UNIVERSITY OF CALIFORNIA

Los Angeles

Essays on the Role of Housing in Household Finance and its Macroeconomic Consequences

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy in Economics

by

Jesper Boejeryd

2024

ABSTRACT OF THE DISSERTATION

Essays on the Role of Housing in Household Finance and its Macroeconomic Consequences

by

Jesper Boejeryd
Doctor of Philosophy in Economics
University of California, Los Angeles, 2024
Professor Lee Ohanian, Chair

This dissertation comprises three chapters that explore the role housing plays in household decision-making and its macroeconomic implications, aiming to improve our understanding of how macroeconomic shocks are amplified or muted due to financial frictions and the special properties of housing through the household side of the economy.

The first chapter investigates how households adjust their car spending in response to a housing wealth shock. Utilizing detailed Swedish administrative data, my co-authors and I study the economic consequences of the unexpected decision to continue operating *Bromma Airport* in Stockholm, Sweden, in 2007. The airport's continuation changed perceptions of housing values in areas exposed to its negative externalities, such as noise and accident risk, causing persistent and heterogeneous changes in housing wealth across households. By using transaction data on homes across differently exposed neighborhoods, we estimated household-level housing wealth losses. We also observe car purchases for all households and can match car purchases to housing wealth losses and other household characteristics which allows us to quantify the effect on car spending in a difference-in-difference setting. The key finding is that households exhibit a muted response in car purchases compared to prior studies. Using loan-level data, we show that the effect is concentrated among homeowners who rely more on mortgage borrowing to finance car purchases; as home prices unexpectedly change, their borrowing ability is impacted, suggesting a significant role for the collateral

channel. By exploiting other dimensions of household heterogeneity, we also conclude that the pure wealth effect is weak.

The second chapter addresses the puzzle of low migration rates from areas experiencing economic decline. Focusing on the labor-market area around Stavanger, Norway, following the 2014 global oil price drop, I empirically document that diminishing home values are associated with homeowners reducing their probability to leave, while renters and homeowners with the least housing wealth exhibit an increase in their leaving probabilities. On net, the out-migration response is economically insignificant while the in-migration response exhibits a strong reduction—the net-migration falls because people stop moving to Stavanger. A life-cycle model with location, housing, and saving decisions explains these results: The reduction in home prices decreases the affordability of housing in potential destinations, thus making migration less attractive—a "housing wealth effect." This finding challenges the prevailing notion that cheaper local housing encourages workers to stay—i.e., benefits stayers—and instead suggests that falling home values act as a barrier to mobility. From a general equilibrium perspective, prices have to fall to clear the housing market. As prices fall, it becomes less attractive for homeowners to move. Potential immigrants are deterred by the fall in potential earnings, and current homeowners end up holding the stock of housing. While effective policy transfers welfare from those impacted by shocks, a policy experiment with moving vouchers shows that the beneficiaries are largely renters, who are already compensated by cheaper rents but still act as if they face lower moving costs.

The third chapter complements the analysis of the first, using the same empirical setting as Chapter 2 but narrowing the sample to government workers who did not suffer differently from the oil shock as a function of their location. However, changes in home prices accompanying the oil shock differed greatly depending on the location's exposure to the petroleum sector. Using Norwegian administrative data matched to high-frequency digital payment data, we document the dynamic response of expenditures by various goods and services as functions of lost housing wealth at the household level. The analysis provides a more detailed understanding of housing wealth shocks' spending dynamics and heterogeneity than previous work. We find that overall spending is affected less than in previous work, with a marginal propensity for expenditures close to 0.02 *kroner* per *krone* housing wealth change. Vehicles and furnishings—durable goods often financed using credit and whose purchases can be postponed—respond more strongly, while food and beverages do not change when contrasting between government workers in Stavanger versus the rest of Norway. We also find that more indebted households and those with less liquid assets respond about twice as strongly than the average. Our findings indicate that the potential spillovers of reduced consumption due to home price shocks can be limited due to the nature of goods more affected, and that the distribution of debt in the economy matters for the aggregate effect.

Collectively, these essays contribute to the fields of household finance and macroeconomics by leveraging comprehensive Scandinavian administrative data and quasi-experiments, offering insights that add to previous studies. The findings underscore the complexity of housing wealth effects on consumption and migration, with policy implications for addressing economic decline and household financial stability.

The dissertation of Jesper Boejeryd is approved.

Andrew Granger Atkeson

Pierre-Olivier Weill

Kyle F. Herkenhoff

Lee Ohanian, Committee Chair

University of California, Los Angeles

2024

I dedicate this work to my greatest champions:

Chandni; Mum and Dad; and Josefine, Julia, Jenny, and Johanna.

TABLE OF CONTENTS

List of Figures . . . xi

List of Tables . . . xiii

Acknowledgments . . . xv

Vita . . . xvii

1 The housing wealth effect: Quasi-experimental evidence . . . 1

1.1 The quasi-experiment . . . 6

1.1.1 History and political governance . . . 6

1.1.2 Data on transactions of houses and apartments . . . 9

1.1.3 Empirical strategy for measuring the house-price shock . . . 11

1.1.4 Estimates of house-price effects . . . 11

1.2 Household data . . . 13

1.2.1 Credit financing of car purchases . . . 17

1.3 Results . . . 20

1.3.1 Reduced form . . . 20

1.3.2 Two-sample instrumental variable estimates . . . 23

1.3.3 The marginal propensity of expenditures (MPX) . . . 25

1.3.4 Robustness . . . 26

1.4 Model . . . 27

1.4.1 Overview . . . 27

1.4.2 Design of the quasi-experiment . . . 29

1.4.3 Consumption response . . . 30

1.4.4 Understanding variation across households . . . 32

1.4.5 Understanding variation across studies . . . 33

1.5 Concluding remarks . . . 37

1.A Appendix: Estimates from previous literature . . . 38

1.B Appendix: Additional empirical results . . . 39

1.C Appendix: Details on the household data set . . . 49

1.D Appendix: The life-cycle model . . . 49

1.D.1 Preferences, endowments and asset markets . . . 49

1.D.2 Dynamic programming problem . . . 52

1.D.3 Solution method . . . 54

1.D.4 Calibration . . . 55

1.D.5 Computing the marginal propensity for car expenditures . . . 59

1.D.6 The role of the measurement period and the regression specification . 60

References . . . 64

2 Should I stay or should I go?
The role of housing in understanding limited inter-regional worker mobility 70

2.1 Data . . . 77

2.1.1 Sample selection and data sources . . . 77

2.1.2 Summary statistics . . . 79

2.2 Economic impact of the 2014 oil price plunge . . . 80

2.2.1 Supply glut and price plunge . . . 82

2.2.2 Impact on the Stavanger economy . . . 83

2.2.3 Post-2014 outlook for oil and Stavanger . . . 88

2.3 A simple theoretical framework of migration 90
2.4 Reduced-form results 93
2.4.1 Changes in the characteristics of Stavanger leavers 94
2.4.2 Changes in leavers' outcomes 103
2.4.3 Changes in the Stavanger arrivals 105
2.5 A life-cycle model with location choices 109
2.5.1 Economic environment 110
2.5.2 The worker's problem 111
2.5.3 Worker skills and wage premia 114
2.5.4 Model estimation 115
2.5.5 A grid-transformation trick 118
2.5.6 Main model results 119
2.6 Conclusion 128
2.A Appendix: Additional empirical results 130
2.B Appendix: Details on estimating worker skills and LMA-specific wage premia 137
2.B.1 Algorithm 137
2.C Properties of Gumbel distributed random variables 140
References 141

3 What do 12 billion card transactions say about house prices and consumption? 145

3.1 Data and sample construction 148
3.1.1 Data sources 148
3.1.2 Sample construction 150
3.1.3 Variable definitions and summary statistics 151

3.2 Research design and identification challenges 152
3.2.1 Experimental setting 153
3.2.2 First test of assumptions 155
3.2.3 Threats to identification 155
3.3 Results 158
3.3.1 IV model 159
3.3.2 Why do home prices affect consumption? 164
3.4 Conclusion 168
3.A Appendix: Comparison of measured expenditures and survey data 170
References 171

LIST OF FIGURES

1.1 Noise propagation around Bromma Airport 8
1.2 Effect on house prices 14
1.3 Model simulations – difference between treatment and control group 31
1.4 The role of the shock magnitude and the measurement period 34
1.B.1 Effect on house prices, quarterly 41
1.B.2 The evolution of economic variables inside and outside the noise area 42
1.B.3 Car outcomes 44

2.1 Brent oil spot prices and futures for June 2023 82
2.2 Impact on labor outcomes of workers in Stavanger 84
2.3 Housing costs 86
2.4 Changes in in- and out-migration trends for Stavanger 90
2.5 Leaving probabilities by housing tenure and housing wealth 96
2.6 Leaving probabilities by age 98
2.7 Leaving probabilities by labor income 102
2.8 Arrival probabilities by age 106
2.9 Arrival probabilities by housing wealth 107
2.10 Arrival probabilities by labor income 108
2.A.1 Stavanger migration in levels 130
2.A.2 Changes in housing costs over time 131
2.A.3 Additional results on the impact on labor outcomes of workers 132
2.A.4 Aggregate outcomes for Rogaland county 133
2.A.5 Additional results for arrival probabilities 134

3.1 The evolution of home price indices across Norway 150

3.2 Distribution of change in home prices 154

3.3 Variation in consumption across time, location, and consumption category . . 157

3.4 The dynamic marginal propensity for spending out of housing wealth shock . 161

3.5 The dynamic marginal propensity for spending out of housing wealth shock for retired homeowners 163

3.6 Variation in consumption across time for renters 164

3.7 Heterogeneity in the MPXs 167

3.8 Relative propensity to borrow 168

LIST OF TABLES

1.1 Summary statistics for single-family house and co-op apartment transactions . 10

1.2 Effect on house prices 12

1.3 Summary statistics for single-family house owners 15

1.4 Marginal propensity to borrow by LTV group 21

1.5 Reduced form results – extensive and intensive margin 22

1.6 Housing wealth elasticities 25

1.A.1 Estimates of housing wealth effects from previous literature 38

1.B.1 Summary statistics for other variables for single-family house and co-op apartment transactions 39

1.B.2 Summary statistics for single-family houses inside and outside the noise area before renewal 40

1.B.3 Summary statistics for co-op apartments inside and outside the noise area before renewal 40

1.B.4 Marginal propensity to borrow for all credit types 43

1.B.5 Triple differences – Effects on the value of new cars by groups (intensive margin) 45

1.B.6 Effects on number new cars bought, by housing tenure (extensive margin) . . 46

1.B.7 Effects on the value of new cars, by housing tenure (intensive margin) 47

1.B.8 Intensity of treatment on home prices and the intensive margin 48

1.C.1 Restrictions on the household data set 49

1.D.1 Calibration 57

1.D.2 Fitted parameters 59

1.D.3 Target moments 59

1.D.4 Reduced form estimates for different post measurement periods (model) . . . 62

1.D.5 MPXs based on different regression specifications 63

2.1 Summary statistics 81

2.2 Regressions with multiple covariates 100

2.3 Outcomes of leavers 104

2.4 Parameter values from external sources 116

2.5 Data moments targeted by model estimation 117

2.6 Model parameter estimates 120

2.7 Changes in simulated migration, by type of regional shock 121

2.8 Changes in simulated migration following an income and home price shock, by type of regional shock and housing tenure 122

2.9 Changes in simulated migration, by housing wealth 123

2.10 Welfare consequences of regional shocks 124

2.11 Corresponding changes in moving costs 125

2.A.1 Role of workers' idiosyncratic moving rates in the unemployed's higher mobility 135

2.A.2 Uptake of social welfare among Stavanger arrivals 136

3.1 Summary statistics 153

3.2 Difference-in-difference model results 156

3.A.1 Comparison to consumer surveys 170

ACKNOWLEDGMENTS

Reflecting on how I got here, I see that it has taken a global village to raise one researcher. I am profoundly grateful to all of you who have been part of this journey.

First, I thank the Swedish House of Finance, the Department of Economics at Stockholm University, the Research Division at Norges Bank, and the Economic Research Department at the Federal Reserve Bank of San Francisco for hosting me during enriching visits. The best part of my job is the exchange of ideas with fellow scholars.

I am also indebted to the Sweden-America Foundation and Handelsbankens forskningsstiftelser for their generous funding, making it possible for me to accept the offer from UCLA and fully focus on my research for several years.

My economics adventure started with Marieke Bos and Peter van Santen, who hired me as their research assistant. They saw some potential in me and encouraged me to pursue a Ph.D. Had they not planted that seed and supported me through applications, rejections, coursework, and the research stage, I would not be where I am today. I am beyond grateful.

Research thrives through collaboration, and I have been blessed to work with fantastic co-authors: Björn Tyrefors, Dany Kessel, Kasper Roszbach, Knut Are Aastveit, Magnus Gulbrandsen, Ragnar Juelsrud, and Roine Vestman. Our research assistants—Karin Ek, Niklas Nordfors, Mambuna Njie, and Yann Cerasi—provided invaluable help. I have learned tremendously from all of you, not just about conducting research but also how to live a good life as a researcher. Special thanks to Kasper and Roine for their boundless support, especially during the demanding final year.

My peers in the program have been indispensable allies. We supported each other through coursework, exams, and comps. Then the pandemic hit, and we helped each other through remote learning, teaching, and starting up our research. The macro group—Akira Ishide, Chang He, Fatih Ozturk, Luis Cabezas, Mariano Palleja, and Seungyub Han—thank you for the enriching reading groups, feedback on my work, and sharing this journey's highs and

lows. Adam Baybutt, Calvin Kuo, Domenico Fabrizi, Nicole Gorton, and Sumit Shinde: Thank you for your friendship.

To my committee: Lee Ohanian, Pierre-Olivier Weill, Andy Atkeson, and Kyle Herkenhoff. Your guidance, insights, and perspectives have been invaluable. You've instilled confidence when I felt lost and inspired me to pursue questions I am truly passionate about. I am forever grateful for your mentorship.

How can I express the importance of my family—my loving parents, sisters, their growing families, and our oblivious-to-economics pets? These years have taught me the roles you all play in my life. Though separated by miles and time zones, you are constantly in my thoughts. I thank you for the values you all teach me that brought me here: to always try my best, cherish the journey over the destination, and that family comes first.

Finally, thank you Chandni. Your love and ability to tell me what I need to hear—sometimes direct and instructive, *always* supportive and caring—has carried me through the Ph.D. I am blessed to have found someone so brilliant and virtuous. I look forward to writing our own chapters next.

VITA

2010–2015	M.Eng. and B.S. (Engineering Physics), M.S. (Applied and Computational Mathematics), KTH Royal Institute of Technology
2015–2017	Research Assistant, Sveriges Riksbank
2017–2018	Research Assistant, Swedish House of Finance, Stockholm School of Economics
2018–2020	M.A. (Economics), UCLA
2019–2022	Teaching Assistant, Economics Department, UCLA
2021	Visiting Research Fellow, Swedish House of Finance, Stockholm School of Economics
2021–2024	Visiting Graduate Student, Department of Economics and CeMoF, Stockholm University
2022–present	Ph.D. Student Intern and Guest Researcher, Research Division, Norges Bank
2023	Thomas J. Sargent Dissertation Fellow, Federal Reserve Bank of San Francisco

PRESENTATIONS

"The housing wealth effect: Quasi-experimental evidence" at

1. Southern California Graduate Conference in Applied Economics 2022, San Diego, USA
2. Research Division at Norges Bank (2022), Oslo, Norway
3. Econometric Society 2023 North America Summer Meeting, Los Angeles, USA
4. ASSA 2024 Annual Meeting, San Antonio, USA

"Should I stay or should I go? The role of housing in understanding limited inter-regional worker mobility" at

1. California State University, Long Beach (2023), Long Beach, USA
2. Stockholm University (2024), Stockholm, Sweden
3. Uppsala University (2024), Uppsala, Sweden
4. Sveriges Riksbank (2024), Stockholm, Sweden
5. Cornerstone Research (2024), Los Angeles, USA
6. 2024 AREUEA National Conference, Washington D.C., USA

CHAPTER 1

The housing wealth effect: Quasi-experimental evidence

with Roine Vestman, Björn Tyrefors, and Dany Kessel[1]

A fundamental topic in economics that has received a great deal of attention since the global financial crisis is how the housing and mortgage markets interact with the macroeconomy. In particular, there is a rich literature on how housing booms and busts affect household consumption—commonly referred to as housing wealth effects.

The early theoretical literature argued that housing was a particular asset that would generate no or small effects. Nevertheless, empirical studies have found mixed results, depending on the use of aggregate data (e.g., Carroll et al., 2011; Case et al., 2013; Guerrieri and Iacoviello, 2017) or household-level data (e.g., Attanasio et al., 2009; Browning et al., 2013; Campbell and Cocco, 2007; Disney et al., 2010) or on the interpretation of estimates.[2] In fact, many of the early contributors to the literature remained skeptical of their estimates due to weak identification.

[1]We thank Aditya Aladangady, Peter Fredriksson, Soren Leth-Petersen, Kurt Mitman, Rodney Ramcharan, Morten Ravn, Kathrin Schlafmann, Amir Sufi, and seminar participants at ASSA 2024, Danmarks Nationalbank, Lund University, Norges Bank, North American Summer Meeting, Stockholm University, the Greater Stockholm Macro Group, the Research Institute of Industrial Economics (IFN), the 2019 SED Annual Meeting, the Swedish House of Finance, Sveriges Riksbank, and UCLA. We are grateful for generous funding from Jan Wallanders och Tom Hedelius stiftelse, Marianne and Marcus Wallenberg Foundation, the Royal Swedish Academy of Sciences, and Vinnova. The computations were enabled by resources in project SNIC 2021/22-584 provided by the Swedish National Infrastructure for Computing (SNIC) at UPPMAX, partially funded by the Swedish Research Council through grant agreement no. 2018-05973. All data used in this research have passed ethical vetting at the Stockholm Ethical Review Board and have been approved by Statistics Sweden. Karin Ek, Mambuna Njie, and Niklas Nordfors provided excellent research assistance.

[2]See Footnote 2 of Berger et al. (2018) for a literature review and Table 1.A.1 for a review of estimates.

To date, the most credible estimates are based on instrumental variable regressions that rely on regional variation in the elasticity of housing supply (Aladangady, 2017; Aruoba et al., 2022; Kaplan et al., 2020; Mian et al., 2013) or city-wide variation in sensitivity to regional house prices (Guren et al., 2021). Such estimates make fairly strong assumptions about consumption demand factors being either observed by the econometrician or uncorrelated with supply elasticities (Davidoff, 2016).

This paper adds to this literature in three ways. The first contribution is that our estimates of the housing wealth effect are based on a novel identification—a quasi-natural experiment. We use unanticipated news from political bargaining in Stockholm, Sweden, regarding the continued operation of the city airport Bromma (or Bromma Airport) to isolate a causal effect from a relative change in house prices, which is a function of distance from the airport's noise contour. It is well-documented that the airport is a negative externality to its closest surroundings, and we show that this is capitalized into house prices within one quarter of the news announcement. Using a data set on all transactions of single-family houses in Stockholm, we document a price divergence, or a relative price fall, of 19.4 percent close to the airport's noise contour relative to further away. Our identification is novel in that it relies on an exogenous change of a negative externality that capitalizes locally into house prices. The source of variation is thus conceptually similar to e.g., Chay and Greenstone (2005) and Currie et al. (2015), but novel in the context of households' consumption response. The announcement is ideal for measuring households' consumption response because the effect is contained in a geographically granular area and the timing does not coincide with any important events that would differentially affect the treatment versus the control group.[3] We cannot find support that any other policy or event would have improved the wealth and income of households in proportion to the distance to the noise zone of the airport. Furthermore, the shock to wealth was as good as permanent because of the length of the extension of the airport. This is important for a wealth effect; it has to be permanent in

[3]We limit our experiment to before the financial crisis became global (before the bankruptcy of Lehman Brothers) to avoid any disruptions in credit conditions that are not due to the collateral value of homes.

nature to influence current consumption. These features of the quasi-experiment set our study apart from other studies.

