the left tail of the distribution. In more than 99.5% of the simulations, the estimated price differential exceeds (i.e., is less negative than) that obtained under the true matching of buyer and seller ethnicity.

A.3 Seller Discrimination and Neighborhood Steering

Most of the existing literature has studied the position of racial or ethnic minorities in the housing market through the lens of discrimination. The price discounts in same-ethnicity transactions that we document may be related

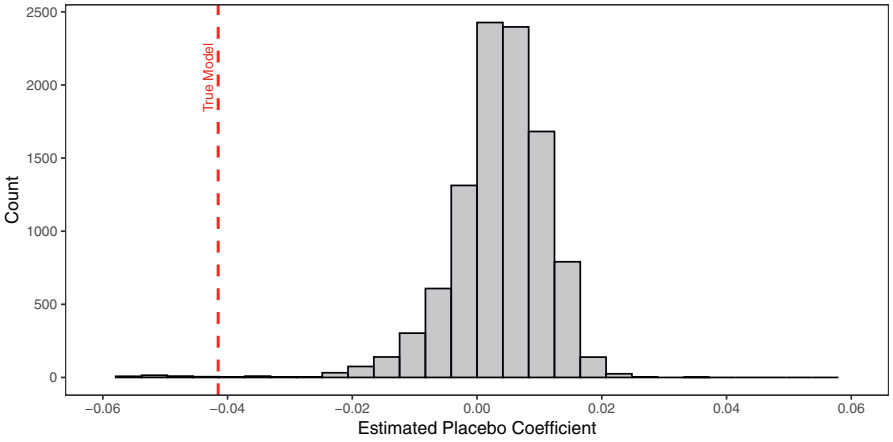


Figure 5.6: Placebo Test: Random Allocation of Matching Dummy

Notes: This figure plots a histogram of estimated price differentials based on random permutations of country pairs that constitute an ethnic match. Each model is based on the baseline model specified in Equation (5.6), including the full set of fixed effects apart from property fixed effects. Coefficients from the "true" model are shown by the dashed red line. Further details are given in the text.

to discrimination. For instance, rather than reflecting a discount granted by co-ethnic sellers, the observed pattern could reflect a premium that migrant buyers pay when transacting with out-group sellers. In order to establish ethnic matching discounts as a phenomenon that is distinct from the various forms of discrimination in the housing market, we therefore carry out an additional analysis.

The most commonly studied forms of discrimination in the market for owner-occupied housing include taste-based discrimination by sellers and discriminatory steering by brokers. Taste-based discrimination predicts that minority buyers have to pay a premium in order to compensate the prejudiced seller for their disutility from interacting with them. Evidence for taste-based discrimination, especially against African American and Hispanic buyers, has been found in the US market (Box-Couillard and Christensen, 2024), though other studies question whether price premia are the result of outright taste-based discrimination (Bayer et al., 2017). Discriminatory steering represents the tendency of buy-side brokers to direct minority clients toward certain neighborhoods. In the US, steering of African American buyers away from White neighborhoods—which often offer better neighborhood amenities—is well doc-

umented (Ondrich et al., 2003). Effectively, by constraining minority buyers' choice set in this way, their effective demand in ethnic neighborhoods is increased. One key empirical prediction of discriminatory steering is that it may lead minority buyers to pay higher prices in ethnic neighborhoods than other buyers, whose choice sets are not constrained.[17]

To investigate the importance of taste-based discrimination and steering in our setting, we extend the baseline empirical model from Equation (5.6) to incorporate information on the neighborhood in which a purchase occurs. Specifically, for each buyer from the largest migrant groups, we compute the share of co-ethnics in the population of the neighborhood where the purchase takes place. We define a neighborhood as "ethnic" for a given group if the share of co-ethnics exceeds the group's national population share in that year. Then, we extend the baseline model as follows:

$$\begin{aligned}\tilde{p}_i = \mu_{b(i)} + \gamma_{s(i)} + \delta_{c(i)} + \beta\, migrant_i + \theta\, match_i + \psi\,(migrant_i \times match_i) \\ + \phi\,(migrant_i \times ethnic_neighborhood_{b(i)}) \\ + \lambda\,(match_i \times ethnic_neighborhood_{b(i)}) \\ + \pi\,(migrant_i \times match_i \times ethnic_neighborhood_{b(i)}) + \varepsilon_i. \end{aligned} \tag{5.8}$$

This extended model allows us to contrast same-ethnicity price discounts with different forms of supply-side discrimination. As before, ψ captures price discounts in transactions between same-ethnicity migrants. β reveals whether migrant buyers have to pay a price premium, regardless of the seller's identity, which would be evidence for taste-based discrimination by sellers. We proxy for discriminatory steering, which concentrates migrant buyers' demand in ethnic neighborhoods, with ϕ. Finally, π, the coefficient on the triple-interaction, captures whether the discount in same-ethnicity transactions among migrants depends on the neighborhood's ethnic composition. This relates to Agarwal et al. (2019), who find that sellers give discounts in same-ethnicity transactions only in non-ethnic neighborhoods. We estimate the model separately for each migrant group relative to indigenous Dutch buyers, since $ethnic_neighborhood_{b(i)}$ varies by buyer ethnicity and is not constant across groups.

[17] Empirically identifying discriminatory steering is complicated by the fact that it is observationally equivalent to neighborhood sorting based on preferences for homophily or amenities. Christensen and Timmins (2022) overcome this challenge by examining paired-actor audit studies.

We visualize the coefficients of interest in Figure 5.7. Across the board, we continue to find a significant price discount in same-ethnicity transactions between buyers and sellers with migration background (black bars). In contrast, there is no direct evidence for taste-based discrimination, as none of the migrant groups appears to be paying a premium when purchasing from an out-group seller (dark blue bars). Similarly, we find no consistent price premia for purchases in ethnic neighborhoods (green bars), indicating that discriminatory steering or homophily play a limited role. Finally, price discounts in same-ethnicity transactions between migrants appear to be somewhat smaller in ethnic neighborhoods (light blue bars), although this relationship is noisier and less consistent. In accordance with the findings of Agarwal et al. (2019), this suggests that sellers make price discounts to their co-ethnics predominantly in neighborhoods where their own group is underrepresented. Overall, these results indicate that ethnic matching discounts among migrants operate independently of other well-established mechanisms, such as taste-based discrimination or steering.