We use the source of house price variation together with a rich household-level data set. This data set includes the geographic location of primary residence, household balance sheet information such as loan-to-value ratio and types of loans, and purchases of new cars at a quarterly frequency. We find that single-family house owners close to the airport reduce their purchase values of new cars by 7.7–8.5 percent relative to homeowners who reside further away.[4] A two-sample IV approach establishes a sizeable elasticity of 0.39 among households that buy a new car, which translates into a marginal propensity for expenditures on cars (what we call a "car MPX") of 2.5 cents per dollar loss in housing wealth per new car purchase.[5] This implies an aggregate MPX in new cars of 0.12 cents per year for each dollar lost in housing wealth and, assuming an equal response in used cars, implies an overall car MPX of 0.38 cents per year per dollar loss. The response of homeowners decreases with the distance to the airport's noise contour, and placebo tests show no response for households that live in apartments—two features that support our identifying assumption.

Our second contribution is that our data set enables analysis of the heterogeneous responses and the financing decisions of car purchases. This allows us to separate between the channels of the housing wealth shock. We find that homeowners with a loan-to-value (LTV) ratio above 50 percent respond twice as much, and that it is households with small bank deposits that respond. This heterogeneity is consistent with general borrowing behavior when purchasing a car. Forty-seven percent of a new car's value is financed with some kind of credit, and mortgages make up about 71 percent of these credits. However, there is substantial cross-sectional variation depending on households' balance sheets. Households with high LTV ratios borrow one-third less (per dollar car), and the difference is almost entirely explained by a reduction in the use of mortgages. These findings support the view that binding

[4] We use the term "homeowner" interchangeably with "single-family house owner."

[5] We adopt the terminology of Laibson et al. (2022) and let MPX stand for the marginal propensity for expenditure and denote by "car MPX" the marginal propensity for expenditure on cars. Analogously, we use the terms non-car MPX and total MPX in the model section. We clarify in the text when we refer to other concepts of marginal propensities. For a recent discussion on these, see Kaplan and Violante (2022).

borrowing constraints and the collateral channel are important for the observed total effect. Our results are complementary to Mian et al. (2013), who find that housing wealth shock responses in autos vary with the level of net worth at the ZIP code level, and to Aladangady (2017), Graham and Makridis (2023), and Aruoba et al. (2022), who find larger responses among credit-constrained households.

Our third contribution is that we relate our empirical findings to economic theory. Our estimates of housing wealth effects are the first ones generated from a quasi-experiment that resembles a partial equilibrium house price shock in the sense of Berger et al. (2018) and Guren et al. (2021). This is because the shock is geographically local (even granular); thus, general equilibrium effects are likely absent.[6] We use a state-of-the-art life-cycle consumption savings model that builds on Berger and Vavra (2015), Berger et al. (2018), McKay and Wieland (2021), and Attanasio et al. (2022).[7] The model includes elements that are relevant to our empirical setting: costly adjustment of cars, long-term mortgages, and an information friction for house prices.

We use the model to investigate responses in total consumption expenditure and cars to a partial equilibrium house-price shock. In simulations, we find that a shock of 19.4 percent to house prices leads to a reduction in the value of cars purchased by 6.1 percent in the next four quarters. This corresponds to a new-car MPX of 0.20 cents for each dollar change in housing wealth. This is close to our empirical estimate of 0.12 cents per dollar and inside the 95-percent confidence interval.

We establish a version of the intertemporal smoothing of durables demonstrated by McKay and Wieland (2021), applied to the context of housing wealth shocks (instead of monetary policy). The response in cars, as a share of the total expenditure response, is high at first:

[6] Our identification is strikingly similar to the one proposed by Carroll et al. (2011), page 71: "[...] to isolate a 'pure' housing wealth effect, one would want data on spending by individual households before and after some truly exogenous change in their house values, caused for example by the unexpected discovery of neighborhood sources of pollution. The perfect experiment observed in the perfect microeconomic dataset is however not available."

[7] Other seminal contributions to the literature on (S, s) models and their applications to households' durable goods or car purchases are Lam (1991), Eberly (1994), Bar-Ilan and Blinder (1992), Caballero (1993), Adda and Cooper (2000), Attanasio (2000), Hassler (2001), Bertola et al. (2005), and Schiraldi (2011).

45–72 percent in the first four quarters after the house-price shock and then gradually falls in the following years. The spending pattern in durables relative to total expenditure is a complementary finding to Laibson et al. (2022), who deduce the relationship between the marginal propensity to consume and the marginal propensity for expenditures.

The explicit modeling of a long-term mortgage enables us to use the model to distinguish between the role of a change in housing wealth and the role of changes to borrowing capacity. We find that changes to borrowing capacity account for 93 percent of the car MPX and 83 percent of the total MPX. In other words, if all households were unconstrained (and remained unconstrained after the shock), the short-term expenditure response would only be one-sixth as large. This finding is consistent with our heterogeneity analysis and several previous studies that have emphasized collateral effects (e.g., Aydin, 2022; DeFusco, 2018; Leth-Petersen, 2010; Sodini et al., 2023). Finally, the model allows us to pinpoint several important aspects of different empirical settings: the shock size, the measurement period, the regression specification, and whether it is conducted in normal times or crisis times, where the latter is likely to influence households' beliefs and awareness. If households are immediately aware of a house price shock, the consumption response is twice as large compared to our setting. In the first year, the total MPX can even be three times as large. These factors are important to reach a consensus view on the housing market's role in fluctuations in aggregate demand.

The rest of the paper is organized as follows. Section 1.1 reviews the institutional setting, the quasi-experiment, and the reaction of house prices to the announcement; Section 1.2 describes the data and discusses the empirical strategy. Section 1.3 reports our empirical results. Section 1.4 presents the model and our insights from it, and Section 1.5 concludes.

1.1 The quasi-experiment

This section describes the political process leading to the renewal of Bromma Airport's operating contract and establishes that it caused a decrease in house prices close to the airport.

1.1.1 History and political governance

Bromma Airport is the city airport of Stockholm. The airport has one runway and is located close to the city in an area that is otherwise dominated by single-family housing. It is Sweden's third-largest airport in terms of takeoffs and landings.[8] Since the late 1990s and early 2000s, however, there was a general perception that Bromma Airport would be closed, at the latest in 2011 when the operating contract would expire.[9]

Stockholm Municipality is the owner of the land on which the airport is located. The municipality has been leasing the land to the airport since 1936. The only political party in the municipality that consistently has been in favor of renewing the contract beyond 2011 is the conservative party (*Moderaterna*). In the 2006 election, this party increased their seats by more than 50 percent in Stockholm Municipality, from 27 to 41 seats out of 101. This was the best result ever for the party. The election outcome boosted the bargaining power of the party in the negotiations with the other parties in the center-right coalition. Rapidly, and behind closed doors, the municipality negotiated a new contract with the airport, extending the airport's operations to 2038. The contract was announced at a press conference in September 2007. All the opposition parties issued minority reports before the municipal

[8] In the years between 2006 and 2015, it had about 60,000 takeoffs and landings per year. In the early years, Bromma was Sweden's largest airport, but after the Arlanda airport opened in 1959, Bromma Airport saw a sharp decrease in traffic. In 1992, the center-right national government deregulated commercial airfare, and Bromma Airport increased in importance again.

[9] There was a series of reports planning for the shutdown. In 1989, Stockholm Municipality presented a major report proposing the closure of the airport by 1996 and the use of the land for housing (Stockholm, 1989). In 1994, the national government put together a commission to evaluate how fast Bromma Airport could be phased out (Kommunikationsdepartementet, 1996). In 2000, the Swedish Civil Aviation Administration presented a report on how it would eventually close Bromma Airport by 2011 (Luftfartsverket, 2000).

council, calling the process a coup. The news about the new contract was widely reported in local media.

The reason for the political controversy over Bromma Airport's existence is its geographic location. The airport is surrounded by residential housing and is a substantial negative externality on its surroundings—not the least in terms of noise. Figure 1.1 displays a map of the noise contour around Bromma Airport as a dark red ellipse. The area inside the noise contour is frequently exposed to noise levels of at least 70 decibels.[10] The contour is regarded as the best approximation of the area that is exposed to hazardous noise, as confirmed in a case in the Land and Environment Court (Miljööverdomstolen, 2010). The court ruled that, within this area, the regulatory agency that oversees aviation in Sweden (LFV) must reimburse sound insulation to homeowners. The measurement error of the border is +/– 100 meters.

Externalities from airports are known to be severe. Many studies find that the noise from aircraft is more hazardous than, for example, noise from trains or cars for a given daily average decibel level (Guski et al., 2017; Miedema and Oudshoorn, 2001).[11] This is because the house facades are not able to shield the noise from above, but also because the aircraft noise is intermittent and unpredictable. Furthermore, there may be fear of dumping of fuel and accidents at takeoff and landing.

We acknowledge that the magnitude of the negative externality is not perfectly proxied by the noise contour that corresponds to 70 decibels. To limit misclassification of the treatment group, we choose to include the area up to 1000 meters away from the border of the contour in our baseline definition of the treated area. Households residing in this area are exposed

[10] The decibel scale is logarithmic. Sixty decibels correspond to a conversation in an office. Seventy decibels are twice as loud. A vacuum cleaner is 70 decibels, which makes it difficult to have a conversation. A Boeing 737 generates 97 decibels before landing at a distance of one nautical mile (1,853 meters). Bromma Airport has permission to service takeoffs and landings that generate more than 70 decibels between 6 a.m. in the morning and 10 p.m. in the evening. Each aircraft type that operates at Bromma Airport must be noise tested, and the upper bound is 89 decibels. A noise level of 89 decibels is well above the threshold to mandate ear protection at Swedish workplaces. It is so loud that people would not be able to have a screaming conversation.

[11] We are not the first to use the fact that a negative externality capitalizes into real estate prices. Chay and Greenstone (2005) and Currie et al. (2015) use it to measure the cost of nearby air pollution and toxic plants.

Figure 1.1: Noise propagation around Bromma Airport

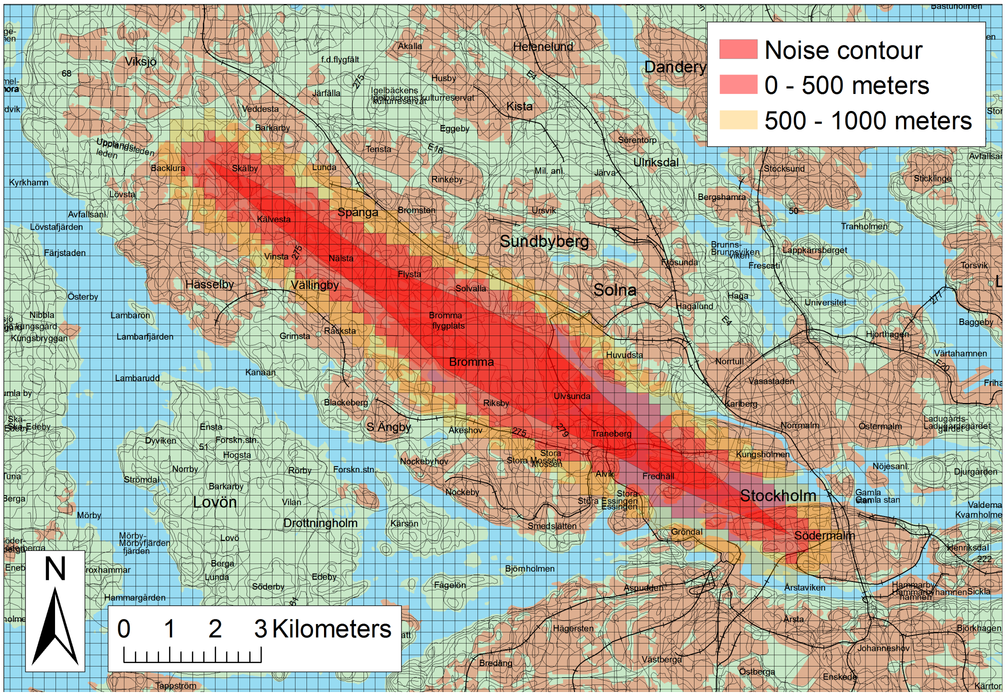

Note: The map shows the noise propagation around Bromma Airport along its runway which stretches from northwest to southeast. The dark red area is referred to as the noise contour. Inside the noise contour, the Land and Environment Court estimated the noise to exceed 70 decibels. The treatment region in our baseline specification is extended also to include the area which is less than 1000 meters away from the noise contour. We call the treatment region the noise area. We locate house transactions and households on the grid. Each square is 250×250 meters. Source: Miljödomstolen (2006) and own analysis.

to at least 60 decibels following the transmission formula for noise:

$$L_2 = L_1 + 10 \times \log_{10}\left(\frac{r_1^2}{r_2^2}\right), \tag{1.1}$$

where L_1 denotes the decibel level at a distance of r_1 meters from the source and L_2 is the decibel level at distance r_2.

The baseline treatment area, marked as yellow in Figure 1.1, is henceforth referred to as "the noise area." In the following, we refer to single-family houses as simply "houses" and households as being located either inside or outside the noise area. We argue that the negative externalities should be negligible outside this area, enabling us to construct a control group of unaffected households.

1.1.2 Data on transactions of houses and apartments

To measure the treatment effect on house prices, we acquire transaction data on houses and apartments. First, we obtain data on all transactions of houses in Stockholm Municipality from the Land Survey Agency (*Lantmäteriet*). The data set covers all transactions and prices in the municipality from January 2004 to December 2012. The data set includes a large number of characteristics for each house, such as the transaction date, the geographic location in the form of GIS coordinates, the area of the building plot, and information about the houses, such as living area, supplementary area, age, and an index of the attractiveness of the location and the house standard. This index is used by the Swedish Tax Agency (*Skatteverket*) to assess the value of the property. The GIS coordinates allow us to compute the distance between the house and the noise contour.

Stockholm Municipality has about 410,000 dwellings, of which approximately 90 percent are apartments, either rentals or co-op shares,[12] and 10 percent are houses. Therefore, we also collect data on transactions of co-op apartments from *Mäklarstatistik* (a data-collection company owned by the Association of Swedish Real Estate Agents). We have data from 2005 to 2010. In this data, we have the transaction date, GIS coordinates, living area, number of rooms, and price.

It is important to distinguish between co-op apartments and single-family houses. Apartments are different because they have no private outdoor area, which makes them less exposed to noise. Furthermore, the co-op buildings are at least three stories high and have thicker walls and more insulation due to fire safety regulations. They are solid concrete buildings and not made of wood like many single-family houses. Therefore apartment owners are expected to suffer less from noise, and hence apartment prices should not react as much upon the announcement of the renewal of the airport contract.

[12] In Sweden, apartment buildings and the units are co-owned through associations. When a household buys an apartment, it buys a co-op share in that association linked to a specific apartment. According to Statistics Sweden, all tenant-owned dwellings in multi-dwelling buildings in Stockholm were co-op apartments during the period of interest. This distinction is not of importance for our analysis.

Table 1.1: Summary statistics for single-family house and co-op apartment transactions

	Panel A: Full sample			
	Single-family houses		Co-op apartments	
	Price	Living area	Price	Living area
Mean	3351	117	2409	62.6
Std. dev.	2244	38.1	1698	29.3
Num. obs.	19,777	19,666	85,168	85,168

	Panel B: Before 2008Q3			
	Single-family houses		Co-op apartments	
	Price	Living area	Price	Living area
Mean	2947	117	2292	62.3
Std. dev.	1955	38.5	1620	29.6
Num. obs.	11,321	11,308	50,312	50,312

Note: Transactions of single-family houses start in 2004Q1 and end in 2012Q4. Transactions of co-op apartments start in 2005Q1 and end in 2010Q4. Amounts are in SEK 1000 and living area in square meters. For additional variables see Table 1.B.1.

Table 1.1 shows summary statistics of the transaction price and the living area, the two variables that exist in both data sets. Panel A reports statistics for the full sample period. The average transaction value of a single-family house is SEK 3.4 million, and the average living area is 117 square meters. Apartments are cheaper, about SEK 2.4 million, and are smaller; the living area is 63 square meters on average.

The full sample period includes the global financial crisis and the transaction volumes dropped at that time. Since we aim to isolate the effect of the airport contract's renewal in September 2007, we focus on the period up until the bankruptcy of Lehman Brothers in 2008Q3. Panel B reports the statistics for this period. The mean and standard deviations of the living areas are similar when comparing the two periods, but prices are lower for the pre-crisis period, which is consistent with the secular increase in Swedish home prices. Table 1.B.1 reports statistics of additional variables in the data sets. Tables 1.B.2 and 1.B.3 compare statistics for transactions inside and outside the noise area before the renewal of the airport contract. Prices of both houses and apartments are similar in the two areas before the renewal.

1.1.3 Empirical strategy for measuring the house-price shock

The identifying variation comes from the unexpected renewal of the airport contract in combination with the location of the residence. Dwellings close to the airport suddenly faced at least another 30 years of negative externalities. Our outcome variable is the natural logarithm of prices of dwellings sold for household i in time period t, $\log(\text{price}_{it})$. We define the treatment period to start on October 1, 2007, also denoted as 2007Q4. The variable noise area_i defines the treatment group in the sense that the variable takes on a value of one if the dwelling is located within 1000 meters from the noise contour (and otherwise zero). The variable post_t is zero up until and including 2007Q3 and one thereafter. A standard difference-in-difference equation reads

$$\log(\text{price}_{it}) = \alpha + \delta \, \text{noise area}_i \times \text{post}_t + \theta \, \text{noise area}_i + \eta \, Z_{it} + \gamma_t + \varepsilon_{it}, \tag{1.2}$$

where α is an intercept. The coefficient γ_t indicates year-quarter-time fixed effects, Z_{it} is a vector of data on the transacted unit, and ε_{it} is an error term. Thus, θ measures the average log point (percent) difference in prices between the dwellings inside and outside the noise area. The coefficient δ is of primary interest. It measures the log point (percent) change in prices due to the renewal of the contract. Standard errors are based on clustering of error terms at the level of the 250×250 meters grid.

1.1.4 Estimates of house-price effects

Table 1.2 reports estimates of equation (1.2). Starting with the effect on prices of single-family houses, Column (1) reports the results for the pre-global financial crisis period with no control variables. The estimated price decrease for the houses within the noise zone is

Table 1.2: Effect on house prices

	Log of house prices			Log of apartment prices		
	(1)	(2)	(3)	(4)	(5)	(6)
noise area$_i$	−0.214***	−0.214***	−0.194***	0.027	0.002	−0.019
× post$_t$	(0.040)	(0.035)	(0.028)	(0.018)	(0.014)	(0.014)
noise area$_i$	0.202***	0.211***	0.222***	0.042	0.135***	0.135***
	(0.043)	(0.035)	(0.034)	(0.037)	(0.031)	(0.032)
Observations	11,321	11,308	19,666	50,312	50,248	85,048
R-squared	0.102	0.330	0.374	0.049	0.430	0.437
Pre-GFC	Yes	Yes	No	Yes	Yes	No
Controls	No	Yes	Yes	No	Yes	Yes

Note: The set of control variables for the housing prices regression includes a polynomial in age, standard, plot area, living area, and non-living area, and the change in the amount of property tax to be paid due to the tax reform coming into effect in 2008. For the apartment prices, regressions controls are living area and number of rooms. Errors cluster robust at the level of the 250×250 meters grid. The baseline time window is 2004Q1–2008Q3 for single-family houses and 2005Q1–2008Q3 for co-op apartments (pre-GFC). The long time window extends the baseline time window to 2012Q4 and 2010Q4, respectively. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

−21.4 percent. Adding control variables has no impact on the estimate.[13] Thus, there is no evidence of any compositional bias (i.e., different types of houses being sold before due to the airport continuation). If we use the longer sample period, we see a slight decrease in the estimate to −19.4 percent (Column (3)). We take this to be our baseline first-stage effect. We conclude that the effect on house prices is highly significant (t-stats > 5 in all specifications) and robust across specifications. Columns (4) to (6) report the same estimates for co-op apartment prices. In contrast to the effects on single-family houses, there are no statistically significant effects, and the estimates are close to zero.