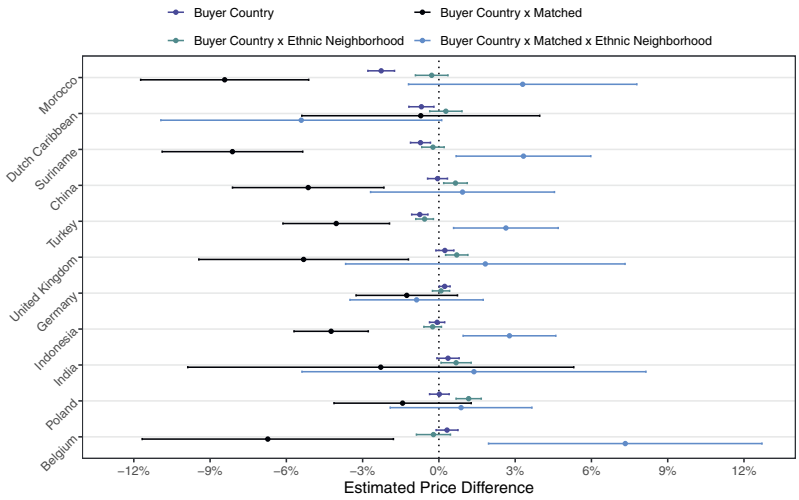


Figure 5.7: Price Discounts and Neighborhood Composition

Notes: This figure contrasts price discounts in same-ethnicity transactions with average price differentials for migrant buyers. In addition, it shows how both factors change depending on whether a purchase is made in an "ethnic" neighborhood or not. Coefficient estimates are based on Equation (5.8); error bands indicate 95% confidence intervals.

Bibliography

Aalbers, M. B. (2007). "Place-Based and Race-Based Exclusion from Mortgage Loans: Evidence from Three Cities in the Netherlands". In: *Journal of Urban Affairs* 29.1, pp. 1–29.

Adermon, A., M. Lindahl, and D. Waldenström (2018). "Intergenerational Wealth Mobility and the Role of Inheritance: Evidence from Multiple Generations". In: *Economic Journal* 128.612, F482–F513.

Agarwal, S., H.-S. Choi, J. He, and T. F. Sing (2019). "Matching in Housing Markets: The Role of Ethnic Social Networks". In: *Review of Financial Studies* 32.10, pp. 3958–4004.

Agarwal, S., Y. Deng, Q. Gu, J. He, W. Qian, and Y. Ren (2022). "Mortgage Debt, Hand-to-Mouth Households, and Monetary Policy Transmission". In: *Review of Finance* 26.3, pp. 487–520.

Ager, P., L. Boustan, and K. Eriksson (2021). "The Intergenerational Effects of a Large Wealth Shock: White Southerners After the Civil War". In: *American Economic Review* 111.11, pp. 3767–3794.

Ahlfeldt, G. M., H. R. Koster, and T. G. Vu (2025). *Mind the Lag: Using Assessed and List Prices as Proxies for Housing Market Values*. Tech. rep. CESifo Working Paper.

Aiello, D., J. D. Kotter, and G. Schubert (2024). "Housing Wealth and Overpayment: When Money Moves In". In: *Working Paper*.

Akerlof, G. A. and R. E. Kranton (2000). "Economics and Identity". In: *Quarterly Journal of Economics* 115.3, pp. 715–753.

Aladangady, A. (2017). "Housing Wealth and Consumption: Evidence from Geographically Linked Microdata". In: *American Economic Review* 107.11, pp. 3415–3446.

Aliprantis, D., D. R. Carroll, and E. R. Young (2024). "What Explains Neighborhood Sorting by Income and Race?" In: *Journal of Urban Economics* 141, p. 103508.

Allen, J., K. Carmichael, R. Clark, S. Li, and N. Vincent (2026). "Housing affordability and parental income support: The role of mortgage co-signing". In: *Journal of Financial Economics* 178, p. 104239.

Allen, J. (2006). "Welfare Regimes, Welfare Systems and Housing in Southern Europe". In: *European Journal of Housing Policy* 6.3, pp. 251–277. DOI: 10.1080/14616710600973102.

Altonji, J. G., F. Hayashi, and L. J. Kotlikoff (1997). "Parental Altruism and Inter Vivos Transfers: Theory and Evidence". In: *Journal of Political Economy* 105.6, pp. 1121–1166.

Ambrose, B. W., J. N. Conklin, and L. A. Lopez (2021). "Does Borrower and Broker Race Affect the Cost of Mortgage Credit?" In: *Review of Financial Studies* 34.2, pp. 790–826.

Andersen, A. L., N. Johannesen, and A. Sheridan (2020). "Bailing Out the Kids: New Evidence on Informal Insurance from One Billion Bank Transfers". In: *Working Paper*.

Arrington, B. T. (2012). "" Free Homes for Free Men": A Political History of the Homestead Act, 1774–1863". In.

Arundel, R. and R. Ronald (2021). "The false promise of homeownership: Homeowner societies in an era of declining access and rising inequality". In: *Urban studies* 58.6, pp. 1120–1140.

Attwood, B. (2015). *Possession: Batman's Treaty and the matter of history*. Melbourne Univ. Publishing.

Avenancio-León, C. F. and T. Howard (2022). "The Assessment Gap: Racial Inequalities in Property Taxation". In: *Quarterly Journal of Economics* 137.3, pp. 1383–1434.

Aydin, E., D. Brounen, and N. Kok (2020). "The Capitalization of Energy Efficiency: Evidence from the Housing Market". In: *Journal of Urban Economics* 117, p. 103243.

Badarinza, C., T. Ramadorai, and C. Shimizu (2022). "Gravity, Counterparties, and Foreign Investment". In: *Journal of Financial Economics* 145.2, pp. 132–152.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2018). "The Economic Effects of Social Networks: Evidence from the Housing Market". In: *Journal of Political Economy* 126.6, pp. 2224–2276.

Baker, J. (2019). "Real Property: Feudal Tenure". In: *Introduction to English Legal History*. 5th ed. Oxford: Oxford University Press. Chap. 13. DOI: 10.1093/oso/9780198812609.001.0001.

Bayer, P., M. Casey, F. Ferreira, and R. McMillan (2017). "Racial and Ethnic Price Differentials in the Housing Market". In: *Journal of Urban Economics* 102, pp. 91–105.

Bayer, P., K. Mangum, and J. W. Roberts (2021). "Speculative Fever: Investor Contagion in the Housing Bubble". In: *American Economic Review* 111.2, pp. 609–651.

Beekers, W. P. (2012). "Het bewoonbare land: Geschiedenis van de volkshuisvestingsbeweging in Nederland". In.

Benetton, M., M. Kudlyak, and J. Mondragon (2024). "Dynastic Home Equity". In: *Working Paper*.

Benmelech, E., A. Guren, and B. T. Melzer (2023). "Making the House a Home: The Stimulative Effect of Home Purchases on Consumption and Investment". In: *Review of Financial Studies* 36.1, pp. 122–154.

Berger, D., N. Turner, and E. Zwick (2020). "Stimulating Housing Markets". In: *Journal of Finance* 75.1, pp. 277–321. DOI: 10.1111/jofi.12847.

Bernstein, A., M. T. Gustafson, and R. Lewis (2019). "Disaster on the Horizon: The Price Effect of Sea Level Rise". In: *Journal of Financial Economics* 134.2, pp. 253–272.