To validate our DiD specification, we test for parallel trends in outcomes prior to 2007Q3. We augment equation (1.2) with yearly time-dummy variables and define them based on time relative to treatment. Since the renewal of the operating contract was disclosed in late September 2007, we define every 12-month period as running from October 1 of year $t-1$ to September 30 of year t. We interact this set of time-dummy variables with noise area$_i$ and

[13] We lose 13 observations due to missing data in the control variables.

omit the dummy variable for 2006Q4–2007Q3 so it serves as the reference level for prices outside the noise area.[14] That is, for every 12-month period, we estimate a treatment effect relative to the 12 months just before the contract renewal.

The top panels of Figure 1.2 show single-family house prices. Panel A displays house-price fluctuations inside and outside the noise area. The price series are indexed to 100 in 2004. The series follows the same trend from 2004Q1 to 2007Q3 but after the new contract, the series diverge. Panel B displays the corresponding event study estimates on house prices using the specification in Column (3) in Table 1.2. Common pre-trends up to 2007Q3 cannot be rejected, and the immediate effect in 2007Q4–2008Q3 is −20 percent. The price difference remains up until 2011. However, since our aim is to isolate the effect of the contract's renewal in September 2007, we focus on the period before the financial crisis became global—that is, from 2004Q1 to 2008Q3 (the area shaded in gray). For completeness, the bottom panels of Figure 1.2 show co-op apartment prices, for which there is no treatment effect.

To support the credibility of our identifying assumption that the announcement was unanticipated, Figure 1.B.1 reports a version of Panel B of Figure 1.2 estimated at quarterly frequency. The figure shows that the response of single-family house prices was immediate in the fourth quarter of 2007.

To conclude, the renewal of the contract had a direct impact on single-family house prices but not on apartment prices. The results are robust when we account for compositional changes in transacted units, and several additional aspects speak in favor of a causal interpretation. Going forward, single-family house owners are the focus of our analysis.

1.2 Household data

Our main analysis uses registry-based panel data provided by Statistics Sweden that covers all households whose residential address is located in the municipality of Stockholm. This data set has information on age, household size, balance sheet items, and car transactions.

[14] In our graphical illustrations, 2006Q4–2007Q3 is referred to as 2007 and so forth unless explicitly stated otherwise.

Figure 1.2: Effect on house prices

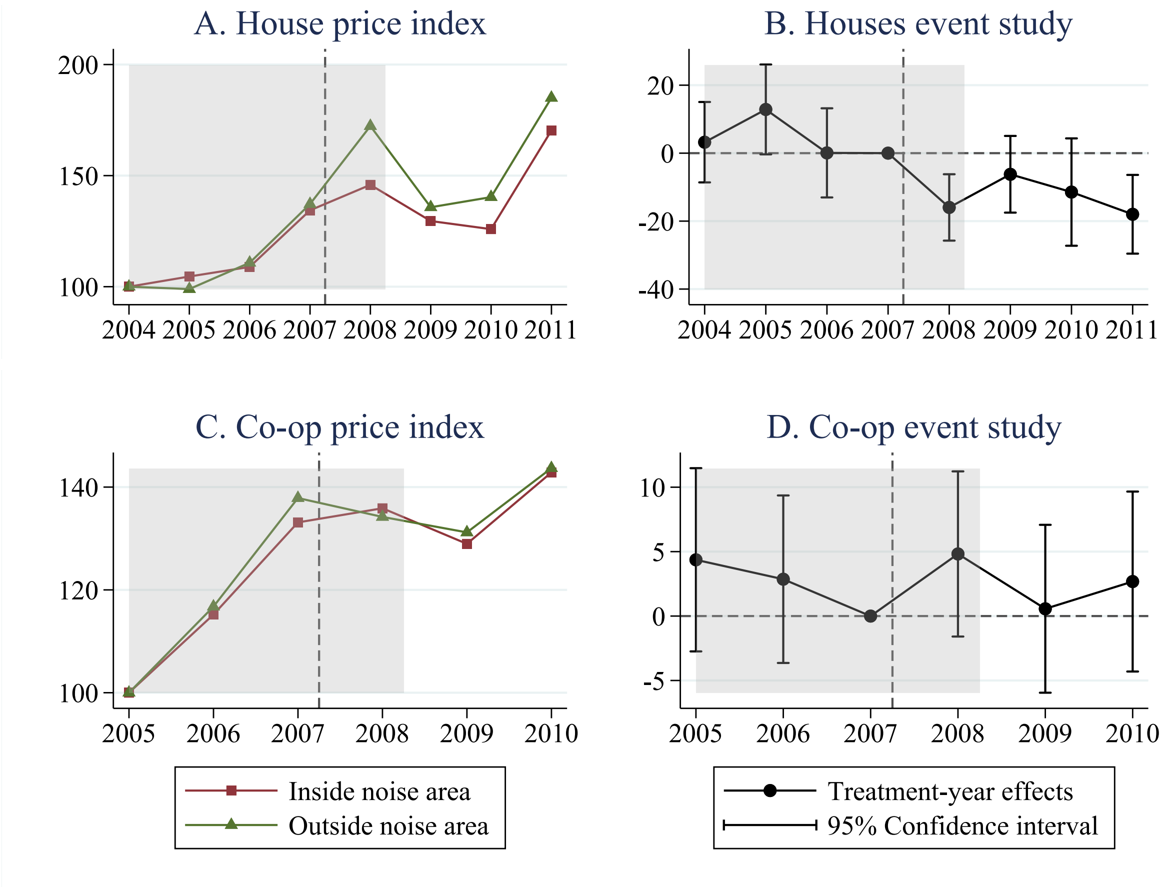

Note: Panels A and C show house and apartment prices indices outside and inside the noise area, respectively. Panels B and D show the corresponding annual treatment effects in log points for house and apartment prices, respectively. The timing on the horizontal axis is shifted by one quarter. That is, 2007 refers to 2006Q4–2007Q3 and so forth. The shaded grey area marks the main sample period (2004Q1–2008Q3). The set of control variables for the housing prices regression includes age, standard, plot area, living area, and non-living area, and the change in the amount of property tax to be paid due to the tax reform coming into effect in 2008. For the apartment prices, regressions controls are living area and number of rooms. The regression specification is an augmented version of equation (1.2) and is described in the text.

Most of the information has an annual frequency, but for cars we have exact transaction dates.

Table 1.3 reports statistics on Stockholm's house owners at the end of 2006, just before the renewal of Bromma Airport's contract.[15] The first two columns report statistics for the full sample. The table also reports the statistics for those outside the noise area (Columns

[15]We impose five sample restrictions and cover 86 percent of the house owners in our analysis. A detailed description of the sample restrictions can be found in Table 1.C.1.

Table 1.3: Summary statistics for single-family house owners

	Full sample		Outside noise area		Inside noise area	
	Mean	SD	Mean	SD	Mean	SD
Panel A: Geographic location and sociodemographics						
Distance from noise contour (meters)	3070	2450	4020	2120	321	338
Age	53.1	13.5	53.3	13.5	52.6	13.4
Household size	2.85	1.33	2.83	1.33	2.89	1.34
Panel B: Balances sheets and income						
Total wealth	5060	9150	5090	9580	4960	7750
Housing wealth	3990	7230	4000	7390	3940	6750
Financial wealth	1070	4620	1090	5080	1020	2880
Bank deposits	306	954	308	1014	308	1010
Total debt	1300	2700	1300	3030	1280	1350
Mortgage	1180	2590	1180	2920	1160	1230
LTV (%)	39.0	75.6	39.3	84.5	38.2	39.6
CLTV (%)	40.5	64.8	40.6	69.6	40.2	48.1
Labor income	605	692	604	734	610	551
Capital income	72.5	990	75.7	935	63.1	1130
Panel C: Car variables						
Ownership (%)	81.7	38.7	81.0	39.2	83.5	37.1
Cars	1.10	0.753	1.08	0.750	1.15	0.759
Purchase freq. (new)	0.049	0.215	0.048	0.214	0.050	0.217
Car value (new)	254	125	255	125	252	124
Households	39,342		29,218		10,124	

Note: The full sample consists of all single-family house owners in Stockholm fulfilling our selection criteria and who are observed in years before 2007. Income variables are for 2006; all amounts are in SEK 1000. The loan-to-value ratio (LTV) is defined as mortgage debt divided by housing wealth and the combined loan-to-value ratio (CLTV) is all debt divided by housing wealth. Net worth is the difference between the value of all the household's assets and its total debt. The exchange rate is approximately 8 SEK/USD. Ownership is a dummy variable that takes on a value of one if the household owns at least one car. The purchase frequency for cars is annual. Car value refers to the value of new cars at the time of purchase.

(3)–(4)) and those inside (Columns (5)–(6)). Panel A reports that house owners inside the noise area have on average 2.89 members. Households outside have marginally fewer (2.83). The average age of the oldest household member is 53 years. Panel B shows that in terms of wealth and other balance sheet metrics, the two groups are very similar. Average financial wealth is SEK 1,020,000 versus SEK 1,090,000. By Swedish standards, these are substantial amounts of financial wealth; Vestman (2019) reports that average financial wealth among house owners equals SEK 448,400 for the period 2000–2007.

Based on Statistics Sweden's appraisal model for single-family houses, both groups' housing wealth is just over SEK 3.9 million. We know the houses' location on the 250×250 meters geographical grid and can thus infer the housing wealth shock implied by the quasi-experiment.

Financial institutions report to the Swedish Tax Agency each credit that households have. We categorize each institution by the type of credit they offer: mortgage, consumer lending, auto lending, agricultural and forestry lending ("Ag lenders", loans collateralized by farmland and property, and forestland), debt collection, and debt to government agencies. Institutions that have not specialized in one type of credit are marked as "mixed" (17 percent of the number of credit items). Unidentified credit suppliers (less than 0.1 percent) are also included in this category.[16] We are able to identify mortgages and auto loans with high precision.[17]

Panel B of Table 1.3 shows that total debt and mortgage debt are similar for the two groups. We define the loan-to-value (LTV) ratio as total mortgage debt divided by housing wealth, and the combined loan-to-value (CLTV) ratio as total debt divided by housing wealth. These

[16] If a financial institution is marked as mixed, we look at its share of outstanding loans that have a balance above SEK 100,000 (approximate USD 12,500). If this share is above 70 percent a year, we categorize the bank as a mortgage lender. This affects about 0.1 percent of all observed credits. We assume that they are not Ag lenders because Ag lenders are highly specialized and easily identified. Student debt falls under debt to government agencies, and this category includes CSN, the agency that provides very advantageous student loans at a low interest rate and with a very long duration. No credit history check is required to get a student loan, and no private alternatives exist in Sweden.

[17] Most Swedish banks have special subsidiaries that manage mortgage lending (so-called *hypoteksbolag*), so we are confident in our identification of mortgage lenders in particular. In December 2007, the outstanding stock of mortgages to Swedish households was SEK 1534 billion (Source: Statistics Sweden); when we use our categorization on the total stock of Swedish household debt, mortgages amount to SEK 1565 billion. Auto lenders are also easy to identify based on their names.

measures of indebtedness are also similar across groups, with averages of 39.0 and 40.5 percent, respectively.

The variables that differ the most between the inside and outside groups are the car ownership rate (81 percent vs. 84 percent) and capital income (SEK 75,700 vs. SEK 63,100). The latter has a very thick right tail, as reflected by the high standard deviation.

Statistics Sweden's car registry contains annual observations of each car with variables such as model, brand, the exact date of purchase, current owner, the two former owners, and the value if it is a new car. We can thus observe essentially every car transaction that involves a household in our sample.[18]

Based on this data, we construct two data sets: an annual panel for all Stockholm homeowners to study the credit financing of car purchases and a quarterly panel for the main analysis to track the response of households to the renewal of Bromma Airport's contract.

1.2.1 Credit financing of car purchases

To set the stage for our main analysis, we use our registry-based data set to document household borrowing behavior at the time of car purchases. It is well-documented that U.S. households frequently use credit to finance car purchases. McCully et al. (2019) find that more than 70 percent of new cars are funded by either home equity or auto loans. We are unaware of any documentation of Swedish households' borrowing behavior in conjunction with car purchases and therefore establish basic facts on borrowing to finance car purchases among Swedish single-family house owners.[19]

We focus on how borrowing behavior varies with the LTV ratio because it is most likely to determine a household's ability to use its house as collateral.[20] To estimate marginal

[18] In cases where the value of a new car is missing, we use data from the Swedish Tax Agency. We match the car models by fuzzy string searching using the user-written Stata command `reclink2` (Wasi and Flaaen, 2015). Fifteen percent of the new cars have been assigned a value in this way.

[19] Grodecka-Messi et al. (2022) use a tailored data set from the Swedish Credit Bureau (UC) to document that Swedish households to a high degree reoptimize their debt portfolio when they make equity withdrawals in response to growth in house prices.

[20] We restrict the sample to house owners that do not move during the year and that purchase at least one car during the year.

propensities to borrow (MPBs), we run the following regression:

$$\begin{aligned}\frac{\Delta \text{credit}_{it}^{k}}{\text{car value bought}_{it}} = {} & \beta_l\, \mathbb{1}(\text{LTV}_{it-1} < 50\%) \\ & + \beta_m\, \mathbb{1}(\text{LTV}_{it-1} \in [50\%, 100\%)) \\ & + \beta_h\, \mathbb{1}(\text{LTV}_{it-1} \geq 100\%) + \eta\, X_{it} + \epsilon_{it},\end{aligned} \tag{1.3}$$

where the outcome variable is the change in the credit of household i from the end of year $t-1$ to the end of year t divided by the value of all the cars bought during t. The k indicates the type of credit, which can be all credit, mortgage, consumer credit, or mixed credit. The dummy variables $\mathbb{1}(\cdot)$ indicate the household's LTV ratio at the end of $t-1$, and X_{it} is a vector of control variables, standardized to mean zero and unit variance. The coefficients of interest are β_l, β_m, and β_h, which are estimates of the MPBs for households in three different groups based on their LTV ratios. The cutoffs for the categorization were chosen to represent groups that are likely to face different borrowing constraints. The first group has an LTV below 50 percent. Households in this group should be able to take out a substantial second mortgage, regardless of house price fluctuations. The second group comprises households with an LTV between 50 and 100 percent and hence are probably close to the maximum they can borrow using a mortgage. We conjecture that this group's borrowing capacity is influenced the most by house price fluctuations. The last group comprises households with a very high LTV; they should be limited in taking out additional mortgage debt regardless of house price fluctuations.[21] Note that the regression is demanding because it requires information about both car values and credit and that the data set has a panel dimension. Throughout our analysis, we focus on the purchase value of the car rather than the net of purchases and sales. This is for three reasons. The first one is conceptual—in the aggregate purchases and sales of used cars would net out and hence it is appropriate to focus on new cars. Secondly, even if one is able to account for sales with an accurate appraisal model for used cars, the timing of purchases and sales of the replacement in the household would

[21] In 2006–2008, banks set their LTV constraints in this range. During this period, there was no legal LTV cap. See Finansinspektionen (2010).

potentially add a lot of noise. Thirdly, this approach is consistent with the majority of micro-level studies and it facilitates comparisons with studies using regional data.

We consider different kinds of credit k to be able to make more precise statements about household borrowing. If $\Delta\text{credit}_{it}^{k}$ is the change in total debt, then MPBs smaller than 1 indicate that households also use proceeds from sales of cars and financial wealth (e.g., bank deposits) to finance their car purchases.[22]

Table 1.4 presents our estimates. We find that the MPB on average is 0.47 (Column (1)). However, Column (2) shows that there is considerable variation in borrowing behavior depending on the LTV ratio. Households in the highest LTV category borrow a substantially smaller share of their cars' value. Central to our analysis is to what extent home equity is used as a source of collateral for mortgage loans. A comparison of Columns (1) and (3) shows that, on average, 71 percent (0.333/0.467) of the borrowing associated with a car purchase is a mortgage. Furthermore, the cross-sectional variation in borrowing propensities is even greater when focusing on mortgage borrowing. Column 4 shows that while households with an LTV ratio below 50 percent have an MPB of 0.35, households with an LTV ratio above 100 percent have no tendency to take on additional mortgage debt.

The estimated MPBs in Column (4) are consistent with the difference in the MPB in total debt; the MPB difference between the lowest LTV groups is insignificant, but the reduction in MPB as we move into the top LTV bin is very similar for total debt and mortgages. We interpret this as binding restrictions on supply as the LTV ratio increases and that households understand that mortgage debt is the cheapest form of credit. In Table 1.B.4, we present other credit types as well, showing that households with below-median bank deposits use more credit financing when buying cars.

The cross-sectional variation in MPBs leads us to conclude that house price fluctuations are likely to influence households' means of financing car purchases. A big divergence in house

[22]Since we only observe the change in debt at year-end, we underestimate debt outtake for the specific purchase; we expect that households amortize a non-trivial amount of the loan over the year, and this introduces a downward bias in the MPBs. On the other hand, we also do not observe amortization payments before the purchase. If households that plan to buy a car pay off more debt leading up to the purchase, then our estimates are upward biased.

prices of the kind we have documented can imply that the borrowing capacity of house owners far away from the airport increases more than the capacity of house owners closer to the airport. In the next section, we estimate the housing wealth effects.[23]

1.3 Results

This section presents estimates of housing wealth effects based on our quasi-experiment. The population function of interest is

$$\text{car outcome}_{it} = \alpha + \beta \log(\text{house price}_{it}) + \varepsilon_{it}, \tag{1.4}$$

where β is an elasticity. It measures how households' spending on cars responds to a change in house prices. The common issue when estimating β is that house prices change when economic circumstances relevant to decisions on cars change. Ordinary least squares therefore lead to biased estimates. We address this by applying a difference-in-difference (DiD) regression and ultimately obtain a two-sample instrumental variables estimate of β. We first present reduced form DiD estimates.

1.3.1 Reduced form

For reduced form estimates, we use the same difference-in-difference model as when we estimated the house price effect:

$$\text{car outcome}_{it} = \phi\, \text{noise area}_i \times \text{post}_t + \rho\, \text{noise area}_i + \eta\, X_{it} + \gamma_t + \varepsilon_{it}, \tag{1.5}$$

[23] An additional implication of the large cross-sectional variation in borrowing behavior is that there may be limitations on analyses that proxy car purchases since borrowing behavior depends on balance sheet characteristics.

Table 1.4: Marginal propensity to borrow by LTV group

	All credit		Mortgage	
	(1)	(2)	(3)	(4)
Intercept	0.467***		0.333***	
	(0.031)		(0.029)	
Low LTV, $< 50\%$		0.458***		0.354***
		(0.033)		(0.031)
Mid LTV, 50–100%		0.524***		0.310***
		(0.044)		(0.040)
High LTV, $> 100\%$		0.323***		0.116
		(0.094)		(0.087)
Low LTV – Mid LTV		−0.066		0.044
High LTV – Mid LTV		−0.201**		−0.194**
Observations	6647	6647	6647	6647
Controls	Yes	Yes	Yes	Yes

Note: This table presents the results from estimating model (1.3) for homeowners in Stockholm who do not move in years of purchasing a new car. The controls are 4th-order polynomials in household size, age, and disposable income. Each control variable is standardized. Standard errors are cluster-robust at the household level. The middle panel displays the F-tests of the differences between the estimated coefficients. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

where ϕ is the coefficient of interest. It measures the change in car outcomes inside the noise area relative to outside after the renewal of the airport contract. The coefficients γ_t indicate year-quarter effects, X_{it} is a vector of household controls, and ε_{it} is an error term.[24]

The results are presented in Table 1.5. Columns (1) and (2) present estimates for the extensive margin effect: the number of cars purchased per household and quarter. On this margin, there is no effect of the renewal. The probability of buying a car in the treatment group is imprecisely estimated at 0.029 percentage points per quarter, which is minuscule relative to a baseline probability of approximately 1.2 percent (4.9 percent per annum). Columns (3) and (4) present effects on the intensive margin: the effect on the log car value conditional on the household buying a new car. Households close to the airport respond by

[24] All estimates are based on the Stata command `reghdfe` by Correia (2016).

Table 1.5: Reduced form results – extensive and intensive margin

	New cars		log car value	
	(1)	(2)	(3)	(4)
noise area$_i$ × post$_t$	0.00029	0.00029	−0.085***	−0.077***
	(0.00054)	(0.00054)	(0.019)	(0.021)
noise area$_i$	0.00014	0.00012	0.000	0.004
	(0.00031)	(0.00031)	(0.012)	(0.012)
Controls	No	Yes	No	Yes
Observations	531,105	531,105	6045	6045
R-squared	0.001	0.001	0.011	0.042
Age	53	53	52	52
LTV (%)	39	39	42	42
Net worth	3400	3400	3860	3860

Note: The table presents estimates of the coefficients in equation (1.5). The outcome variable in Columns (1) and (2) is the number of new cars bought (i.e., extensive margin response). The outcome variable in Columns (3) and (4) is the log value of these cars (i.e., intensive margin response). We control for household size, age of household head, disposable income, and net worth, all in 2006. All regressions include year-quarter fixed effects. Errors are two-way cluster robust at the household and year-quarter level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

reducing car values by 8.5 log points. Adding control variables reduces the estimate only a little, to 7.7 log points.[25]

The lack of an extensive margin response combined with a substantial intensive margin response is somewhat surprising but not implausible. A strong extensive margin response is considered standard in (S, s) models (e.g., Bar-Ilan and Blinder, 1992; Caballero, 1993; Eberly, 1994). However, Attanasio et al. (2022) illustrate that different combinations of shocks can lead to such responses, for instance in the U.S. during the Great Recession. When we present the results from our model in Section 1.4, we will return to this discussion.

Having established that there are household responses to the renewal, we turn to IV estimation to obtain elasticity estimates.

[25] Figure 1.B.3 reports tests of parallel pre-trends of outcome variables.

1.3.2 Two-sample instrumental variable estimates

Reduced form estimates from the difference-in-difference model established that house owners respond to the renewal. We now estimate the elasticity to housing wealth using the two-sample instrumental variable approach of Angrist and Krueger (1992).[26] Obtaining a credible partial equilibrium estimate of β has a strong appeal because it is closely connected to economic theory, as discussed in Berger et al. (2018) and Guren et al. (2021). Our quasi-experimental design, with granular geographic divergence in house prices, makes our estimates particularly credible. For instance, Orchard et al. (2022) argues that households' responses are dampened in general equilibrium because of price effects, implied by non-perfectly elastic car supply (see also Gavazza and Lanteri, 2021, who analyze the price effect of shocks to credit market conditions). But it is unlikely that our geographically granular setting gives rise to such effects (see also Section 1.3.4). Furthermore, the econometric issues concerning difference-in-difference regressions raised by Orchard et al. (2022) and Borusyak et al. (2024) do not apply.[27]

Table 1.6 reports the instrumental variable estimates for the subset of households that purchase a new car. We focus on the intensive margin responses since we found no responses on the extensive margin. Columns (1) and (2) report baseline elasticities. Without additional control variables, the elasticity is 0.398, which is statistically significant at the one-percent level. Adding control variables barely reduces the estimate, which remains at 0.39.

There is a strong theoretical foundation for the amplification of the housing wealth effect due to balance sheet characteristics. Berger et al. (2018) and Kaplan et al. (2020) derive analytical expressions for the response to non-durable consumption in standard incomplete

[26] Inoue and Solon (2010) show that the inference problem is solved using a generated regressor correction as proposed by Murphy and Topel (2002). We follow the implementation of Fredriksson and Öckert (2013). We first estimate the equation (1.2) in the transaction data set. We then use the estimates to predict the quasi-random variation in house prices in the household data set. The predicted house prices are used as a covariate in equation (1.4). Different functional forms for the predictive variables are considered. The main specification utilizes noise area$_i$ and noise area$_i$ $\times$ post$_t$, but we have also considered the continuous "dose specification"—see Section 1.3.4.3. The IV estimates, $\widehat{\beta}$, may deviate somewhat from $\hat{\phi}/\hat{\delta}$ because the set of control variables in the two specifications differ. We thank Peter Fredriksson and Björn Öckert for sharing their Stata code with us.

[27] Carroll et al. (2011) page 71 also discuss the ideal quasi-experimental design.

market models; the response is a function of the household's marginal propensity to consume out of transitory income shocks, which depends on the household's liquid savings buffer and the ratio between housing wealth and net worth.

Since our data set includes detailed balance sheet items and car transactions, we are in a good position to explore these relationships in the context of durable goods. We split households into two groups based on their balance sheet characteristics and explore variation in the intensive margin response. First, we split the households by their LTV ratios. Given our findings of differential MPBs across the LTV distribution, it seems plausible that a house price shock that affects LTV ratios will amplify the house owners' response on car consumption. We find strong heterogeneity in responses. Columns (3) and (4) of Table 1.6 show that house owners with an LTV ratio greater than 50 percent respond almost twice as strongly to the house price shock than house owners with an LTV ratio below 50 percent (0.526 versus 0.269). We also split house owners by their amount of bank deposits, which is the most liquid form of financial wealth and hence a good proxy for the households' marginal propensity to consume out of transitory income shocks. For house owners with small deposits, the response is 0.694, whereas for the other house owners, the response is only one-fifth as large. Such strong heterogeneity in responses along this dimension is consistent with previous general findings on responses to income shocks (e.g., Baker, 2018, who find particular strong heterogeneity in responses of durables) or lottery prize gains (Fagereng et al., 2021) but has to our knowledge not been documented before in the context of a well-identified partial equilibrium house price shock.[28] Section 1.4 compares these findings to a quantitative model adapted specifically to our setting. Before turning to the model, we convert our estimates to MPXs.

[28] We present related reduced form estimates from triple differences in Table 1.B.5. This allows us to formally test the statistical significance of differences between groups. We also investigate differential responses along additional dimensions. Broader measures of financial wealth (that include mutual funds and direct stock ownership) reveal no significant difference.

Table 1.6: Housing wealth elasticities

	Full sample		LTV		Bank deposits	
			$\leq 50\%$	$> 50\%$	$\leq$ P50	$>$ P50
	(1)	(2)	(3)	(4)	(5)	(6)
log(house price$_{it}$)	0.398***	0.393***	0.269**	0.526***	0.694***	0.123
	(0.108)	(0.124)	(0.124)	(0.188)	(0.183)	(0.138)
Controls	No	Yes	Yes	Yes	Yes	Yes
Observations	6045	6045	3945	2100	2748	3297
Age	52	52	56	45	50	54
LTV	42	42	22	80	51	35
Net worth	3860	3860	4950	1830	2440	5050
Financial wealth	1517	1517	1860	879	508	2360

Note: The table presents the second-stage two-sample IV estimates of β in equation (1.4). The dependent variable is the log of car values. Standard errors are corrected for first-stage estimation in the house transaction data set. In the first stage, we control for taxation value, building age, living area, and non-living area, and a control variable for the change in the property tax code in 2007. In the second stage we control for household size, age of household head, disposable income, and net worth, all in 2006. Errors are two-way cluster robust at the household and year-quarter levels. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

1.3.3 The marginal propensity of expenditures (MPX)

What are the aggregate implications of our findings in terms of car MPXs? The loss in housing wealth for households inside the noise area is approximately equal to SEK 774,060 (SEK $3,990,000 \times 19.4\%$). The elasticity among purchasers of new cars implies that the reduction in spending is SEK 19,061 ($0.393 \times 19.4\% \times 250,000$). This is a substantial response among house owners that purchase a new car; a car MPX of 0.025 (i.e., 2.5 cents per dollar). To compare this MPX to estimates that rely on geographically aggregated data, we must adjust for the number of new cars households buy per year, which is 0.049. This implies that the MPX on cars among all house owners is 0.0012 (i.e., 0.12 cents per dollar). Mian et al. (2013) estimate an MPX on cars of 1.8–2.3 cents per dollar and Aruoba et al. (2022) 1.2 cents per dollar.[29]

[29] We report estimates from Figure 4 and Table 5 in Mian et al. (2013), consistent with Aruoba et al. (2022). Berger et al. (2018) start from elasticities to housing net worth shocks in Table 2 of Mian et al. (2013) and scale them by 0.25–0.33 to obtain housing wealth elasticities.

1.3.4 Robustness

We have undertaken some tests to further strengthen the credibility of our quasi-experimental design.

1.3.4.1 Tests of the exclusion restriction

For the IV estimates to be consistent, the exclusion restriction must hold. What circumstances would lead to violations?

One such circumstance would be if the renewal of the airport contract affects car purchases for other reasons. Since the obvious alternative use of the land is to build residential housing and connect public transport to the neighborhood, the demand for cars among prospective residents could conceivably be affected. However, we argue that the plans to convert the land were too diffuse at the time. For instance, due to soil pollution, it was estimated that it would take many years to clean the area, which would delay new construction. Another such circumstance would be tax reforms, but they had already been announced and both groups benefited equally. A third circumstance would be if the income of current residents is affected. This could happen in two ways. It could be the result of an indirect effect (in the sense of Guren et al., 2021) where demand adjustments of households spill over to income or the result of a direct effect on job opportunities at the airport. However, Figure 1.B.2 shows that there are no such effects. Indeed, any income effects would seem implausible given both the geographical granularity of the house price shock relative to the overall Stockholm labor market and that we study responses over just five quarters. We find neither an effect on the probability to move in the short or long run, so it seems unlikely that any housing externalities, in the sense of Rossi-Hansberg et al. (2010), materialize.[30]

[30] The baseline probability of moving in our sample is 8.3 percent per year and the treatment effect is estimated to be -0.0024 and is insignificant.

1.3.4.2 Placebo tests

We have confirmed that co-op owners and renters respond differently than house owners. These reduced form results are reported in Table 1.B.6 and 1.B.7.

1.3.4.3 Dose response

We have also performed additional analyses that strengthen the credibility of our quasi-experimental design. If noise is the fundamental cause of the local divergence in house prices, there is an expectation that the house price shock would be muted monotonically with distance to the noise contour. Table 1.B.8 confirms a monotone relationship in terms of relative house prices and distance but also the house owners' responses and distance.

1.4 Model

We set up a model to compare empirical estimates of MPXs out of housing wealth with their theoretical counterpart. We also use the model to generalize the insights of Berger et al. (2018) on responses to house price shocks, taking into account durable goods, long-term mortgages, and information frictions. We provide an informal overview of our model and refer to Appendix 1.D for details.

1.4.1 Overview

Our model is closest to Attanasio et al. (2022) and Berger et al. (2018). Time is discrete and at a quarterly frequency. Households live from age 30 to 85 and retire at 65.[31] Income is exogenous, has a hump-shaped profile during working life, and is exposed to transitory i.i.d. shocks. Upon retirement, the household receives a predetermined fixed income that corresponds on average to a replacement rate of 70 percent.

[31] We motivate the age span based on our empirical setting. Only 2.1 percent of our sample of house owners have a household head that is younger or older.

Households have preferences for a non-durable and a durable good, where the latter is the car. The goods form a Cobb-Douglas consumption basket from which the household receives utility. When the household purchases a car, it pays an adjustment cost and chooses a new level of its car stock. Between time periods, the car depreciates. With a probability that we estimate in the model, the match-quality falls to zero and the household sells the durables and picks a new level of it.

From the start of life, the household is endowed with a house and a long-term fixed-rate mortgage (FRM). The house is used as collateral for the mortgage. In the last period of life, the house is sold and its value added to cash-on-hand.[32] In the last period, the household consumes as per usual and then receives utility from a bequest motive of what remains of their wealth. The car depreciates one more period, is sold and is added to the bequest. At any point in life, the household can adjust its mortgage amount subject to satisfying a down-payment constraint and pay an adjustment cost. The sum of interest and amortization is constant per period, as in, for instance, Campbell and Cocco (2003).

There is a second financial asset that households use to save or borrow. Borrowing in this asset does not require collateral and is limited by an exogenous borrowing constraint. There are interest rate spreads: the interest rate for borrowing is greater than for saving. Furthermore, the interest rate on the mortgage is lower than the interest on other borrowing. Unlike Attanasio et al. (2022), cars cannot be used as collateral for a loan. Instead, the household can use equity withdrawal through mortgage adjustments.

The adjustment costs for cars and mortgages make adjustments infrequent. In particular, there are inaction regions in the state space; (S, s) bounds. This feature can give rise to strong extensive margin responses to large shocks. However, recent additions to the literature have argued that strong immediate extensive margin responses do not always occur in the market for cars. This was the case in the Great Recession (see in particular Attanasio et al., 2022). There may also be a delay in response to aggregate shocks, such as monetary policy (McKay and Wieland, 2021). We therefore allow for an information friction. Households

[32] It is not possible for households to sell the house and become renters. We motivate this simplification from a low moving rate and the absence of treatment effect.

are not immediately aware of the spatial divergence in house prices. They only conceive the change to their housing wealth once they attempt to adjust their mortgage or trade cars or if they are hit by a random shock to their information set. We argue that this information friction is reasonable, in particular in "normal" economic times when our quasi-experiment occurred (recall that this was just before the U.S. financial crisis became a global crisis). Furthermore, we find it plausible that households do not frequently stay up to date on aggregate house prices as long as they do not fall. In this perspective, the friction appears to be a weaker assumption than in e.g., McKay and Wieland (2021), where households are unaware of aggregate variables.[33] We return to the implications of the information friction and its plausibility in different settings in Section 1.4.5.

1.4.2 Design of the quasi-experiment

We set up a replica of the quasi-experiment as follows. We calibrate the model to the sample of single-family house owners in Stockholm in 2006.[34] We simulate life-cycle paths for 200,000 households in the treatment and control groups. Each household has its identical twin in the other group, meaning that the twins have the same initial states and experience an identical sequence of idiosyncratic shocks to income. The difference between them is that in a random quarter, households in one group unexpectedly receive an increase of 19.4 percent in their housing wealth (from w^h to $w^{h'}$); an MIT shock.[35] The housing wealth shock is permanent. The shock affects both end-of-life wealth and the household's capacity to increase its mortgage balance. The spending response of the household depends on its other state variables: age, mortgage balance, cash-on-hand, and the current stock of cars. Our analysis of the consumption responses to this kind of partial equilibrium house price

[33] Carroll et al. (2020) and Auclert et al. (2020) have similar frictions to produce hump-shaped aggregate responses.

[34] To limit the support and size of the state space, we remove the top and bottom quartiles with respect to labor income when we compute our target moments. Therefore, several numbers will not coincide with the summary statistics in Table 1.3.

[35] We label the group that experiences the shock as the control group. As in our empirical setting, the treatment group (close to the noise area) experienced flat prices and the control group (i.e., further away) experienced increasing prices.

shock resembles the analysis of Berger et al. (2018), though our model includes endogenous responses to a durable good in the presence of long-term mortgages. These aspects are the focus of our analysis.

1.4.3 Consumption response

Figure 1.3 presents the average difference in consumption and spending responses (treatment group minus control group). Panel A illustrates the fall in total expenditures, non-durable consumption, and cars. Despite the Cobb-Douglas expenditure share for cars being only 5.5 percent, the reduction in expenditure in the short term is mostly due to a response in cars. Conditional on a car purchase in the first four quarters, homeowners reduce the car value by SEK 8,700–9,900. This translates to an unconditional car expense of SEK 1000 per quarter, which in turn amounts to 45–72 percent of the total response. The total MPX in the first year is 1.2 cents per dollar and the car MPX is 0.64 cents per dollar which implies a new car MPX of 0.20 cents per dollar (see Appendix 1.D.5).

Panels B, C, and D show other aspects of the response: the cumulative probability of a car purchase; the log car expense conditional on a purchase; and the log of non-durable consumption. We first focus on the solid lines, labeled "Total effect." Two features are noteworthy. First, Panel B shows that the extensive margin response is gradual as opposed to most (S, s) models. The second feature is illustrated by the solid line in Panel C. On impact, the response is 6.6 log points, and the average over the first four quarters is 6.1 log points. This translates to an elasticity of 0.31 (0.061/0.194), which is close to our empirical estimate and well inside the 95-percent confidence interval.[36] But the elasticity depends heavily on the measurement period. For instance, eight quarters after the shock, it is 20 percent smaller.[37] Finally, Panel D illustrates that the maximum elasticity in non-durables is attained only after ten quarters when it is 0.079 (0.0153/0.194). This translates to an MPC in non-durables of 1.1 cents per dollar.

[36] Based on Column (2) of Table 1.6, the confidence interval is 0.150–0.636.

[37] See Table 1.D.4 for average responses over measurement periods of 4, 16, and 32 quarters.