Bernstein, A. and P. Koudijs (2024). "The Mortgage Piggy Bank: Building Wealth Through Amortization". In: *Quarterly Journal of Economics* 139.3, pp. 1767–1825. DOI: 10.1093/qje/qjae011.

Best, M. C. and H. J. Kleven (2018). "Housing Market Responses to Transaction Taxes: Evidence from Notches and Stimulus in the UK". In: *Review of Economic Studies* 85.1, pp. 157–193. DOI: 10.1093/restud/rdx032.

Biesenbeek, C., M. Mastrogiacomo, R. Alessie, and J. de Haan (2022). "The Effect of Introducing a Loan-to-Value Limit on Home Ownership". In: *De Nederlandsche Bank Working Paper*.

Black, S. E., P. J. Devereux, F. Landaud, and K. G. Salvanes (2025). "The (Un)Importance of Inheritance". In: *Journal of the European Economic Association* 23.3, pp. 1060–1094.

Black, S. E., P. J. Devereux, P. Lundborg, and K. Majlesi (2020). "Poor Little Rich Kids? The Role of Nature Versus Nurture in Wealth and Other Economic Outcomes and Behaviours". In: *Review of Economic Studies* 87.4, pp. 1683–1725.

Blickle, K. and M. Brown (2019). "Borrowing Constraints, Home Ownership and Housing Choice: Evidence from Intra-Family Wealth Transfers". In: *Journal of Money, Credit and Banking* 51.2-3, pp. 539–580.

Bloomberg (2024). *Harris to Propose $25,000 Assistance for First-Time Homeowners*. www.bloomberg.com/news/articles/2024-08-15/harris-proposes-25-000-assistance-for-stubborn-housing-prices.

Boelhouwer, P. (2017). "The Role of Government and Financial Institutions During a Housing Market Crisis: A Case Study of the Netherlands". In: *International Journal of Housing Policy* 17.4, pp. 591–602.

Boileau, B. and D. Sturrock (2023). *Who Gives and Receives Substantial Financial Transfers in Britain?* Tech. rep. IFS Working Papers.

Bos, J. W., N. Kok, and J. Wogh (2024). "Bidder Beware: Intergenerational Wealth Transfers and Overpayment in Housing Markets". In: Mimeo. DOI: 10.2139/ssrn.4737438.

Boserup, S. H., W. Kopczuk, and C. T. Kreiner (2016). *Intergenerational Wealth Formation over the Life Cycle: Evidence from Danish Wealth Records 1984–2013*. Tech. rep. Working Paper, University of Copenhagen.

— (2018). "Born with a Silver Spoon? Danish Evidence on Wealth Inequality in Childhood". In: *Economic Journal* 128.612, F514–F544.

Boustan, L. P. (2010). "Was Postwar Suburbanization "White Flight"? Evidence from the Black Migration". In: *Quarterly Journal of Economics* 125.1, pp. 417–443.

Box-Couillard, S. and P. Christensen (2024). *Racial Housing Price Differentials and Neighborhood Segregation*. Tech. rep. National Bureau of Economic Research.

Brandsaas, E. E. (2021). *Illiquid Homeownership and the Bank of Mom and Dad*. Tech. rep. Working Paper.

Burnside, C., M. Eichenbaum, and S. Rebelo (2016). "Understanding Booms and Busts in Housing Markets". In: *Journal of Political Economy* 124.4, pp. 1088–1147.

Bush, G. W. (June 2002). *Remarks by the President on Homeownership*. Speech delivered at St. Paul AME Church, Atlanta, Georgia. Office of the Press Secretary, The White House. Available at `https://georgewbush-whitehouse.archives.gov/news/releases/2002/06/20020617-2.html`.

Card, D., A. Mas, and J. Rothstein (2008). "Tipping and the Dynamics of Segregation". In: *Quarterly Journal of Economics* 123.1, pp. 177–218.

Carozzi, F., C. A. Hilber, and X. Yu (2024). "On the Economic Impacts of Mortgage Credit Expansion Policies: Evidence from Help to Buy". In: *Journal of Urban Economics* 139, p. 103611. DOI: `10.1016/j.jue.2023.103611`.

Carry, P. M., L. A. Vanderlinden, F. Dong, T. Buckner, E. Litkowski, T. Vigers, J. M. Norris, and K. Kechris (2021). "Inverse Probability Weighting Is an Effective Method to Address Selection Bias During the Analysis of High Dimensional Data". In: *Genetic Epidemiology* 45.6, pp. 593–603.

CBS (2025). *Immigratie*. URL: `https://www.cbs.nl/nl-nl/visualisaties/dashboard-bevolking/bevolkingsgroei/immigratie` (visited on 03/17/2026).

— (2026). *Hoeveel Inwoners Hebben een Herkomst Buiten Nederland*. URL: `https://www.cbs.nl/nl-nl/dossier/dossier-asiel-migratie-en-integratie/hoeveel-inwoners-hebben-een-herkomst-buiten-nederland` (visited on 03/17/2026).

Cetre, S., Y. Algan, G. Grimalda, F. Murtin, D. Pipke, L. Putterman, U. Schmidt, and V. Siegerink (2024). "Ethnic Bias, Economic Achievement and Trust Between Large Ethnic Groups: A Study in Germany and the US". In: *Journal of Economic Behavior & Organization* 224, pp. 996–1021.

Chade, H., J. Eeckhout, and L. Smith (2017). "Sorting Through Search and Matching Models in Economics". In: *Journal of Economic Literature* 55.2, pp. 493–544. DOI: 10.1257/jel.20150777.

Chambers, D. N. (1992). "The Racial Housing Price Differential and Racially Transitional Neighborhoods". In: *Journal of Urban Economics* 32.2, pp. 214–232.

Charles, K. K. and E. Hurst (2003). "The Correlation of Wealth Across Generations". In: *Journal of Political Economy* 111.6, pp. 1155–1182.

Chen, J. I., A. Foster, and L. Putterman (2019). "Identity, Trust and Altruism: An Experiment on Preferences and Microfinance Lending". In: *European Economic Review* 120, p. 103304.

Chen, Y. and S. X. Li (2009). "Group Identity and Social Preferences". In: *American Economic Review* 99.1, pp. 431–457.

Chetty, R. (2009). "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods". In: *Annual Review of Economics* 1, pp. 451–488. DOI: 10.1146/annurev.economics.050708.142910.

Christensen, P. and C. Timmins (2022). "Sorting or Steering: The Effects of Housing Discrimination on Neighborhood Choice". In: *Journal of Political Economy* 130.8, pp. 2110–2163.

Colmsjö, E. (2024). *A Flying Start: The Long-Run Effects of Inter Vivos Transfers*. Tech. rep. CEBI Working Paper Series.

Constant, A. F. and K. F. Zimmermann (2008). "Measuring Ethnic Identity and Its Impact on Economic Behavior". In: *Journal of the European Economic Association* 6.2-3, pp. 424–433.