Figure 1.3: Model simulations – difference between treatment and control group

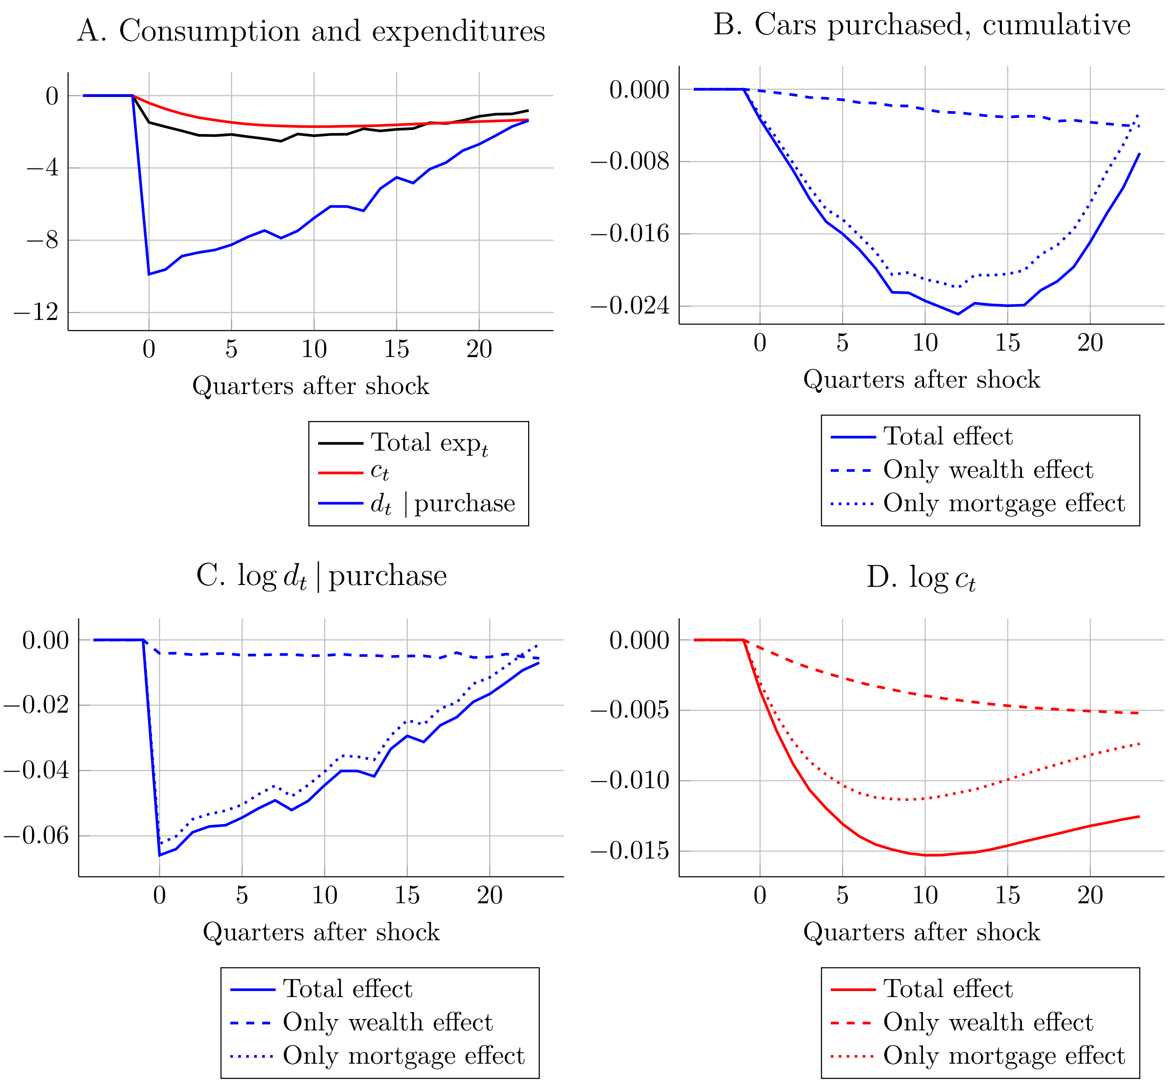

Note: The figure shows households' responses to the shock to housing wealth at event time $t = 0$. The plots are based on simulations of 200,000 households. The shock happens at a random age and every shocked household is compared to its benchmark counterpart. For each household, the response is calculated as the difference between the choice in the presence of the shock minus the choice in the absence of the shock. The values in the top left panel are in SEK 1000. The other panels display a decomposition of the shock into shocks to wealth and the mortgage cap separately. The top-right panel shows the difference in the cumulative number of cars purchased (extensive margin); the bottom-left panel shows the difference in the car value bought conditional on buying a car (intensive margin, in logs); and the bottom-right panel shows the difference in non-car consumption (in logs).

There are two main takeaways from our discussion so far. First, the total MPX and the MPC in non-durables are consistent with many estimates in the literature (e.g., Disney et al., 2010; Graham and Makridis, 2023; Guren et al., 2021), albeit in the lower range. Second, we think that the dynamic aspect of the responses and its implications for the empirical design have not been given sufficient attention when comparing empirical estimates. We return to this matter in Section 1.4.5.

1.4.4 Understanding variation across households

Motivated by the uncovered heterogeneity in responses from different balance sheet characteristics, we use the model to decompose the total response into a wealth effect and a collateral effect. We modify the baseline experiment and run it twice. In the first modification, to investigate the role of wealth, we shock the housing value but hold the borrowing capacity constant. In the second modification, to investigate the role of housing as collateral, we do not shock house prices but only households' borrowing constraints so the borrowing capacity in mortgages increases by as much as in the main experiment.[38] The results are presented as the dashed and dotted lines in Panels B, C, and D of Figure 1.3.

We find that the collateral effect dominates, and it is particularly strong for cars. The response in car expenditures (Panel C) in the first four quarters is 93 percent of the baseline response, whereas in the case of a pure wealth shock, the response is 7.5 percent of the baseline. After 20 quarters, the collateral effect is 69 percent of the baseline while the wealth effect is 32 percent. For non-durable consumption, the difference in force between the wealth and the collateral channel is somewhat less stark. The short-run collateral effect is 81 percent of the baseline, and the short-run wealth effect is 19 percent. In summary, the effects of wealth and collateral are consistent with the differential response in the data. Our empirical and theoretical results are also consistent with important contributions that emphasize collateral effects (e.g., Aydin, 2022; DeFusco, 2018; Leth-Petersen, 2010; Sodini et al., 2023).

From the point of view of macroeconomic stabilization policy, total expenditure is perhaps the most interesting, and the dynamics of cars strongly influence the short-run expenditure

[38] We shock the downpayment constraint from $\phi = 0.85$, which corresponds to a minimum downpayment of 15 percent, to ϕ'. In the first experiment $\phi' = 0.7119$. In the second experiment $\phi' = 1.0149$.

statistics: the collateral effect on total expenditure is 83 percent of the baseline in the first four quarters, and the wealth effect on total expenditure is 19 percent. This illustrates how credit market freezes can have strong immediate effects on aggregate demand (see e.g., Benmelech et al., 2017; Bernanke, 2018).

Relatedly, it is worth pointing out that while our model exhibits intertemporal shifting as in McKay and Wieland (2021), it is not due to monetary policy but a shift in households' credit portfolios. In our model, homeowners' effective interest rates on loans increase as they shift from mortgage financing to uncollateralized borrowing, and households gradually recover access to cheap credit, and at that point the value of their stock of cars recovers, too.[39]

Another important takeaway from this analysis is that the characteristics of consumption goods matter for households' expenditure responses to shocks. One such characteristic is the ability to postpone purchases—intertemporal shifting due to durability. Another one is if the good is credit-financed. Under a restriction of studying short-term housing wealth effects based on a single consumption item, cars are the preferable item.

1.4.5 Understanding variation across studies

We have set up a state-of-the-art model that matches our empirical evidence. However, the empirical literature has found a wide range of estimates of elasticities and MPXs.[40] Based on the combination of our model analysis and the conclusions of previous studies, we identify three factors that contribute to this variation. Figure 1.4 illustrates these factors.

1.4.5.1 Empirical setting and regression specification.

Several features related to the empirical studies' design can give rise to variation in estimates. One such feature is the shock magnitude. Panels A and B of Figure 1.4 plot the Total MPX and Car MPX in our model for different shock magnitudes in the interval from -30 percent

[39] Attanasio et al. (2022) achieve a similar effect by shocking the risk premium on auto loans.
[40] Table 1.A.1 reports car MPXs and car elasticities from previous housing wealth studies.

Figure 1.4: The role of the shock magnitude and the measurement period

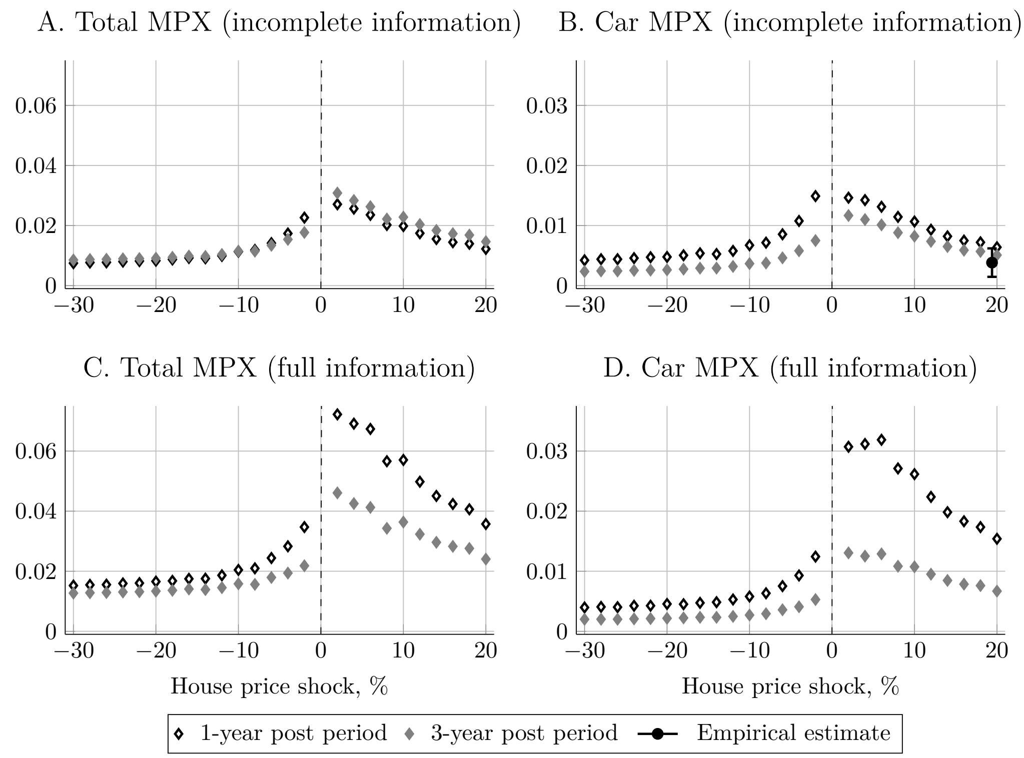

Note: The figure shows households' responses to different partial equilibrium housing wealth shocks. Each diamond represents the average across 200,000 simulated households. We mimic our other simulations by shocking households at a random stage in the life-cycle and compare them to identical households that experience no shock. The left panels display the marginal propensity for total expenditures. The right panels display the marginal propensity for car expenditures. Panels A and B present our baseline model with an information friction (incomplete information). Panels C and D present responses in the absence of this friction (full information). The black circle in Panel B indicates the empirical estimate from our quasi-experiment, adjusted so that it reflects the total of new and used car purchases. The error bars around the circle indicate a confidence interval of 95 percent. See Appendix 1.D.5 for details.

to +20 percent (i.e., shocks to w^h—recall that our experiment implies +19.4 percent—see the point estimate and confidence interval in Panel B). The panels show that the greater the magnitude, the smaller the partial equilibrium MPX. Car and total MPXs may be in the range of 1 to 2 cents per dollar. The reason is that as the shock size increases, fewer and fewer households' immediate consumption decisions are affected as they become constrained. Another empirical feature that gives rise to variation is the measurement period relative to the shock. The differently colored diamonds illustrate this. Short-run MPXs (say, 1-year: white diamonds) are greater than long-run ones (say, 3-year: gray diamonds).

This is particularly relevant for the lumpy spending on durables (Panel B), but also for total spending (Panel A). An aspect related to the measurement period is the regression specification of different studies. Because of the intertemporal dynamics of durable spending, the use of "long differences" in regressions—that is, outcome variables based on differences over several years—implies a downward bias in estimates. Such regression specifications do not capture that the initial response is greater, whereas a difference-in-difference specification identifies the average effect over the post-period. We illustrate this in Appendix 1.D.6 using data generated by our simulation. We find that for estimates of the car MPX, the bias can be 50 percent or more. The bias on total MPX is much smaller, in particular for large housing wealth shocks.

The state of the economy. The focus of our analysis is a partial equilibrium shock to house prices, but the underlying reasons behind the movement in house prices matter for the response. We have documented ample use of mortgage debt among car purchasers and matched that feature in our model.[41] Consequently, many households in our simulation are borrowing constrained and respond strongly to small positive house price shocks. In contrast, the U.S. financial crisis can be thought of as a classic boom-bust episode where credit supply was ample in the boom phase (Adelino et al., 2016; Mian and Sufi, 2018). In such an episode, the bust would imply that many homeowners become constrained, and thus the MPXs would be greater for negative shocks than positive ones. That is, such a state would yield the opposite asymmetry compared to Figure 1.4. Indeed, Guerrieri and Iacoviello (2017) report a car elasticity of 0.24 in the boom phase of the U.S. financial crisis and 0.49 in the bust phase. It is also plausible that households are more attentive to housing market developments in bad times and to falls in house prices (rather than increases). This reasoning is in line with Kaplan et al. (2020), who argue that changing household beliefs was a critical component of the U.S. financial crisis. We consider the effect of higher awareness by relaxing the information friction in our model: Panels C and D of Figure 1.4 report the

[41] The product of the subjective discount factor and the rate of return is less than 1, which implies that households borrow to front-load consumption.

resulting MPXs. The 1-year MPX to small positive shocks is nearly three times as large, with a car MPX of 3 cents per dollar and a total MPX well above 6 cents. Additional factors can amplify further the variation across studies: Attanasio et al. (2022) argue that factors such as shocks to expected income growth and to the risk premium on auto loans mattered for the responses of U.S. households during the crisis.[42]

General equilibrium effects. Finally, we wish to highlight the role of general equilibrium effects, discussed extensively in Guren et al. (2021). We argue that the longer the measurement period, the more likely it is that estimates are biased upward due to feedback loops and spill-over from other markets to the housing market. In combination with the use of long differences in car spending as the outcome variable, the net effect is unclear.

With these factors in mind, how do our empirical and model results relate to the literature?[43] Our empirical estimate of the car MPX and the fitted model's MPXs is clearly in the lower range relative to the literature. Yet, housing wealth MPCs of this magnitude are not unique; Browning et al. (2013) report an MPC of 0.0003–0.05 and Graham and Makridis (2023) an interval of 0.0078–0.0092. Adam Guren and his co-authors explicitly adjust for general equilibrium effects and report a partial equilibrium MPC of 0.018 and an MPC including equilibrium effects of 0.033 (Guren et al., 2021a,b).[44] If anything, later studies tend to report smaller estimates.

We see two main reasons why our estimates are on the low side. The first reason is that we view our estimate as a "normal" times estimate. Notice that the time window of our analysis entails no absolute fall in house prices but a mere price divergence, or relative house price fall. House prices in the noise area do not decrease during the five quarters that we consider. They remain flat, while house prices outside the noise area continue to increase. In contrast, responses in crisis times may be different as households become more attentive

[42] Relatedly, Berger and Vavra (2015) argue that durable expenditures do not respond as strongly to stimulus in recessions as in expansions. The mechanism is that more households remain inside their (S, s) bound even after the stimulus.

[43] See Table 1.A.1 for a literature review.

[44] Many studies report only elasticities and not MPCs. We omit them from this discussion, although it is noteworthy that some of those studies report low elasticities.

to economic conditions.[45] The second reason is that our experimental design is probably most conceptually similar to recent theory-based partial equilibrium experiments. The U.S. financial crisis includes a credit supply effect as well as an employment effect.[46]

1.5 Concluding remarks

Long before the global financial crisis, economists debated the extent to which the housing market influences the evolution of the macroeconomy through its roles as a store of value and collateral for borrowing. Recent studies have found such housing wealth effects using a variety of methods. However, there is substantial variation in empirical estimates.

We bring to the table a novel identification method—based on a quasi-experiment and household level data—that enables us to identify a partial equilibrium housing wealth effect, as recently defined and examined in Berger et al. (2018); Guren et al. (2021), with the important difference that our outcome variable is car expenditure.

Our estimated elasticity implies a new car MPX of 0.12 cents per dollar, which is small. Yet, conditional on purchasing a new car, house owners' response is substantial, around 8 percent, and house owners with an LTV ratio above 50 percent or with bank deposits below the median value respond much stronger than other house owners. We use a state-of-the-art model to argue that it is consistent with a short-run response to durable consumption in "normal times" in the presence of information frictions and absent general equilibrium effects. Further, the model enables us to highlight several factors that most likely explain some of the variation in estimates: measurement period, regression specification, and the state of the economy. To make progress on a consensus view on the housing market's role in aggregate demand fluctuations, future empirical studies should highlight each of these aspects.

[45] Guren et al. (2021) and Guerrieri and Iacoviello (2017) indeed argue that the housing wealth effect has varied over time.

[46] See Mian and Sufi (2014). See also page 3434 in Aladangady (2017) and footnote 12 of Berger et al. (2018) for discussions of general equilibrium effects. Also, Charles et al. (2019) document strong interaction effects between the housing and labor markets.

1.A Appendix: Estimates from previous literature

Table 1.A.1: Estimates of housing wealth effects from previous literature

Study, country	Elasticity	MPC	Car elasticity	Car MPX
Aladangady (2017),[a] U.S.	–	0.047	–	–
Andersen and Leth-Petersen (2020),[b] DK	–	0.03–0.05	–	–
Aruoba et al. (2022), U.S.	–	–	–	0.012
Attanasio et al. (2009),[c] U.K.	0.0	–	–	–
Browning et al. (2013),[d] DK	0.0–0.13	0.003–0.05	–	–
Calomiris et al. (2013), U.S.	0.163–0.270	0.049–0.081	–	–
Campbell and Cocco (2007),[e] U.K.	0–1.7	–	–	–
Carroll et al. (2011),[f] U.S.	–	0.02–0.09	–	–
Case et al. (2013), U.S.	0.065–0.068	–	–	–
Cloyne et al. (2019),[g] U.K.	0.2–0.3	–	–	–
Cooper (2013), U.S.	0.06	0.06	–	–
DeFusco (2018),[h] U.S.	–	0.04–0.13	–	–
Disney et al. (2010),[i] U.K.	0.087–0.120	0.01	–	–
Graham and Makridis (2023), U.S.	0.10	0.0078–0.0092	–	–
Guerrieri and Iacoviello (2017),[j] U.S.	–	–	0.24–0.49	–
Guren et al. (2021),[k] U.S.	0.040	0.018	–	–
Guren et al. (2021),[l] U.S.	0.072	0.033	–	–
Kaplan et al. (2020),[m] U.S.	0.06–0.12	–	–	–
Mian et al. (2013),[n] U.S.	0.13–0.26	0.054	0.33–0.43	0.018–0.023

Note: The table presents estimates from previous studies. Elasticities and MPCs either refer to total expenditure or non-durables.

[a] Aladangady in addition reports a zero effect for renters.
[b] Estimates a marginal propensity to borrow on mortgage debt.
[c] Estimates reported in Table 1 are positive but the authors refuse to interpret them as casual.
[d] Significant effect only for subsamples (young and constrained), i.e., no pure wealth effect. Estimates should be scaled with the proportion of such households to attain comparable aggregate estimates. Elasticity baseline=0.08 and MPC baseline 0.03 (Table 3, page 415).
[e] 1.7 is for older homeowners, 0 for young renters. I.e., pure wealth effect.
[f] The lower end of the estimates is the direct effect and the upper end is long-run estimate.
[g] The reported number is an elasticity of borrowing with respect to house prices.
[h] Same interpretation as above.
[i] For elasticities, the lower end is for old homeowners and the upper is for young homeowners. They also report results for young renters.
[j] They report two estimates of car elasticities: 0.24 in 2002–2006, and 0.49 in 2006–2010.
[k] The estimates are the P.E. housing wealth effect computed using results in Guren et al. (2021).
[l] We only report the estimates of their sensitivity instrument. See the paper for other estimates.
[m] Elasticities have been multiplied by the mean housing wealth to net worth ratio (H/NW); see footnote 12 in Berger et al. (2018). The elasticities are for non-durable consumption.
[n] Rescaled by H/NW as above.

1.B Appendix: Additional empirical results

Table 1.B.1: Summary statistics for other variables for single-family house and co-op apartment transactions

	Panel A: Full sample					
	Single-family houses					Co-op apartments
	Tax value	Age	Score	Lot area	Non-living area	Rooms
Mean	1751	77.2	28.4	540	47.7	2.30
Std. dev.	823	210	4.32	367	59.7	1.07
Num. obs.	19,777	19,666	19,666	19,777	19,666	85,048
	Panel B: Before 2008Q3					
	Single-family houses					Co-op apartments
	Tax value	Age	Score	Lot area	Non-living area	Rooms
Mean	1762	71.4	28.6	541	47.8	2.27
Std. dev.	858	189	4.32	349	35.9	1.07
Num. obs.	11,321	11,308	11,308	11,321	11,308	50,248

Note: This table complements Table 1.1. Transactions of single-family houses start in 2004Q1 and end in 2012Q4. Transactions of co-op apartments start in 2005Q1 and end in 2010Q4. All amounts are in SEK 1000.

Table 1.B.2: Summary statistics for single-family houses inside and outside the noise area before renewal

	Panel A: Inside noise area						
	Price	Living area	Tax value	Age	Score	Lot area	Non-living area
Mean	2807	120	1729	51.5	29.0	563	48.2
Std. dev.	1517	36.5	725	61.0	4.39	329	40.1
Num. obs.	2330	2329	2330	2329	2329	2330	2329
	Panel B: Outside noise area						
	Price	Living area	Tax value	Age	Score	Lot area	Non-living area
Mean	2672	116	1766	78.2	28.3	533	48.2
Std. dev.	1893	39.3	837	271	4.17	363	34.4
Num. obs.	6926	6926	6926	6926	6926	6926	6926

Note: Transactions of single-family houses start in 2004Q1 and end in 2007Q2. Tax value refers to the value assigned by The Swedish Tax Authority. Score is a hedonic variable that determines the tax value. Amounts are in SEK 1000 and areas are in square meters.

Table 1.B.3: Summary statistics for co-op apartments inside and outside the noise area before renewal

	Panel A: Inside noise area		
	Price	Living area	Rooms
Mean	2172	56.8	2.07
Standard deviation	1282	26.1	0.982
Num. observations	9764	9764	9740
	Panel B: Outside noise area		
	Price	Living area	Rooms
Mean	2240	64.3	2.33
Standard deviation	1700	31.2	1.10
Num. observations	26,608	26,608	26,579

Note: Transactions of apartments start in 2005Q1 and end in 2007Q2. Values in SEK 1000 and area in square meters.

Figure 1.B.1: Effect on house prices, quarterly

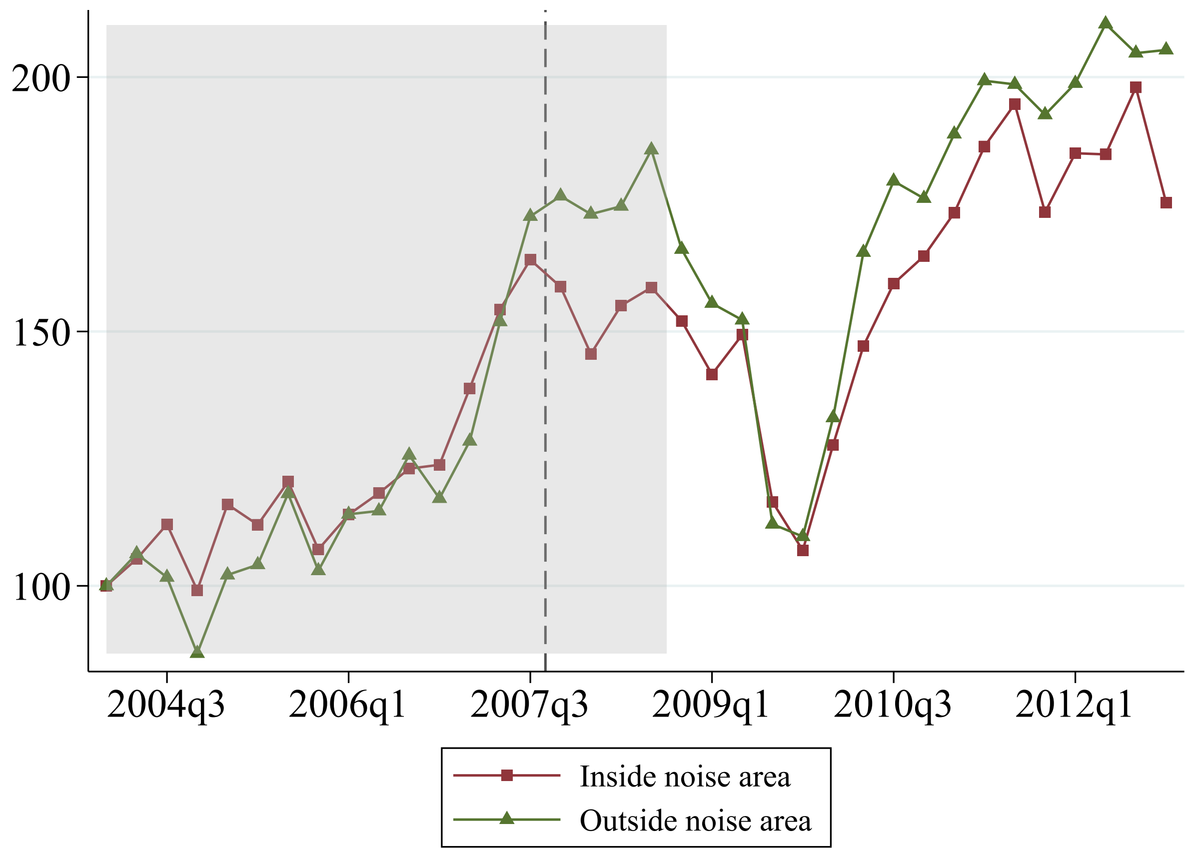

Note: This figure plots the quarterly house price indices outside and inside the noise area, respectively. No shifting of quarters is used. The gray-shaded region indicates the time period used in the empirical analysis.

Figure 1.B.2: The evolution of economic variables inside and outside the noise area

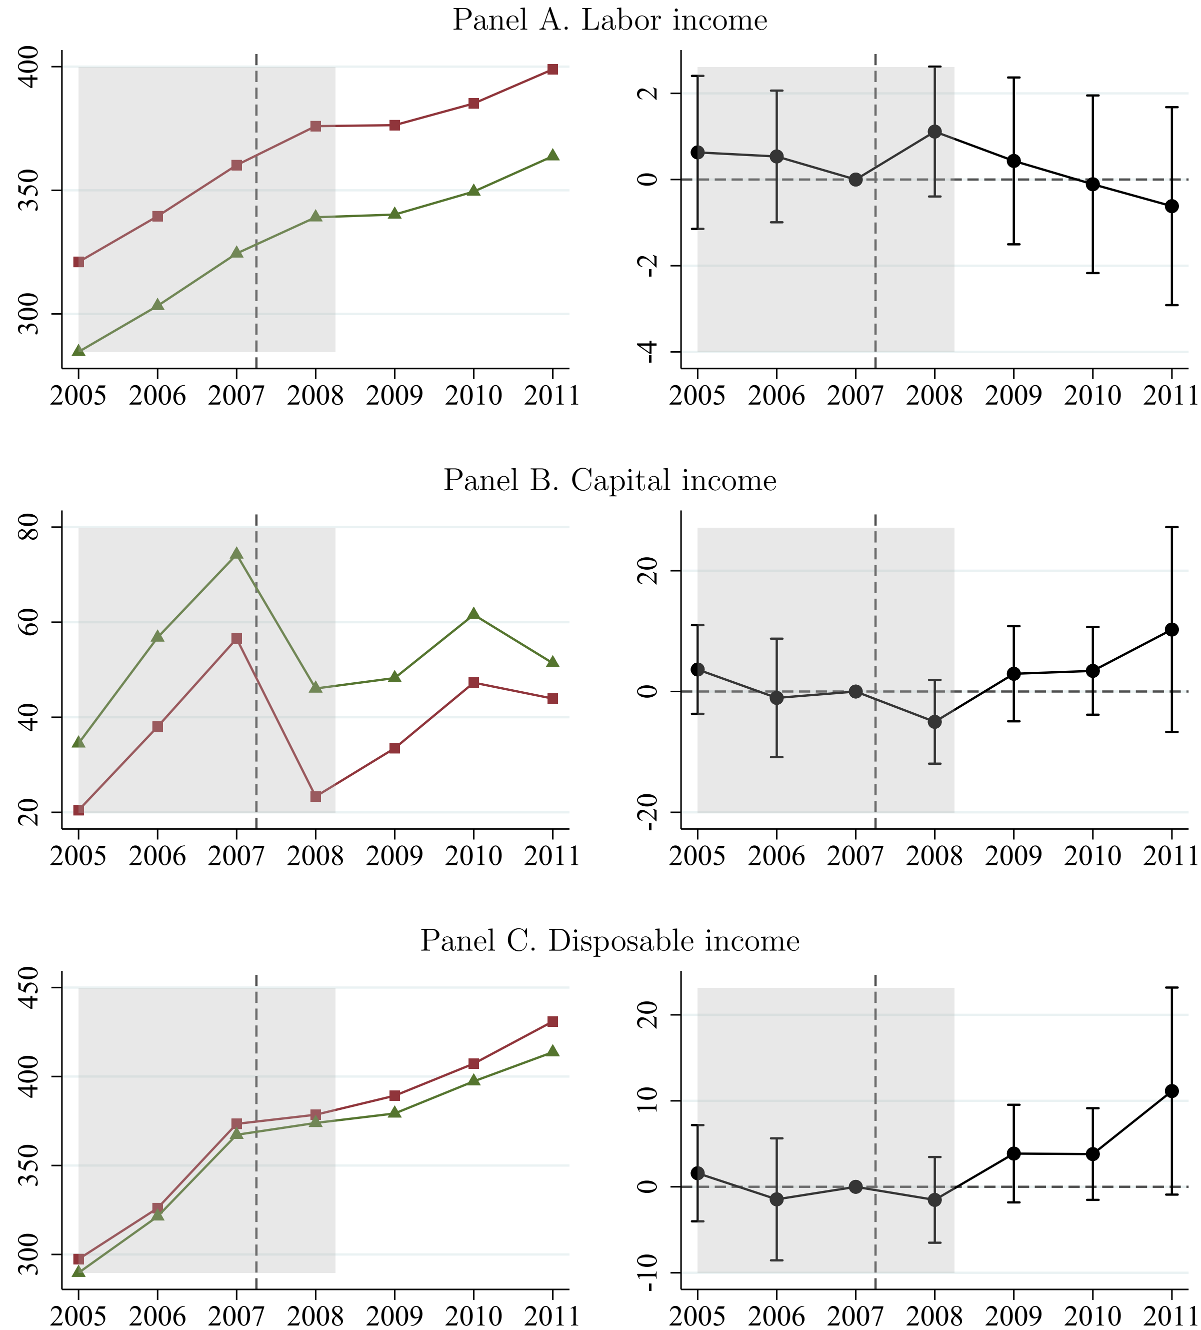

Note: This figure presents different income measures for single-family homeowners living inside and outside the noise area. All amounts are in SEK 1000. The left graphs plot the mean level of each measure by group and the right graphs plot the differences between the groups, using a regression model as in Figure 1.2. No additional household controls are used in the regression but household fixed effects are included. Errors are cluster-robust at the household level. The gray-shaded region indicates the time period used in the empirical analysis.

Table 1.B.4: Marginal propensity to borrow for all credit types

	Panel A: By LTV group			
	All credit (1)	Mortgage (2)	Consumer credit (3)	Mixed credit (4)
Low LTV, < 50%	0.458***	0.354***	0.019***	0.018**
	(0.033)	(0.031)	(0.003)	(0.007)
Mid LTV, 50–100%	0.524***	0.310***	0.036***	0.030***
	(0.044)	(0.040)	(0.005)	(0.010)
High LTV, > 100%	0.323***	0.116	0.020**	0.047***
	(0.094)	(0.087)	(0.009)	(0.017)
Low LTV – Mid LTV	−0.066	0.044	−0.017***	−0.012
High LTV – Mid LTV	−0.201**	−0.194**	−0.016*	0.017
	Panel B: By bank-deposit group			
	All credit (5)	Mortgage (6)	Consumer credit (7)	Mixed credit (8)
Bank deposits, P0–P25	0.579***	0.389***	0.029***	0.031***
	(0.041)	(0.038)	(0.005)	(0.010)
Bank deposits, P25–P50	0.619***	0.436***	0.037***	0.033***
	(0.040)	(0.039)	(0.005)	(0.009)
Bank deposits, P50–P75	0.497***	0.369***	0.019***	0.024***
	(0.040)	(0.038)	(0.004)	(0.009)
Bank deposits, P75–P100	0.199***	0.149***	0.009**	0.003
	(0.044)	(0.042)	(0.004)	(0.008)
Controls	Yes	Yes	Yes	Yes
Num. observations	6647	6647	6647	6647

Note: This table presents the results from estimating model (1.3) for homeowners in Stockholm who don't move in years when they purchase a car. The controls are 4th-order polynomials in household size, age, and disposable income. Each control variable is standardized. Errors are cluster-robust at the household level. The lower part of panel A displays the F-tests of the differences between the estimated coefficients above. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure 1.B.3: Car outcomes

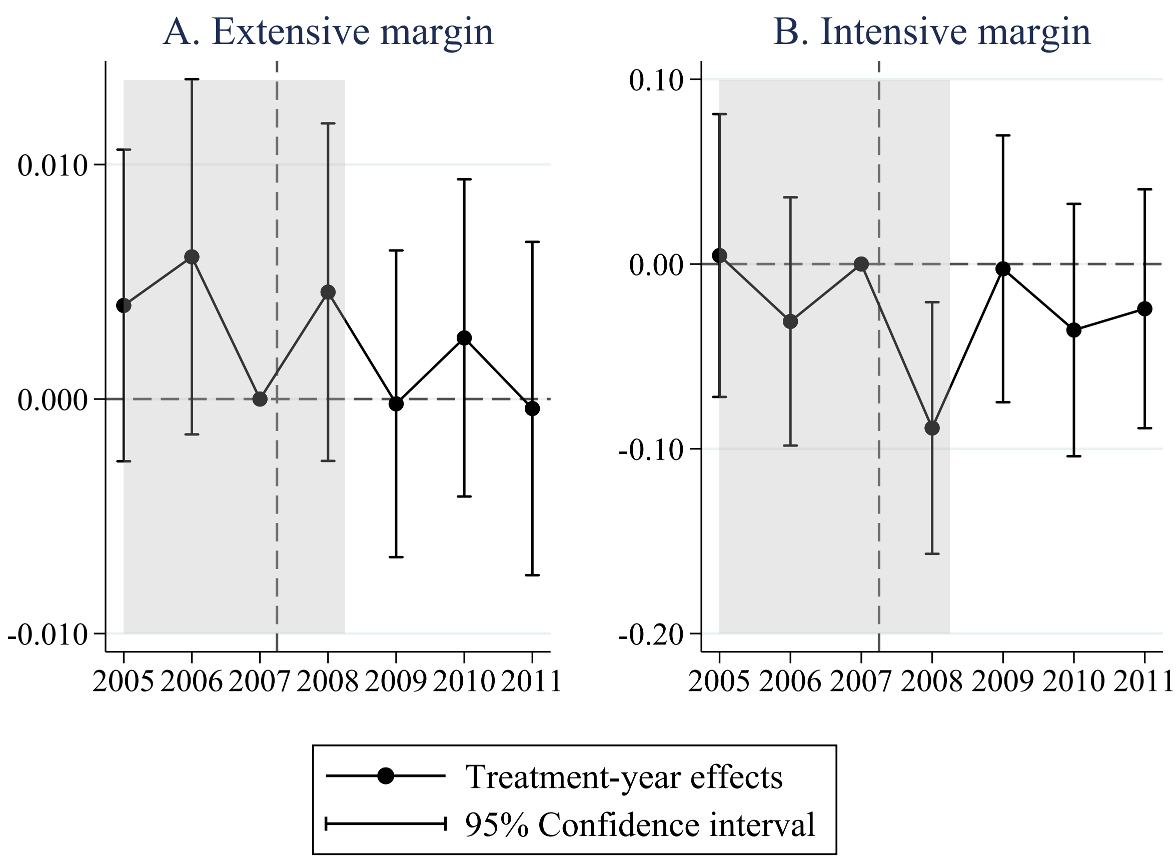

Note: This figure presents the difference in car outcomes for homeowners inside and outside the noise zone. To the left is the difference in the extensive margin, the car buying probability. To the right is the difference in the intensive margin, the logarithm of the car value at purchase. The error bands indicate the 95-percent confidence intervals.

Table 1.B.5: Triple differences – Effects on the value of new cars by groups (intensive margin)

	Loan-to-value (1)	Bank deposits (2)	Net worth (3)	Age (4)	Financial wealth (5)
noise area$_i$ × post$_t$	−0.054**	−0.023	−0.063***	−0.069**	−0.053
	(0.022)	(0.027)	(0.024)		(0.041)
noise area$_i$ × post$_t$	−0.061*				
× $\mathbb{1}(\text{LTV}_{it} \geq 50\%)$	(0.029)				
noise area$_i$ × post$_t$		−0.115***			
× $\mathbb{1}(\text{BD}_{it} \leq \text{P50})$		(0.031)			
noise area$_i$ × post$_t$			−0.058*		
× $\mathbb{1}(\text{NW}_{it} \leq \text{P25})$			(0.029)		
noise area$_i$ × post$_t$				−0.027	
× $\mathbb{1}(\text{Age}_{it} \leq \text{P25})$				(0.040)	
noise area$_i$ × post$_t$					−0.051
× $\mathbb{1}(\text{FW}_{it} \leq \text{P50})$					(0.048)
Controls	Yes	Yes	Yes	Yes	Yes
Num. obs.	6045	6045	6045	6045	6045
R-squared	0.043	0.047	0.043	0.043	0.048

Note: This table presents the estimates of our main specification (1.5) with log(car value$_{it}$) as the dependent variable. BD, NW, and FW are short for bank deposits, net worth, and financial wealth. We control for household size, household head age, and labor income, all measured in 2006. All non-collinear interactions of noise area$_i$, post$_t$, and $\mathbb{1}(\cdot)$ are included in the model but not presented above. P25 and P50 stand for percentile 25 and 50, respectively. The standard errors are two-way cluster-robust at the household and quarter levels. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.B.6: Effects on number new cars bought, by housing tenure (extensive margin)

	Single-family home owners		Co-op owners		Renters	
	(1)	(2)	(3)	(4)	(5)	(6)
noise area$_i$ × post$_t$	0.00029	0.00029	0.00059	0.00059	0.00001	0.00001
	(0.00054)	(0.00054)	(0.00038)	(0.00038)	(0.00010)	(0.00010)
noise area$_i$	0.00014	0.00012	−0.00077***	−0.00050**	−0.00047***	−0.00030***
	(0.00031)	(0.00031)	(0.00021)	(0.00020)	(0.00009)	(0.00010)
Controls	No	Yes	No	Yes	No	Yes
Num. obs.	531,105	531,105	1,837,905	1,837,905	2,959,005	2,959,005
R-squared	0.001	0.001	0.000	0.001	0.000	0.002

Note: This table presents the estimates of our main specification (1.5) with the number of cars bought as the dependent variable. We control for household size, household head age, and labor income, all measured in 2006. Errors are two-way cluster-robust at the household and quarter levels. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.B.7: Effects on the value of new cars, by housing tenure (intensive margin)

	Single-family home owners		Co-op owners		Renters	
	(1)	(2)	(3)	(4)	(5)	(6)
noise area$_i$ × post$_t$	−0.085***	−0.077***	0.007	0.009	−0.012	−0.006
	(0.019)	(0.021)	(0.019)	(0.018)	(0.039)	(0.038)
noise area$_i$	0.000	0.004	−0.029*	−0.026*	−0.018	−0.008
	(0.012)	(0.012)	(0.015)	(0.015)	(0.013)	(0.011)
Controls	No	Yes	No	Yes	No	Yes
Observations	6045	6045	11,065	11,065	9334	9334
R-squared	0.011	0.042	0.007	0.043	0.006	0.051

Note: This table presents the estimates of our main specification (1.5) with log car value as the dependent variable. We control for household size, household head age, and labor income, all measured in 2006. Errors are two-way cluster-robust at the household and quarter levels. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.B.8: Intensity of treatment on home prices and the intensive margin

	Single-family homes			
	log single-family home prices		log car value	
	(1)	(2)	(3)	(4)
$\text{post}_t \times 0$	-0.288^{***}		-0.105^{***}	
	(0.041)		(0.028)	
$\text{post}_t \times (0, 500]$	-0.216^{***}		-0.052^{**}	
	(0.055)		(0.023)	
$\text{post}_t \times (500, 1000]$	-0.143^{***}		-0.069	
	(0.052)		(0.042)	
$\text{post}_t \times -\log_{10}((1 + \text{dist}_i)^2)$		-0.053^{***}		-0.015^{***}
		(0.007)		(0.003)
Controls	Yes	Yes	Yes	Yes
Num. obs.	11,308	11,308	6045	6045
R-squared	0.331	0.334	0.043	0.042

Note: This table presents the estimates of our main specifications (1.2) and (1.5). Columns (1) and (2) use samples of house transactions, while Columns (3) and (4) use samples of households that buy cars. When estimating model (1.2) we use the controls as listed in Table 1.2, and in model (1.5) as listed in Table 1.5. Errors are cluster-robust as in Tables 1.2 and 1.5. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

1.C Appendix: Details on the household data set

Table 1.C.1: Restrictions on the household data set

	Sample restrictions	Num. households
1.	Start with households that are coded as living in Stockholm municipality and own their main property in Stockholm at the end of 2006.	43,975
2.	Keep households that live in Stockholm throughout 2004–2011.	39,700
3.	Keep households we cannot geo-locate.	39,578
4.	Keep households who own at least 50% of their home that has a positive taxation value.[a]	37,105
5.	Keep only households whose housing wealth is tied mostly in their home in Stockholm (above 50%).[b]	35,612
6.	Keep households who own less than 6 cars in one quarter.[c]	35,407

Note: The table presents the number of households that remain after imposing consecutive sample restrictions.
[a] Non-positive taxation value can indicate that the house is not used for residential purposes.
[b] This makes more certain that the household can be influenced by a wealth or collateral effect and that the household does not own significant housing outside of Stockholm.
[c] In Sweden so-called "goalkeepers" are registered owners of many cars that are in reality owned and used by other people. By being the registered owner, the goalkeeper is responsible to pay car-related taxes, fees, and fines and usually does not have income that can be confiscated to finance the costs; thus, the actual owner can escape these costs.