Cooper, A. and P. Kurzer (2023). "Similar Origins–Divergent Paths: The Politics of German and Dutch Housing Markets". In: *German Politics* 32.2, pp. 341–360.

Cox, D. (1990). "Intergenerational Transfers and Liquidity Constraints". In: *Quarterly Journal of Economics* 105.1, pp. 187–217.

Cvijanović, D. and C. Spaenjers (2021). ""We'll Always Have Paris": Out-of-Country Buyers in the Housing Market". In: *Management Science* 67.7, pp. 4120–4138.

Daysal, N. M., M. F. Lovenheim, and D. N. Wasser (2023). *The Intergenerational Transmission of Housing Wealth*. Tech. rep. National Bureau of Economic Research.

Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017). "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging". In: *American Economic Review* 107.11, pp. 3550–3588.

Diefendorf, J. M. (1993). *In the wake of war: the reconstruction of German cities after World War II*. Oxford University Press.

DiPasquale, D. and E. L. Glaeser (1999). "Incentives and social capital: Are homeowners better citizens?" In: *Journal of urban Economics* 45.2, pp. 354–384.

Disney, R., J. Gathergood, S. Machin, and M. Sandi (2023). "Does homeownership reduce crime? A radical housing reform from the UK". In: *The Economic Journal* 133.655, pp. 2640–2675.

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". In: *Journal of Urban Economics* 145.103717. DOI: 10.1016/j.jue.2024.103717.

Dröes, M. I. and H. R. Koster (2016). "Renewable Energy and Negative Externalities: The Effect of Wind Turbines on House Prices". In: *Journal of Urban Economics* 96, pp. 121–141.

Eibich, P. and T. Siedler (2020). "Retirement, Intergenerational Time Transfers, and Fertility". In: *European Economic Review* 124, p. 103392.

Elinder, M., O. Erixson, and D. Waldenström (2018). "Inheritance and Wealth Inequality: Evidence from Population Registers". In: *Journal of Public Economics* 165, pp. 17–30.

Elsinga, M., H. Priemus, and P. Boelhouwer (2016). "Milestones in housing finance in the Netherlands, 1988–2013". In: *Milestones in European housing finance*, pp. 255–272.

Engelhardt, G. V. and C. J. Mayer (1998). "Intergenerational Transfers, Borrowing Constraints, and Saving Behavior: Evidence from the Housing Market". In: *Journal of Urban Economics* 44.1, pp. 135–157.

Escobar, S., H. Ohlsson, and H. Selin (2023). "Giving to the Children or the Taxman? Lessons from a Swedish Inheritance Tax Loophole". In: *European Economic Review* 153, p. 104382.

Fagereng, A., M. Mogstad, and M. Rønning (2021). "Why Do Wealthy Parents Have Wealthy Children?" In: *Journal of Political Economy* 129.3, pp. 703–756.

Fan, Y. and A. Yavas (2020). "How Does Mortgage Debt Affect Household Consumption? Micro Evidence from China". In: *Real Estate Economics* 48.1, pp. 43–88.

Favilukis, J., P. Mabille, and S. Van Nieuwerburgh (2023). "Affordable Housing and City Welfare". In: *Review of Economic Studies* 90.1, pp. 293–330.

Fehr, E. and K. M. Schmidt (1999). "A Theory of Fairness, Competition, and Cooperation". In: *Quarterly Journal of Economics* 114.3, pp. 817–868.

Fershtman, C. and U. Gneezy (2001). "Discrimination in a Segmented Society: An Experimental Approach". In: *Quarterly Journal of Economics* 116.1, pp. 351–377.

Fetter, D. K. (2013). "How Do Mortgage Subsidies Affect Home Ownership? Evidence from the Mid-Century GI Bills". In: *American Economic Journal: Economic Policy* 5.2, pp. 111–147. DOI: 10.1257/pol.5.2.111.

Fisman, R., D. Paravisini, and V. Vig (2017). "Cultural Proximity and Loan Outcomes". In: *American Economic Review* 107.2, pp. 457–492.

Floetotto, M., M. Kirker, and J. Stroebel (2016). "Government Intervention in the Housing Market: Who Wins, Who Loses?" In: *Journal of Monetary Economics* 80, pp. 106–123. DOI: 10.1016/j.jmoneco.2016.04.005.

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. DOI: 10.18128/D030.V13.0.

Fong, C. M. and E. F. P. Luttmer (2009). "What Determines Giving to Hurricane Katrina Victims? Experimental Evidence on Racial Group Loyalty". In: *American Economic Journal: Applied Economics* 1.2, pp. 64–87.

Francke, M., L. Hans, M. Korevaar, and S. van Bekkum (2024). "Buy-to-Live vs. Buy-to-Let: Home Ownership, Property Prices and Neighborhoods". In: *Working Paper*.

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". In: Mimeo. DOI: 10.2139/ssrn.4480261.

Gale, W. G. and J. K. Scholz (1994). "Intergenerational Transfers and the Accumulation of Wealth". In: *Journal of Economic Perspectives* 8.4, pp. 145–160.

Ganong, P. and P. Noel (2020). "Liquidity Versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession". In: *American Economic Review* 110.10, pp. 3100–3138.

Glaeser, E. L., D. I. Laibson, J. A. Scheinkman, and C. L. Soutter (2000). "Measuring Trust". In: *Quarterly Journal of Economics* 115.3, pp. 811–846.

Glaeser, E. L. and J. D. Shapiro (2003). "The Benefits of the Home Mortgage Interest Deduction". In: *Tax Policy and the Economy* 17, pp. 37–82. DOI: 10.1086/tpe.17.20140504.

Gobillon, L. and D. Le Blanc (2008). "Economic Effects of Upfront Subsidies to Ownership: The Case of the Prêt à Taux Zéro in France". In: *Journal of Housing Economics* 17.1, pp. 1–33. DOI: 10.1016/j.jhe.2007.06.003.

Goldsmith-Pinkham, P. and K. Shue (2023). "The Gender Gap in Housing Returns". In: *Journal of Finance* 78.2, pp. 1097–1145.

Goodwin, K. and L. Zumpano (2011). "The Home Buyer Tax Credit of 2009 and the Transition to Homeownership". In: *Journal of Housing Research* 20.2, pp. 211–224. DOI: 10.1080/10835547.2011.12092040.

Gruber, J., A. Jensen, and H. Kleven (2021). "Do People Respond to the Mortgage Interest Deduction? Quasi-Experimental Evidence from Denmark". In: *American Economic Journal: Economic Policy* 13.2, pp. 273–303. DOI: 10.1257/pol.20170366.

Guiso, L. and T. Jappelli (2002). "Private Transfers, Borrowing Constraints and the Timing of Homeownership". In: *Journal of Money, Credit and Banking*, pp. 315–339.

Guiso, L., P. Sapienza, and L. Zingales (2009). "Cultural Biases in Economic Exchange?" In: *Quarterly Journal of Economics* 124.3, pp. 1095–1131.

Gupta, A. (2019). "Foreclosure Contagion and the Neighborhood Spillover Effects of Mortgage Defaults". In: *Journal of Finance* 74.5, pp. 2249–2301.

Habakkuk, H. J. (1940). "English Landownership, 1680-1740". In: *The Economic History Review* 10.1, pp. 2–17.

Hacamo, I. (2024). "Interest Rates and the Distribution of House Prices". In: *Review of Economic Studies*. Accepted. DOI: 10.2139/ssrn.2687277.

Haerpfer, C., R. Inglehart, A. Moreno, C. Welzel, K. Kizilova, J. Diez-Medrano, M. Lagos, P. Norris, E. Ponarin, and B. Puranen (2020). *World Values Survey: Round Seven – Country-Pooled Datafile*. Madrid, Spain & Vienna, Austria: JD Systems Institute & WVSA Secretariat. DOI: 10.14281/18241.1.

Harding, J. P., S. S. Rosenthal, and C. F. Sirmans (2003). "Estimating Bargaining Power in the Market for Existing Homes". In: *Review of Economics and Statistics* 85.1, pp. 178–188.

Haurin, D. R., T. L. Parcel, and R. J. Haurin (2002). "Does homeownership affect child outcomes?" In: *Real Estate Economics* 30.4, pp. 635–666.

Hausman, N., T. Ramot-Nyska, and N. Zussman (2022). "Homeownership, Labor Supply, and Neighborhood Quality". In: *American Economic Journal: Economic Policy* 14.2, pp. 193–230. DOI: 10.1257/pol.20200177.

Hembre, E. (2018). "An Examination of the First-Time Homebuyer Tax Credit". In: *Regional Science and Urban Economics* 73, pp. 196–216. DOI: 10.1016/j.regsciurbeco.2018.10.002.

Hilber, C. A. and T. M. Turner (2014). "The Mortgage Interest Deduction and Its Impact on Homeownership Decisions". In: *Review of Economics and Statistics* 96.4, pp. 618–637. DOI: 10.1162/REST_a_00427.

Himmelberg, C., C. Mayer, and T. Sinai (2005). "Assessing high house prices: Bubbles, fundamentals and misperceptions". In: *Journal of Economic Perspectives* 19.4, pp. 67–92.

House of Commons (Jan. 1980). *Housing Bill*. House of Commons Debate, Hansard. Accessed: September 19, 2026. URL: https://hansard.parliament.uk/commons/1980-01-15/debates/0d61c425-da31-47ba-af78-4b446c4a7227/HousingBill.

Iacus, S. M., G. King, and G. Porro (2012). "Causal Inference Without Balance Checking: Coarsened Exact Matching". In: *Political Analysis* 20.1, pp. 1–24.

Ihlanfeldt, K. and T. Mayock (2009). "Price Discrimination in the Housing Market". In: *Journal of Urban Economics* 66.2, pp. 125–140.

Jordà, Ò., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019). "The Rate of Return on Everything, 1870–2015". In: *Quarterly Journal of Economics* 134.3, pp. 1225–1298.

Kermani, A. and F. Wong (2026). "Racial disparities in housing returns". In: *American Economic Review* 116.1, pp. 287–331.

Kiel, K. A. and J. E. Zabel (1996). "House Price Differentials in US Cities: Household and Neighborhood Racial Effects". In: *Journal of Housing Economics* 5.2, pp. 143–165.

Kleven, H. (2021). "Sufficient Statistics Revisited". In: *Annual Review of Economics* 13 (1), pp. 515–538. DOI: 10.1146/annurev-economics-060220-023547.

Kohl, S. (2014). "Homeowner nations or nations of tenants: how historical institutions in urban politics, housing finance and construction set Germany, France and the US on different housing paths". PhD thesis. University of Cologne Cologne.

— (2017). *Homeownership, Renting and Society: Historical and Comparative Perspectives*. Vol. 212. Routledge Advances in Sociology. London: Routledge.

— (2018). "A small history of the homeownership ideal". In.

Kopczuk, W. (2007). "Bequest and Tax Planning: Evidence from Estate Tax Returns". In: *Quarterly Journal of Economics* 122.4, pp. 1801–1854.

Kopczuk, W. and D. Munroe (2015). "Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate Market". In: *American Economic Journal: Economic Policy* 7.2, pp. 214–257. DOI: 10.1257/pol.20130361.

Koster, H. R. and J. Rouwendal (2024). "Housing Market Discount Rates: Evidence from Bargaining and Bidding Wars". In: *International Economic Review* 65.2, pp. 955–1002.

Kudashvili, N. and P. Lergetporer (2022). "Minorities' Strategic Response to Discrimination: Experimental Evidence". In: *Journal of Public Economics* 208, p. 104630.

Kvaerner, J. S. (2023). "How Large Are Bequest Motives? Estimates Based on Health Shocks". In: *Review of Financial Studies* 36.8, pp. 3382–3422.

Landen Mammos, N. (2025). "The Bank of Mom and Dad: Intergenerational Transfers and Macroprudential Regulation". In: *Working Paper*.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015). "The Housing Market(s) of San Diego". In: *American Economic Review* 105.4, pp. 1371–1407. DOI: 10.1257/aer.20111662.

Lane, T. (2016). "Discrimination in the Laboratory: A Meta-Analysis of Economics Experiments". In: *European Economic Review* 90, pp. 375–402.

Lee, H., D. Myers, G. Painter, J. Thunell, and J. Zissimopoulos (2020). "The Role of Parental Financial Assistance in the Transition to Homeownership by Young Adults". In: *Journal of Housing Economics* 47, p. 101597.

Leib, M., N. C. Köbis, M. Francke, S. Shalvi, and M. Roskes (2021). "Precision in a Seller's Market: Round Asking Prices Lead to Higher Counteroffers and Selling Prices". In: *Management Science* 67.2, pp. 1048–1055.

Leopold, T. and T. Schneider (2011). "Family Events and the Timing of Intergenerational Transfers". In: *Social Forces* 90.2, pp. 595–616.

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

Locke, J. (1690). *Two Treatises of Government*. London: Awnsham Churchill.

Määttänen, N. and M. Terviö (2014). "Income Distribution and Housing Prices: An Assignment Model Approach". In: *Journal of Economic Theory* 151, pp. 381–410. DOI: 10.1016/j.jet.2014.01.003.

— (2022). "Welfare Effects of Housing Transaction Taxes: A Quantitative Analysis with an Assignment Model". In: *Economic Journal* 132 (644), pp. 1566–1599. DOI: 10.1093/ej/ueab080.

Mantilla, C., L. Zhou, C. Wang, D. Yang, S. Shen, and P. Seabright (2021). "Favoring Your In-Group Can Harm Both Them and You: Ethnicity and Public Goods Provision in China". In: *Journal of Economic Behavior & Organization* 185, pp. 211–233.