1.D Appendix: The life-cycle model

This section provides a mathematical formulation of the model in Section 1.4, our solution method, and the calibration, and uses the model to illustrate insights discussed in the paper.

1.D.1 Preferences, endowments and asset markets

The household lives from $t = 1$ to $t = 221$, where t denotes a quarter. Beginning of life should be thought of as 30 years of age. The household is born with a house worth w_1^h. The household has a per-period utility function

$$U(c_t, d_t) = \frac{({c_t}^{\nu} d_t^{1-\nu})^{1-\gamma}}{1-\gamma},$$

where c_t denotes non-durable consumption goods, d_t the value of cars, ν the Cobb-Douglas expenditure share on non-durables, and γ the coefficient of relative risk aversion.

Disposable income is exogenous and stochastic and follows a standard process, akin to Carroll and Samwick (1997). A working-age individual receives disposable income y_t that follows a deterministic hump-shaped life-cycle trend, $\bar{y}_t$, and is exposed to a transitory idiosyncratic income shock, ω_{it}. Disposable income cannot be less than $\underline{y}$, which is a parsimonious way to account for welfare and transfers. For $t < 140$ (for ages younger than 65),

$$y_{it} = \bar{y}_t \exp(\omega_{it}). \tag{1.D.1}$$

The transitory random variable ω_{it} is distributed

$$\omega_{it} \sim N\left(-\sigma_\omega^2/2, \sigma_\omega^2\right). \tag{1.D.2}$$

Notice that the process abstracts from permanent income shocks to economize on state variables.

Upon retirement, which happens at $t = 140$, individuals have a safe pension income. It is modeled as a deterministic replacement rate, κ, relative to permanent labor income at 64.75 years (139 quarters):

$$y_{it} = \kappa \times \bar{y}_{139}, \quad t \geq 140. \tag{1.D.3}$$

There are two financial assets: a mortgage and liquid financial savings. The current period control variables associated with these assets are denoted by $\tilde{m}_t$ and $\tilde{s}_t$, respectively. A negative value for $\tilde{s}_t$ means uncollateralized borrowing. The beginning-of-period balance, including returns, is denoted by m_t and s_t. A mortgage payment has to be made at the end of the period. The amount depends on the end-of-period mortgage balance, $\tilde{m}_t$. The function $mp_t(m)$, defined below, determines the balance.

During the period, the household chooses to adjust their mortgage and/or stock of cars and then derives utility from non-durable consumption and cars. There is an exogenous state variable that indicates a match-quality shock denoted by ζ_t that follows a Bernoulli distribution. The probability of the shock happening is ρ. The shock makes the household want to adjust d_t. After its current period decision on the car, the car value is denoted by $\tilde{d}_t$. The cost of transacting the car is $A^d(d_t)$.

The laws of motion for the endogenous state variables are

$$d_{t+1} = \tilde{d}_t(1-\delta), \tag{1.D.4}$$

$$m_{t+1} = \tilde{m}_t(1+r^m) - mp_t(\tilde{m}_t), \tag{1.D.5}$$

$$s_{t+1} = (1+r^s(\tilde{s}_t))\tilde{s}_t, \tag{1.D.6}$$

where the installment $mp_t(m)$ covers both interest and amortization so the mortgage is zero in the last period T, as in Campbell and Cocco (2003):

$$mp_t(m) = m\left(\sum_{j=1}^{T+1-t}\left(\frac{1}{1+r^m}\right)^j\right)^{-1} = m\frac{r^m(1+r^m)^{T+1-t}}{(1+r^m)^{T+1-t}-1}. \tag{1.D.7}$$

Furthermore, there are return spreads between uncollateralized borrowing, savings, and mortgage borrowing. The interest rates are ordered $r^m < r^s < r^b$, where r^m is the interest rate on the mortgage, r^s is the return on savings ($s_t > 0$), and r^b is the cost of borrowing ($s_t < 0$).

The borrowing constraints are

$$\tilde{s}_t \geq \underline{s}, \tag{1.D.8}$$

$$\tilde{m}_t \in [0, \phi\, w_t^h], \tag{1.D.9}$$

where (1.D.8) holds in every time period and $\underline{s}$ is a borrowing limit. Equation (1.D.9) holds only at the time of refinancing. The variable w_t^h is the value of the house in period t, and ϕ is a requirement on home equity at the time of refinancing. Refinancing the mortgage is associated with cost $A^m(m_t)$.

1.D.2 Dynamic programming problem

The household maximizes the value function $V_t(d_t, m_t, s_t, y_t, \zeta_t)$, which is defined by four cases:

1. No adjustment of either cars or the mortgage (no adj.),
2. Adjustment of the car stock (d adj.),
3. Adjustment of the mortgage (m adj.), and
4. Adjustment of both car and mortgage (d&m adj.).

Each case is associated with its own value function and Bellman equation.

The value function also depends on whether the household is hit by the car-purchase shock. Let $V_t(d_t, m_t, s_t, y_t, 0)$ denote the value function if the household is not hit by this shock:

$$\begin{aligned} V_t(d_t, m_t, s_t, y_t, 0) = \max \big\{ & V_t^{\text{no adj.}}(d_t, m_t, s_t, y_t),\ V_t^{d \text{ adj.}}(d_t, m_t, s_t, y_t), \\ & V_t^{m \text{ adj.}}(d_t, m_t, s_t, y_t),\ V_t^{d\&m \text{ adj.}}(d_t, m_t, s_t, y_t) \big\}. \end{aligned}$$

Let $V_t(d_t, m_t, s_t, y_t, 1)$ denote the value function if the household is hit by the shock:

$$V_t(d_t, m_t, s_t, y_t, 1) = \max \big\{ V_t^{d \text{ adj.}}(d_t, m_t, s_t, y_t),\ V_t^{d\&m \text{ adj.}}(d_t, m_t, s_t, y_t) \big\}.$$

The Bellman equations and associated budget constraints are as follows for periods $t < T$:

1.D.2.1 No adjustment

$$V_t^{\text{no adj.}}(d_t, m_t, s_t, y_t) = \max_{\tilde{c}_t, \tilde{s}_t} \frac{(\tilde{c}_t^{\nu} d_t^{1-\nu})^{1-\gamma}}{1-\gamma} + \beta\, \mathbb{E}_t[V_{t+1}(d_{t+1}, m_{t+1}, s_{t+1}, y_{t+1}, \zeta_{t+1})]$$

subject to (1.D.1–1.D.8) and $\tilde{c}_t + \tilde{s}_t = s_t + y_t - mp_t(m_t)$.

1.D.2.2 Adjusting only the car

$$V_t^{d\text{ adj.}}(d_t, m_t, s_t, y_t) = \max_{\tilde{c}_t, \tilde{d}_t, \tilde{s}_t} \frac{(\tilde{c}_t^{\nu} \tilde{d}_t^{1-\nu})^{1-\gamma}}{1-\gamma} + \beta\, \mathbb{E}_t[V_{t+1}(d_{t+1}, m_{t+1}, s_{t+1}, y_{t+1}, \zeta_{t+1})],$$

subject to (1.D.1–1.D.8) and $\tilde{c}_t + \tilde{s}_t + p\,\tilde{d}_t = s_t + y_t - mp_t(m_t) + p\,d_t - A^d(d_t)$.

1.D.2.3 Adjusting only the mortgage

$$V_t^{m\text{ adj.}}(d_t, m_t, s_t, y_t) = \max_{\tilde{c}_t, \tilde{m}_t, \tilde{s}_t} \frac{(\tilde{c}_t^{\nu} d_t^{1-\nu})^{1-\gamma}}{1-\gamma} + \beta\, \mathbb{E}_t[V_{t+1}(d_{t+1}, m_{t+1}, s_{t+1}, y_{t+1}, \zeta_{t+1})],$$

subject to (1.D.1–1.D.9) and $\tilde{c}_t + \tilde{s}_t = s_t + y_t - m_t + \tilde{m}_t - mp_t(\tilde{m}_t) - A^m(m_t)$.

1.D.2.4 Adjusting both the car and the mortgage

Lastly, we consider the case to adjust both the stock of cars and the mortgage:

$$V_t^{d\&m\text{ adj.}}(d_t, m_t, s_t, y_t) = \max_{\tilde{c}_t, \tilde{d}_t, \tilde{m}_t, \tilde{s}_t} \frac{(\tilde{c}_t^{\nu} \tilde{d}_t^{1-\nu})^{1-\gamma}}{1-\gamma} + \beta\, \mathbb{E}_t[V_{t+1}(d_{t+1}, m_{t+1}, s_{t+1}, y_{t+1}, \zeta_{t+1})],$$

subject to (1.D.1–1.D.9) and

$$\begin{aligned}\tilde{c}_t + \tilde{s}_t + p\,\tilde{d}_t = s_t + y_t + p\,d_t - A^d(d_t) \\ - m_t + \tilde{m}_t - A^m(m_t) - mp_t(\tilde{m}_t).\end{aligned}$$

1.D.2.5 The last time period

In the last time period, $t = T$, the household derives utility from a bequest motive, $\Psi(b)$. The household also sells the house:

$$\Psi(b) = \Psi_0 \frac{b^{1-\gamma}}{1-\gamma}.$$

It is not possible to take out a mortgage, and the car can be used in the last period and is then sold (depreciated by one more period). The value functions are as follows:

No adjustment in the last time period

$$\begin{aligned} V_T^{\text{no adj.}}(d_T, m_T, s_T, y_T) = \max_{\tilde{c}_T, \tilde{b}_T} \; & \frac{(\tilde{c}_T^{\nu} d_T^{1-\nu})^{1-\gamma}}{1-\gamma} + \Psi(\tilde{b}_T + p\, d_T(1-\delta)), \\ \text{subject to} \quad & \tilde{c}_T + \tilde{b}_T = s_T + y_T - mp_T(m_T) + w_T^h. \end{aligned}.$$

Adjusting the car in the last time period

$$\begin{aligned} V_T^{d\text{ adj.}}(d_T, m_T, s_T, y_T) = \max_{\tilde{c}_T, \tilde{d}_T, \tilde{b}_T} \; & \frac{(\tilde{c}_T^{\nu} \tilde{d}_T^{1-\nu})^{1-\gamma}}{1-\gamma} + \Psi(\tilde{b}_T + p\, \tilde{d}_T(1-\delta)), \\ \text{subject to} \quad & \tilde{c}_T + p\, \tilde{d}_T + \tilde{b}_T = s_T + y_T + p\, d_T - mp_T(m_T) + w_T^h - A^d(d_T). \end{aligned}$$

1.D.3 Solution method

We solve the model using nested value function iteration (Druedahl, 2021).[47] The technique exploits a nested structure of value functions: once we have solved for $V_t^{\text{no adj.}}(d_t, m_t, s_t, y_t, \zeta_t)$ over a grid, we can quickly compute $V_t^{d\text{ adj.}}(d_t, m_t, s_t, y_t, \zeta_t)$ by solving

$$V_t^{d\text{ adj.}}(d_t, m_t, s_t, y_t, \zeta_t) = \max_{\tilde{d}} \; V_t^{\text{no adj.}}(\tilde{d}, m_t, s_t + p\, d_t - A^d(d_t) - p\, \tilde{d}, y_t, \zeta_t)$$

[47] Because it is a life-cycle model, we actually use backward induction and not iteration, but the method is known as nested value function *iteration*.

by interpolation. If the optimal new car level is $\tilde{d}^*$, the other optimal decisions are given by the interpolations

$$\begin{aligned}
\tilde{c}^{d\text{ adj.}}(d_t, m_t, s_t, y_t, \zeta_t) &= \tilde{c}^{\text{no adj.}}(\tilde{d}^*, m_t, s_t + p\, d_t - A^d(d_t) - p\, \tilde{d}^*, y_t, \zeta_t), \qquad \forall t, \\
\tilde{s}^{d\text{ adj.}}(d_t, m_t, s_t, y_t, \zeta_t) &= \tilde{s}^{\text{no adj.}}(\tilde{d}^*, m_t, s_t + p\, d_t - A^d(d_t) - p\, \tilde{d}^*, y_t, \zeta_t), \qquad \forall t, \\
\tilde{b}^{d\text{ adj.}}(d_T, m_T, s_T, y_T, \zeta_T) &= \tilde{b}^{\text{no adj.}}(\tilde{d}^*, m_T, s_T + p\, d_T - A^d(d_T) - p\, \tilde{d}^*, y_T, \zeta_T).
\end{aligned}$$

We can solve the case of adjusting the mortgage and the case of adjusting both cars and the mortgage in a similar way.

Furthermore, we follow Druedahl (2021) in that we reduce the dimensionality of the state space by summing savings, uncollateralized borrowing, and income into the variable cash-on-hand, $x_{t+1} \equiv \tilde{s}_t(1 + r(\tilde{s}_t)) + y_{t+1}$. Each case of adjusting one variable at a cost also makes it possible to reduce the state space by one state: when changing the car stock, the car is sold, and the adjustment cost is paid (since it only depends on today's initial car stock), and what is left is added to that period's cash-on-hand, x_t^d. In this case, we solve over a grid over (x_t^d, m_t); the initial d_t only affects the decision $\tilde{d}_t$ through its contribution to the budget constraint. Correspondingly, when adjusting the mortgage, we solve over a grid (x_t^m, d_t) where $x_t^m \equiv x_t - m_t - A^m(m_t)$. In the case of adjusting both the level of cars and the mortgage, we define $x_t^{d,m} \equiv x_t - m_t - A^m(m_t) + p\, d_t - A^d(d_t)$ and solve over a grid over only this measure of cash-on-hand (i.e., both states d_t and m_t can be dropped).

1.D.4 Calibration

For most parameters, we use standard values or values consistent with the Swedish institutional setting. Those parameter values are reported in Table 1.D.1. The adjustment costs for the durable good and the mortgage are worth discussing.

For cars, we have studied the pricing model of used car dealers and noted that they discount their procurement price by a relative amount.[48] Thus, we assume that the adjustment cost

[48] Our source is a 2006 catalog produced by Autodata that published valuations of used cars in Sweden.

of cars is proportional to the car value, $A^d(d) = \eta_1 d$, and we set that value to 15 percent. For mortgages, we assume that the adjustment cost is constant as in e.g., Berger et al. (2021); Eichenbaum et al. (2022): $A^m(m) = \mu_0 =$ SEK 18,400.

1.D.4.1 Fitted parameters

We fit five parameter values to target moments. The time discount factor (β) and the bequest parameter (Ψ_0) are chosen to fit liquid financial wealth (i.e., financial wealth outside the pension system) during the working phase and the retirement phase. The Cobb-Douglas expenditure share (ν) is chosen to match an approximate value of the car stock, including used and new cars. Based on car advertisements, we approximate the average value of used cars in 2006 to SEK 75,274. For new cars, our micro data set indicates a value of SEK 220,200. Given the average transaction frequencies of used and new cars, we match a target value of SEK 100,030. We compute the car transaction probability in the data set, 5.365 percent per quarter, and set the probability of exogenous match quality shocks, ρ, to match the overall frequency in the data. About half of the car adjustments are determined by the exogenous shock. The probability of information arrival, λ, is set in accordance with McKay and Wieland (2021) and Coibion and Gorodnichenko (2012).

The parameter values are reported in Table 1.D.2 and the target moments in Table 1.D.3. Below, we provide further details on how we fit the model.

1.D.4.2 Method

Fitting the model is a two-stage process. The first stage consists of the following steps:

We make an initial guess of the parameters (step 0), solve the model at the initial housing wealth w^h (step 1), and simulate the model for a large sample of households (step 2). After this, we compute the moments of the simulated data that we are targeting (step 3), and test if the total error is below a tolerance level (step 4). If it is, we continue to stage 2; otherwise, we make a new guess (step 5).

Table 1.D.1: Calibration

Description	Parameter	Value	Source
Panel A: Preferences and income			
Risk aversion	γ	2.0	Standard
Income profile	$\bar{y}_t$	–	Dahlquist et al. (2018)
Transitory income risk	σ_ω	0.178	Micro data
Panel B: Financial savings and uncollateralized borrowing			
Interest rate on savings	r^s	1.337%	Sveriges Riksbank
Interest rate on borrowing	r^b	1.692%	Statistics Sweden
Borrowing constraint	$\underline{s}$	−200.0	Sodini et al. (2023)
Panel C: Cars			
Relative price	p	1.0	Attanasio et al. (2022)
Depreciation rate	δ	3.599%	Autodata
Adjustment cost	η_1	0.15	Autodata
Panel D: Mortgages			
Interest rate	r^m	1.103%	Statistics Sweden
Borrowing constraint	ϕ	85%	Swedish bank norm
Adjustment cost	μ_0	18.4	Eichenbaum et al. (2022) ($2,100) and Berger et al. (2021) ($2,500)
Panel E: Housing wealth			
Before announcement	w^h	3218.0	Microdata
After announcement	$w^{h'}$	3842.3	Microdata

Note: This table presents parameter values determined by historical values and institutional features. The average income per period, $\bar{y}_t$, follows the age profile in Dahlquist et al. (2018), scaled to match the average of our sample. We calculate the adjustment cost and depreciation rate for cars from Autodata price catalogs. Interest rates are reported at a quarterly frequency. Amounts are in SEK 1000.

The new guess is computed by studying the simulated error of each moment and the percentage deviation. We associate each moment with a particular parameter. The level of savings of younger households are associated with the time discount factor β; the savings of older households are associated with the bequest motive Ψ_0; the average level of the car stock with the non-car preference parameter ν; and the car purchasing frequency with the match-quality shock frequency ρ. If the average car stock is above the life-cycle target, we increase $\nu^{\text{new guess}}$ by some relative amount proportional to the percentage deviation of the moment. (An increase in the non-car preference parameter makes cars less desirable, and the household will thus on average buy less cars). Likewise, if young savings are too high, we decrease the time discount factor β; if the saving of the old is too high, we decrease the bequest motive Ψ_0; if the frequency of car purchases is too high, we decrease ρ.