McCartney, W. B. and A. M. Shah (2022). "Household Mortgage Refinancing Decisions Are Neighbor Influenced, Especially Along Racial Lines". In: *Journal of Urban Economics* 128, p. 103409.

McGarry, K. (1999). "Inter Vivos Transfers and Intended Bequests". In: *Journal of Public Economics* 73.3, pp. 321–351.

— (2016). "Dynamic Aspects of Family Transfers". In: *Journal of Public Economics* 137, pp. 1–13.

Milgrom, P. and I. Segal (2002). "Envelope Theorems for Arbitrary Choice Sets". In: *Econometrica* 70.2, pp. 583–601. DOI: 10.1111/1468-0262.00296.

Miller, D. L. (2023). "An Introductory Guide to Event Study Models". In: *Journal of Economic Perspectives* 37.2, pp. 203–230.

Myers, C. K. (2004). "Discrimination and Neighborhood Effects: Understanding Racial Differentials in US Housing Prices". In: *Journal of Urban Economics* 56.2, pp. 279–302.

NARvis (2021). *Aboriginal Custodianship – Background Information*. https://narvis.com.au/the-region/aboriginal-custodianship-background-information/. Natural Resource Management Vision for the Northern Agricultural Region of Western Australia. Accessed: 24 March 2026.

Nekoei, A. and D. Seim (2023). "How Do Inheritances Shape Wealth Inequality? Theory and Evidence from Sweden". In: *Review of Economic Studies* 90.1, pp. 463–498.

Ockenfels, A. and P. Werner (2014). "Beliefs and Ingroup Favoritism". In: *Journal of Economic Behavior & Organization* 108, pp. 453–462.

Ondrich, J., S. Ross, and J. Yinger (2003). "Now You See It, Now You Don't: Why Do Real Estate Agents Withhold Available Houses from Black Customers?" In: *Review of Economics and Statistics* 85.4, pp. 854–873.

Oswald, A. (1999). "The housing market and Europe's unemployment: a non-technical paper". In: *Homeownership and the labour Market in Europe*.

Palomino, J. C., G. A. Marrero, B. Nolan, and J. G. Rodríguez (2022). "Wealth Inequality, Intergenerational Transfers, and Family Background". In: *Oxford Economic Papers* 74.3, pp. 643–670.

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". In: *American Economic Review* 98.2, pp. 84–89. DOI: 10.1257/aer.98.2.84.

Redfin (2024). *Nepo-Homebuyers: More Than One-Third of Gen Z and Millennial Homebuyers Plan to Use Family Money for Down Payment*. www.redfin.com/news/gen-z-millennial-down-payment-family-help/.

Reher, M. and R. Valkanov (2024). "The Mortgage-Cash Premium Puzzle". In: *Journal of Finance* 79.5, pp. 3149–3201.

Rohe, W. M., S. Van Zandt, and G. McCarthy (2002). "Home ownership and access to opportunity". In: *Housing Studies* 17.1, pp. 51–61.

Ronald, R. and K. Dol (2011). "Housing in the Netherlands before and after the global financial crisis". In: *Housing markets and the global financial crisis*. Edward Elgar Publishing.

Rosen, S. (1974). "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition". In: *Journal of Political Economy* 82.1, pp. 34–55.

Rouwendal, J. and A. Petrat (2022). "Mortgage Underwriting and House Prices: Evidence from the 2011 Change in the Dutch Code of Conduct for Mortgage Loans". In: *Real Estate Economics* 50.4, pp. 1141–1159.

Saez, E. and S. Stantcheva (2016). "Generalized Social Marginal Welfare Weights for Optimal Tax Theory". In: *American Economic Review* 106.1, pp. 24–45. DOI: 10.1257/aer.20141362.

Schulz, J. F., D. Bahrami-Rad, J. P. Beauchamp, and J. Henrich (2019). "The Church, Intensive Kinship, and Global Psychological Variation". In: *Science* 366.6466, eaau5141.

Slemrod, J., C. Weber, and H. Shan (2017). "The Behavioral Response to Housing Transfer Taxes: Evidence from a Notched Change in DC Policy". In: *Journal of Urban Economics* 100, pp. 137–153. DOI: 10.1016/j.jue.2017.05.005.

Solinge, F. van and C. Fiedler (2020). "Evaluatie Verlaging Overdrachtsbelasting 2011–2012". In: *CPB Report*.

Sommer, E. (2017). *Wealth Transfers and Tax Planning: Evidence for the German Bequest Tax*. Tech. rep. IZA Discussion Papers.

Statistics Netherlands (2019). *Schenkingen Vooral van Ouders aan Vermogende Kinderen*. https://www.cbs.nl/nl-nl/nieuws/2019/11/schenkingen-vooral-van-ouders-aan-vermogende-kinderen.

— (2021). *Minder Geschonken Onder Verruimde Vrijstelling voor de Eigen Woning*. www.cbs.nl/nl-nl/nieuws/2021/09/minder-geschonken-onder-verruimde-vrijstelling-voor-de-eigen-woning.

Tajfel, H. and J. C. Turner (1979). "An Integrative Theory of Intergroup Conflict". In: *The Social Psychology of Intergroup Relations*. Ed. by W. G. Austin and S. Worchel. Monterey, CA: Brooks/Cole, pp. 33–37.

Thaler, R. H. and C. R. Sunstein (2021). *Nudge: The Final Edition*. Penguin.

The Economist (June 2025). "How to Invest Your Enormous Inheritance". In: *The Economist*. Accessed: 2025-09-12. URL: https://www.economist.

com/finance-and-economics/2025/06/12/how-to-invest-your-enormous-inheritance.

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.

Van Deursen, H. (2023). "The People's Housing: Woningcorporaties and the Dutch Social Housing System-Part 1: The History". In: *Jt Cent Hous Stdies Harv Univ*.

Van Dijk, W. (2019). "The Socio-Economic Consequences of Housing Assistance". In: *University of Chicago Kenneth C. Griffin Department of Economics Job Market Paper* 36.

Van Kempen, R. and H. Priemus (2002). "Revolution in social housing in the Netherlands: Possible effects of new housing policies". In: *Urban Studies* 39.2, pp. 237–253.

Vangeel, W., L. Defau, and L. De Moor (2023). "Young households' diminishing access to homeownership attainment in Europe". In: *Sustainability* 15.8, p. 6906.

Waldfogel, J. (2008). "The Median Voter and the Median Consumer: Local Private Goods and Population Composition". In: *Journal of Urban Economics* 63.2, pp. 567–582.

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.

West, E. G. (2003). "Property rights in the history of economic thought". In: *Property rights: Cooperation, conflict, and law* 20.

Wold, E. G., K. A. Aastveit, E. E. Brandsaas, R. E. Juelsrud, and G. J. Natvik (2024). *The Housing Channel of Intergenerational Wealth Persistence*. 16/2023. Working Paper.

Wong, M. (2013). "Estimating Ethnic Preferences Using Ethnic Housing Quotas in Singapore". In: *Review of Economic Studies* 80.3, pp. 1178–1214.

Wooldridge, J. M. (2002). "Inverse Probability Weighted M-Estimators for Sample Selection, Attrition, and Stratification". In: *Portuguese Economic Journal* 1.2, pp. 117–139.

Yang, S. and M. Ripoll (2023). "Financial Transfers from Parents to Adult Children". In: *Journal of Economic Behavior & Organization* 208, pp. 286–303.

Zorlu, A., C. H. Mulder, and R. Van Gaalen (2014). "Ethnic Disparities in the Transition to Home Ownership". In: *Journal of Housing Economics* 26, pp. 151–163.

Summary

In many countries, homeownership is seen as the desirable tenure form both from an individual and from a societal perspective. This dissertation examines various aspects that influence households' ability to become homeowners beyond their own financial resources.

Chapter 2 studies the effectiveness of policies that support young homebuyers financially. Relying on the Dutch transfer tax exemption for homebuyers under 35, we assess the impact of such policies on various individual- and market-level outcomes. Most importantly, we find that the exemption led to a significant increase in home purchases by eligible buyers, which did not, however, translate into a lasting ownership effect. This is because additional transactions generated by the policy mostly reflected repeat purchases by existing owners, as well as timing responses by individuals who were already planning their first purchase. In addition, most of the incidence of the subsidy fell on sellers, rather than buyers, since between 60% and 90% of the tax exemption was captured by increased prices for eligible units. We combine all of these empirical findings in a simple welfare analysis, calculating the utility consequences of the policy for every member of the Dutch adult population. While the tax exemption mostly benefited those in the upper-middle part of the income distribution, those with low incomes lost out, as access to homeownership continued to be out of reach for them even with a tax incentive.

Chapter 3 studies the influence of parental cash transfers on young first-time homebuyers. Concretely, it investigates how homebuyers who receive parental support use the funds in their immediate housing decisions, as well as in their longer-run consumption and savings decisions. I find that during the home purchase, (close to) every euro of parental transfer is used to deleverage, by increasing the size of the initial down payment. This mechanically lowers the monthly costs that homebuyers face due to mortgage interest payments. However, rather than investing these savings elsewhere, transfer recipients use their higher residual income for consumption. The results therefore suggest that targeted parental wealth transfers are not only used to accumulate long-run wealth, but also to help the recipients smooth consumption in the middle of the life cycle.

Chapter 4 studies the broader housing market implications of parental cash

transfers. During a period of special tax incentives for cash transfers, we compare the bidding behavior of all homebuyers who do and do not receive parental support to fund their purchase. At the individual level, we find that transfer recipients make between 0.5% and 2% higher bids for homes of the same quality, compared to other buyers. Further analysis indicates that overbidding by transfer recipients occurs in those local markets where competition is tighter. At the market level, overbidding has a spillover effect: in postcode areas where many homebuyers receive a parental transfer, we find that other buyers start to make higher bids as well, driving up average prices.

Chapter 5 studies the role of ethnicity, which has been shown to constrain access to homeownership for prospective buyers belonging to ethnic minorities. Using information on the origin country of all homebuyers and homesellers in the Dutch residential housing market, we study the interaction between buyer and seller ethnicity. We find that, even within small sub-markets, transactions between buyers and sellers of the same migrant group occur more frequently than predicted by random matching. This indicates preferences for ethnic matching in housing transactions. Crucially, ethnic matching preferences have an effect on sale prices. Relative to transactions involving other buyer-seller combinations, those involving same-ethnicity buyers and sellers sell at a discount of around 4%. This discount is not explained by differences in house and neighborhood quality. Excluding transactions that likely did not occur at arm's length—such as those between family members and those involving buyers and sellers living in the same neighborhood—reduces the estimated discount to 2%. Various additional sub-sample analyses provide evidence that the discount is driven by ethnic in-group preferences, which induce sellers to forgo a financial return when negotiating with same-ethnicity buyers.

Samenvatting

In veel landen wordt het eigenwoningbezit zowel vanuit individueel als vanuit maatschappelijk perspectief gezien als de wenselijke eigendomsvorm. Dit proefschrift onderzoekt verschillende aspecten die, naast de eigen financiële middelen van huishoudens, van invloed zijn op hun mogelijkheden om een eigen woning te verwerven.

Hoofdstuk 2 onderzoekt de effectiviteit van beleid dat jonge woningkopers financieel ondersteunt. Aan de hand van de Nederlandse vrijstelling van overdrachtsbelasting voor kopers jonger dan 35 jaar bepalen we het effect van dergelijk beleid op verschillende uitkomsten op individueel en op marktniveau. Het belangrijkste resultaat is dat de vrijstelling leidde tot een significante toename van het aantal woningaankopen door kopers die ervoor in aanmerking kwamen, wat zich echter niet vertaalde in een blijvend effect op het eigenwoningbezit. Dit komt doordat de extra transacties die door het beleid werden gegenereerd grotendeels bestonden uit herhaalaankopen door bestaande eigenaren, en uit timingreacties van personen die hun eerste aankoop toch al van plan waren. Daarnaast sloeg de incidentie van de subsidie vooral neer bij verkopers in plaats van bij kopers, aangezien tussen de 60 en 90% van de vrijstelling werd afgeroomd door hogere prijzen voor woningen die onder de regeling vielen. We combineren al deze empirische bevindingen in een eenvoudige welvaartsanalyse, waarin we de nutseffecten van het beleid berekenen voor elk lid van de Nederlandse volwassen bevolking. Hoewel waar de vrijstelling vooral ten goede kwam aan degenen in het hogere middensegment van de inkomensverdeling, waren degenen met lage inkomens de verliezers, omdat het eigenwoningbezit voor hen ook met een fiscale stimulans buiten bereik bleef.

Hoofdstuk 3 onderzoekt de invloed van ouderlijke schenkingen op jonge starters. Concreet wordt onderzocht hoe kopers die ouderlijke steun ontvangen deze middelen inzetten in hun directe woonbeslissingen, en in hun consumptie- en spaarbeslissingen op de langere termijn. Ik vind dat bij de woningaankoop (vrijwel) elke euro aan ouderlijke schenking wordt gebruikt om de schuld te verlagen, door de omvang van de initiële aanbetaling te verhogen. Dit verlaagt mechanisch de maandlasten die kopers dragen als gevolg van hypotheekrentebetalingen. In plaats van deze besparing elders te beleggen, gebruiken ont-

vangers van een schenking hun hogere restinkomen echter voor consumptie. De resultaten suggereren daarmee dat gerichte ouderlijke vermogensoverdrachten niet zozeer worden gebruikt om op de lange termijn vermogen op te bouwen, maar vooral om de ontvangers te helpen hun consumptie in het midden van de levenscyclus glad te strijken.

Hoofdstuk 4 onderzoekt de bredere gevolgen van ouderlijke schenkingen voor de woningmarkt. Voor een periode met bijzondere fiscale stimulansen voor schenkingen vergelijken we het biedgedrag van alle woningkopers die wel en niet ouderlijke steun ontvangen om hun aankoop te financieren. Op individueel niveau vinden we dat ontvangers van een schenking tussen de 0,5% en 2% hoger bieden op woningen van dezelfde kwaliteit dan andere kopers. Nadere analyse laat zien dat overbieden door ontvangers van een schenking zich voordoet in die lokale markten waar de concurrentie krapper is. Op marktniveau heeft overbieden een spillover-effect: in postcodegebieden waar veel woningkopers een ouderlijke schenking ontvangen, zien we dat ook andere kopers hoger gaan bieden, waardoor de gemiddelde prijzen stijgen.

Hoofdstuk 5 onderzoekt de rol van etniciteit. Daarvan is aangetoond dat zij de toegang tot een koopwoning beperkt voor kopers met een etnische minderheidsachtergrond. Aan de hand van informatie over het herkomstland van alle kopers en verkopers op de Nederlandse koopwoningmarkt onderzoeken we de wisselwerking tussen de etniciteit van koper en verkoper. We vinden dat, zelfs binnen kleine deelmarkten, transacties tussen kopers en verkopers uit dezelfde migrantengroep vaker voorkomen dan op grond van willekeurige matching te verwachten valt. Dit wijst op voorkeuren voor etnische matching bij woningtransacties. Van belang is dat deze voorkeuren ook effect hebben op de verkoopprijzen. Ten opzichte van transacties met andere koper-verkopercombinaties worden woningen bij transacties tussen kopers en verkopers van dezelfde etniciteit verkocht met een korting van ongeveer 4%. Deze korting wordt niet verklaard door verschillen in woning- en buurtkwaliteit. Wanneer transacties worden uitgesloten die waarschijnlijk niet zakelijk tot stand zijn gekomen—zoals transacties tussen familieleden en transacties waarbij koper en verkoper in dezelfde buurt wonen—daalt de geschatte korting tot 2%. Verschillende aanvullende analyses op deelsteekproeven leveren bewijs dat de korting wordt gedreven door voorkeuren voor de eigen etnische groep, die verkopers ertoe brengen af te zien van financieel rendement bij onderhandelingen met kopers van dezelfde etniciteit.

Impact

The own home is by far the largest asset in the balance sheet of the average household in most countries. This suggests that access to ownership is of first-order importance for wealth accumulation. Meanwhile, subsidies for homeownership are among the most significant items on many governments' budget sheets. In 2026, the Dutch government lost around €11 billion in tax revenue due to mortgage interest deductions from income tax alone, a sum that exceeds the total budget of the country's entire Ministry of Housing and Spatial Planning.

For these reasons, it is important to ensure equitable access to homeownership across the population. This dissertation contributes to the discussion around this issue, by considering various external factors which may influence a household's chances of accessing ownership. In chapter 2, we investigate a targeted subsidy for young homebuyers, among whom ownership rates have been declining substantially in recent decades. We study the exemption of the transfer tax for buyers under 35, which, while formally not a subsidy, reflects the deliberate political decision to redistribute public funds from one group to another. Our finding that the tax exemption does not increase lasting ownership rates among the targeted population should give pause to policymakers planning similar schemes in the future. In addition, the simulation analysis we carry out offers a simple way to think about the welfare consequences of targeted subsidies. Rather than creating a sizable welfare gain across the population, the transfer tax exemption—and likely other comparable policies too—redistributes resources from those at the lower end of the income distribution to those in the upper half. At the very least, governments should offer more transparency about these trade-offs when discussing similar policies.

In chapters 3 and 4, we consider the role of parental wealth in facilitating access to ownership. This has gained increased public attention in recent years, as house prices in many tight markets now extend far beyond the borrowing capacity of the average household. Importantly, debates around parental influence in the housing market, like all debates concerning the intergenerational transfer of wealth, are normative in nature. Side-stepping this distributional discussion as much as possible, the chapters in this dissertation instead document the effects that parental transfers for homebuyers have on the recipients

and on the market. The findings of chapter 3 indicate that homebuyers benefit from parental assistance not only because it allows for an earlier purchase. By investing the transfer into reducing mortgage debt, the recipients effectively annuitize their parents' wealth, turning it from a lump-sum amount to a monthly benefit. From a life-cycle perspective, this has key benefits, as it allows for smoother consumption exactly in those years when households are typically most constrained in their budget.

Chapter 4 extends beyond the immediate recipient and asks how parental transfers affect other actors in the market. The finding that buyers who receive a transfer increase their bid for a given home indicates that at least a small fraction of the benefit is captured by sellers. Our analysis suggests, however, that the effect of parental transfers extends to other buyers in the market, who respond by increasing their bids even though their budgets are unchanged. While a more formal analysis of the welfare effects of such spillovers is missing, this suggests that the impacts of bilateral wealth transfers are far-reaching. This should be borne in mind by policymakers, who have indirect control over the inflow of parental wealth to the housing market, not only by taxing inter vivos transfers, but also by regulating other forms of intergenerational support such as family mortgages or mortgage co-signing.

Finally, chapter 5 studies the role of ethnicity, which has deservedly received wide attention in the academic debate around homeownership. We add more nuance to this debate, by studying a mechanism that has not been identified previously in housing market transactions: preferences for selling to ethnic in-group members. The implications of these preferences are twofold. At the individual level, migrant homebuyers may be behaving rationally by self-selecting into neighborhoods where they are likely to meet an in-group seller. Not only are ethnic neighborhoods often desirable for migrant buyers due to homophily and shared amenities, but purchasing from an in-group seller may offer a direct financial benefit due to the discount on the purchase price. At the societal level, however, ethnic in-group bias may contribute to neighborhood segregation, with possible disadvantages due to reduced interactions across members of different ethnic groups.

About the author

Jonas Wogh was born on April 6, 1995, in Munich. He has earned a Bachelor's degree in Economics from LMU Munich, and a Master's degree in Economic and Financial Research from Maastricht University. In 2021, he began a PhD program at the Finance department of Maastricht University.

In his research, he uses administrative data to study questions related to household finance and housing markets. He typically employs quasi-experimental methods to identify causal relationships in household behavior or to evaluate the effects of government policies.

He has presented his research at international conferences, including the Allied Social Science Association (ASSA) annual meeting, the annual conference of the Royal Economic Society (RES), and the European meeting of the Urban Economics Association (UEA). Moreover, he gave seminar presentations at various universities, including Cambridge, Tilburg and Konstanz. In the fall of 2024, he conducted a research visit at the University of Colorado at Boulder.

In 2026, Jonas began a postdoctoral fellowship at Tilburg University and the University of Colorado, where he will be part of the Center for Research on Consumer Financial Decision Making (CFDM).