Then we return to step 1 and continue from there. Once the test in step 4 is passed, the first stage is completed, and we continue to stage 2.

In stage 2, we set up a simplex in the parameter space around the set of values obtained from stage 1. We pass the simplex to an optimizer that uses the Nelder-Mead method that goes through steps 1 to 4 above using a lower tolerance level. Step 5 is replaced by the internal updating method of the Nelder-Mead method.

The value for one parameter requires an even more elaborate process. The frequency for information arrival, λ, requires a change in the house price to be identified. Using the bisection method, we compute the λ that makes the households' average time to update their information six quarters (as in Coibion and Gorodnichenko, 2012; McKay and Wieland, 2021). At a random time, we shock the household by changing the current housing wealth from the initial value to $w^{h'}$; the household is not aware of this until they update their information set. There are now three cases when they update: when deciding it is optimal and using their 'un-updated' information set, (1) to buy a new car, (2) to refinance the mortgage, or (3) when they are hit by the information-arrival process. The parameter λ controls the frequency of the latter and is tuned to make the average time for an update

Table 1.D.2: Fitted parameters

Description	Parameter	Value
Time discount factor	β	0.981
Bequest parameter	Ψ_0	533.2
Preference parameter	ν	0.945
Probability of match shock	ρ	0.029
Probability of information arrival	λ	0.063

Note: This table presents the fitted parameters (see Section 1.D.4.1).

Table 1.D.3: Target moments

Summary statistic	Target value	Simulated value
Financial wealth (s_t), age 35–65	918.59	922.88
Financial wealth (s_t), age 65–85	1478.12	1468.77
Value of car stock (d_t)	100.30	100.09
Car purchasing freq. (%)	5.385	5.365
Time to update after house price shock	6.000	6.006

Note: This table presents the targets of the estimation using the method outlined in Section 1.D.4.1. Savings and car stock are in SEK 1000; the car purchasing frequency is quarterly; and the time to update information after the house price shock is stated in quarters and represents the average time. This statistic is based on McKay and Wieland (2021) and Coibion and Gorodnichenko (2012).

six quarters after the shock. Given that the quasi-experiment occurs in normal times, just before the U.S. financial crisis spread globally, we find this frequency to be reasonable.

1.D.5 Computing the marginal propensity for car expenditures

While the model makes no distinction between new and used cars, we only observe values for new cars in our empirical analysis. Therefore, we have to make an assumption about the new-car share of total car expenditures to be able to compare the new car MPX in the data with the car MPX in the model.

We assume that households reduce expenditures of each type of car proportionally to their expenditure share. To approximate the share of expenditure spent on new cars we rely on statistics from webcar2000.com.[49] We find that the average value in used car ads in 2006 was SEK 75,274. New cars were about SEK 220,200 (computed from our micro data set). Given the probabilities of buying used and new cars, the total share of new cars out of all car expenditures is 0.315.

The black circle in Panel B of Figure 1.4 relies on this share. It is a conversion of our new-car MPX estimate to an all-car MPX of 0.38 cents per dollar (0.0012/0.315). The 95-percent confidence interval is 0.146–0.620 cents per dollar, which we obtain from a conversion of the 95-percent confidence interval of the elasticity estimate, 0.150–0.636, in Column (2) of Table 1.6.

1.D.6 The role of the measurement period and the regression specification

In the main text, we emphasize that the temporal dynamics of car purchases have important implications for the appropriate empirical design. Below, we expand on this.

1.D.6.1 The role of the measurement period

To illustrate the role of the length of the measurement period, we estimate reduced form responses based on the simulated paths:

$$y_{it} = \alpha + \phi\,\text{treated}_i \times \text{post}_t + \gamma\,\text{post}_t + \rho\,\text{treated}_i + \varepsilon_{it}, \tag{1.D.10}$$

where y_{it} denotes one of the outcome variables, displayed in Figure 1.3. This regression specification corresponds to the empirical difference-in-difference model (1.5). We present estimates of ϕ in Table 1.D.4 for different lengths of the post-period. Panel A is based on a post-period of four quarters, as in our empirical setting. Columns (1) and (2) show that the total spending response in the first year is SEK 1,836 per quarter and the car spending

[49]Source: http://www.webcar2000.com/countries/sweden/car/statistics.phtml.

response is SEK 996 per quarter, respectively. Despite cars' small expenditure share on average (6.9 percent in the simulation), they dominate the change in household expenditures in the first year, with a response share of 54 percent. Column (3) shows a small extensive margin response (-0.3 percentage points per quarter relative to a baseline probability of 5.37 percent per quarter).[50] We also see in Column (4) that the intensive margin response over the four following quarters is 6.1 log points, close to our empirical estimates of 7.7–8.5. The house price divergence between the treatment and control groups corresponds to a relative loss in housing wealth of SEK 624,200. This implies that the (per-year) car MPX in the first year is 0.64 cents per dollar ($996 \times 4/624{,}200$). Given that 31.5 percent of all car expenditure is on new cars (see Appendix 1.D.5), the model gives a new-car MPX in the first year of 0.20 (0.64×0.315) cents per dollar. This is close to our empirical estimate of 0.12 cents per dollar.

The four-quarter results of Panel A are the most relevant for a comparison with our quasi-experiment. However, to complete the picture of the model dynamics, Panel B of Table 1.D.4 reports estimates for longer post-periods. The estimates align with the dynamics in Figure 1.3; the change in total expenditures does not peak until 6–8 quarters after the shock, while cars' expenditure share falls. The reversion of car expenditures is driven by the intensive margin as extensive margin responses are minute relative to the baseline. We will discuss this reversal below.

1.D.6.2 The role of the regression specification

The regression specification of Mian et al. (2013) and Aruoba et al. (2022) is different from ours. The outcome variable is the three-year difference (2009 versus 2006) in county-level expenditures: $\Delta_q c_{jt_0} \equiv c_{jt_0+q} - c_{jt_0}$, where j denotes the county and q the length between

[50] We find no extensive margin response in the empirical analysis. This can be due to a lack of statistical power. If we compute the standard error of the model estimate for a simulated sample of the same size as the empirical sample, the t-statistic of a Welch's t-test is -1.84. We estimate model (1.5) with the outcome being a dummy variable $\mathbb{1}(\text{num. bought cars}_{it} > 0)$. The estimate is 0.00071, and the standard error is 0.00109. The re-scaled standard error from the model when using a sample of the size of the empirical sample is 0.00171, which is similar. The empirical outcome is not unlikely enough to be rejected by the model outcome using conventional levels of statistical significance.

Table 1.D.4: Reduced form estimates for different post measurement periods (model)

	Panel A: 4 quarters post period			
	Total expenditure (1)	Car expenditure (2)	Prob. to buy car (3)	log car value | car purch. (4)
$\text{treated}_i \times \text{post}_t$	-1.836*** (0.085)	-0.996*** (0.081)	-0.003*** (0.001)	-0.061*** (0.001)
Num. observations	3,200,000	3,200,000	3,200,000	173,181

	Panel B: Longer post periods			
	16 quarters		32 quarters	
	Total expenditure (5)	Car expenditure (6)	Total expenditure (7)	Car expenditure (8)
$\text{treated}_i \times \text{post}_t$	-2.071*** (0.068)	-0.643*** (0.064)	-1.556*** (0.063)	-0.168*** (0.063)
Num. observations	8,000,000	8,000,000	14,400,000	14,400,000

Note: This table presents the estimated coefficients of model (1.D.10) using simulated data. The length of the post-period used after the shock to house prices is indicated by the column head; the length of the pre-period is irrelevant but we use 4 periods. For each outcome, we simulate 200,000 households twice (one treated- and one control-copy of each). All amounts are in SEK 1000. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.D.5: MPXs based on different regression specifications

	Panel A: Total expenditures					
	Housing wealth shock -30%			Housing wealth shock -10%		
Post-period length q	4	8	16	4	8	16
$\text{treated}_i \times \text{post}_{t,q}/\Delta_q w^h$	0.0075	0.0079	0.0091	0.0113	0.0111	0.0116
$\Delta_q w^h_{jt_0}$	0.0065	0.0089	0.0110	0.0099	0.0115	0.0109

	Panel B: Car expenditures					
	Housing wealth shock -30%			Housing wealth shock -10%		
Post-period length q	4	8	16	4	8	16
$\text{treated}_i \times \text{post}_{t,q}/\Delta_q w^h$	0.0042	0.0029	0.0021	0.0067	0.0045	0.0030
$\Delta_q w^h_{jt_0}$	0.0016	0.0014	0.0010	0.0032	0.0020	-0.0001

Note: This table presents the different estimates of our model (1.D.10), and the alternative model (1.D.11), using simulated data with different housing wealth shocks and different post-period lengths.

the observations at t_0 and $t_0 + q$. Their regression specification can be written

$$\Delta_q c_{jt_0} = \beta \, \Delta_q \text{home prices}_{jt_0} + \eta \, X_{jt_0} + \varepsilon_{jt_0},$$

where X_{jt_0} is a set of control variables and β is an estimate of the MPC. In Mian et al. (2013), $t_0 = 2006$, $q = 12$ (i.e., 3 years). They also instrument for $\Delta_q \text{home prices}_{jt_0}$. In our model-based regression, the specification translates to

$$\Delta_q c_{t_0} = \beta \, \Delta_q w^h_{t_0} + \varepsilon_{t_0}. \tag{1.D.11}$$

In contrast, our specification measures the average change over the measurement period. Table 1.D.5 illustrates the role of the regression specification by reporting estimates on simulated data based on regression specification (1.D.10) and (1.D.11).

References

Adda, Jérôme and Russell Cooper (2000) "Balladurette and Juppette: A discrete analysis of scrapping subsidies," *Journal of Political Economy*, 108 (4), 778–806.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino (2016) "Loan originations and defaults in the mortgage crisis: The role of the middle class," *Review of Financial Studies*, 29 (7), 1635–1670.

Aladangady, Aditya (2017) "Housing wealth and consumption: Evidence from geographically linked microdata," *American Economic Review*, 107 (11), 3415–3446.

Andersen, Henrik Yde and Søren Leth-Petersen (2020) "Housing wealth or collateral: How home value shocks drive home equity extraction and spending," *Journal of the European Economic Association*, 19 (1), 403–440, 10.1093/jeea/jvz083.

Angrist, Joshua and Alan B. Krueger (1992) "The effect of age at school entry on educational attainment: An application of instrumental variables with moments from two samples," *Journal of the American Statistical Association*, 87, 328–336.

Aruoba, S. Borağan, Ronel Elul, and Ṣebnem Kalemli-Özcan (2022) "Housing wealth and consumption: The role of heterogeneous credit constraints," Federal Reserve Bank Philadelphia Working Paper No. 22-34.

Attanasio, Orazio, Kieran Larkin, Morten O Ravn, and Mario Padula (2022) "(S)cars and the Great Recession," *Econometrica*, 90 (5), 2319–2356.

Attanasio, Orazio P. (2000) "Consumer durables and inertial behaviour: Estimation and aggregation of (S, s) rules for automobile purchases," *Review of Economic Studies*, 67 (4), 667–696.

Attanasio, Orazio P, Laura Blow, Robert Hamilton, and Andrew Leicester (2009) "Booms and busts: Consumption, house prices and expectations," *Economica*, 76 (301), 20–50.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2020) "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," *NBER Working Paper* Number 26647.

Aydin, Deniz (2022) "Consumption response to credit expansions: Evidence from experimental assignment of 45,307 credit lines," *American Economic Review*, 112 (1), 1–40, 10.1257/aer.20191178.

Baker, Scott R. (2018) "Debt and the response to household income shocks: Validation and application of linked financial account data," *Journal of Political Economy*, 126 (4), 1504–1557.

Bar-Ilan, Avner and Alan S. Blinder (1992) "Consumer durables: Evidence on the optimality of usually doing nothing," *Journal of Money, Credit and Banking*, 24 (2), 258–272.

Benmelech, Efraim, Ralf R. Meisenzahl, and Rodney Ramcharan (2017) "The real effects of liquidity during the Financial Crisis: Evidence from automobiles," *Quarterly Journal of Economics*, 132 (1), 317–365.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra (2018) "House prices and consumer spending," *The Review of Economic Studies*, 85 (3), 1502–1542.

Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra (2021) "Mortgage prepayment and path-dependent effects of monetary policy," *American Economic Review*, 111 (9), 2829–78.

Berger, David and Joseph Vavra (2015) "Consumption dynamics during recessions," *Econometrica*, 83 (1), 101–154.

Bernanke, Ben S. (2018) "The real effects of disrupted credit: Evidence from the Global Financial Crisis," *Brookings Papers on Economic Activity*, Fall, 251–342.

Bertola, Giuseppe, Luigi Guiso, and Luigi Pistaferri (2005) "Uncertainty and consumer durables adjustment," *Review of Economic Studies*, 72 (4), 973–1007.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess (2024) "Revisiting event study designs: Robust and efficient estimation," *Review of Economic Studies*, rdae007.

Browning, Martin, Mette Gørtz, and Søren Leth-Petersen (2013) "Housing wealth and consumption: A micro panel study," *The Economic Journal*, 123, 401–428.

Caballero, Ricardo J. (1993) "Durable goods: An explanation for their slow adjustment," *Journal of Political Economy*, 101 (2), 351–384.

Calomiris, Charles W, Stanley D Longhofer, William Miles et al. (2013) "The housing wealth effect: The crucial roles of demographics, wealth distribution and wealth shares," *Critical Finance Review*, 2 (1), 049–099.

Campbell, John Y and Joao F Cocco (2003) "Household risk management and optimal mortgage choice," *The Quarterly Journal of Economics*, 118 (4), 1449–1494.

Campbell, John Y. and Joao F. Cocco (2007) "How do house prices affect consumption? Evidence from micro data," *Journal of Monetary Economics*, 54, 591–621.

Carroll, Christopher D, Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N White (2020) "Sticky expectations and consumption dynamics," *American Economic Journal: Macroeconomics*, 12 (3), 40–76.

Carroll, Christopher D, Misuzu Otsuka, and Jiri Slacalek (2011) "How large are housing and financial wealth effects? A new approach," *Journal of Money, Credit and Banking*, 43 (1), 55–79.

Carroll, Christopher D. and Andrew A. Samwick (1997) "The nature of precautionary wealth," *Journal of Monetary Economics*, 40, 41–71.

Case, Karl E., John M. Quigley, and Robert J. Shiller (2013) "Wealth effects revisited 1975–2012," *Critical Finance Review*, 2, 101–128.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J Notowidigdo (2019) "Housing booms, manufacturing decline and labour market outcomes," *The Economic Journal*, 129 (617), 209–248.

Chay, Kenneth Y. and Michael Greenstone (2005) "Does air quality matter? Evidence from the housing market," *Journal of Political Economy*, 113 (2), 376–424.

Cloyne, James, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven (2019) "The effect of house prices on household borrowing: A new approach," *American Economic Review*, 109 (6), 2104–36.

Coibion, Olivier and Yuriy Gorodnichenko (2012) "What can survey forecasts tell us about information rigidities?" *Journal of Political Economy*, 120 (1), 116–159.

Cooper, Daniel (2013) "House price fluctuations: The role of housing Wealth as borrowing collateral," *Review of Economics and Statistics*, 95 (4), 1183–1197.

Correia, Sergio (2016) "Linear models with high-dimensional fixed effects: An efficient and feasible estimator," Working Paper, http://scorreia.com/research/hdfe.pdf, Board of Governors of the Federal Reserve System.

Currie, Janet, Lucas Davis, Michael Greenstone, and Reed Walker (2015) "Environmental health risks and housing Values: Evidence from 1,600 toxic plant openings and closings," *American Economic Review*, 105 (2), 678–709.

Dahlquist, Magnus, Ofer Setty, and Roine Vestman (2018) "On the asset allocation of a default pension fund," *The Journal of Finance*, 73 (4), 1893–1936.

Davidoff, Thomas (2016) "Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors," *Critical Finance Review*, 5 (2), 177–206.

DeFusco, Anthony (2018) "Homeowner borrowing and housing collateral: New evidence from expiring price controls," *Journal of Finance*, 73 (2), 523–573.

Disney, Richard, John Gathergood, and Andrew Henley (2010) "House price shocks, negative equity, and household consumption in the United Kingdom," *Journal of the European Economic Association*, 8 (6), 1179–1207.

Druedahl, Jeppe (2021) "A guide on solving non-convex consumption-saving models," *Computational Economics*, 58 (3), 747–775.

Eberly, Janice C. (1994) "Adjustment of consumers' durables stocks: Evidence from automobile purchases," *Journal of Political Economy*, 102 (3), 403–436.

Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong (2022) "State-dependent effects of monetary policy: The refinancing channel," *American Economic Review*, 112 (3), 721–61.

Fagereng, Andreas, Martin B Holm, and Gisle J Natvik (2021) "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 13 (4), 1–54.

Finansinspektionen (2010) "Den svenska bolånemarknaden och bankernas kreditgivning," Government report.

Fredriksson, Peter and Björn Öckert (2013) "Life-cycle effects of age at school start," *The Economic Journal*, 124, 977–1004.

Gavazza, Alessandro and Andrea Lanteri (2021) "Credit shocks and equilibrium dynamics in consumer durable goods markets," *The Review of Economic Studies*, 88 (6), 2935–2969.

Graham, James and Christos A. Makridis (2023) "House prices and consumption: A new instrumental variables approach," *American Economic Journal: Macroeconomics*, 15 (1), 411–443.

Grodecka-Messi, Anna, Jieying Li, and Xin Zhang (2022) "Riding the housing wave: Home equity withdrawal and consumer debt composition,"Technical report, The Riksbank, Working Paper.

Guerrieri, Luca and Matteo Iacoviello (2017) "Collateral constraints and macroeconomic asymmetries," *Journal of Monetary Economics*, 90, 28–49.

Guren, Adam M, Alisdair McKay, Emi Nakamura, and Jón Steinsson (2021) "Housing wealth effects: The long view," *The Review of Economic Studies*, 88 (2), 669–707.

Guren, Adam, Alisdair McKay, Emi Nakamura, and Jón Steinsson (2021) "What do we learn from cross-regional empirical estimates in macroeconomics?" *NBER Macroeconomics Annual*, 35 (1), 175–223.

Guski, Rainer, Dirk Schreckenberg, and Rudolf Schuemer (2017) "WHO environmental noise guidelines for the European region: A systematic review on environmental noise and annoyance," *International journal of environmental research and public health*, 14 (12), 1539.

Hassler, John (2001) "Uncertainty and the timing of automobile purchases," *Scandinavian Journal of Economics*, 103 (2), 351–366.

Inoue, Atsushi and Gary Solon (2010) "Two-sample instrumental variables estimators," *The Review of Economics and Statistics*, 92 (3), 557–561.

Kaplan, Greg, Kurt Mitman, and Giovanni L Violante (2020) "Non-durable consumption and housing net worth in the Great Recession: Evidence from easily accessible data," *Journal of Public Economics*, 189, 104176.

Kaplan, Greg and Giovanni L Violante (2022) "The marginal propensity to consume in heterogeneous agent models," *Annual Review of Economics*, 14, 747–775.

Kommunikationsdepartementet (1996) "Förutsättningar för avveckling av Bromma flygplats," Regeringskansliets Offsetcentral, DS 1996:12.

Laibson, David, Peter Maxted, and Benjamin Moll (2022) "A simple mapping from MPCs to MPXs," NBER Working Paper No. 29664.

Lam, Pok-Sang (1991) "Permanent income, liquidity, and adjustments of automobile stocks: Evidence from panel data," *Quarterly Journal of Economics*, 106 (1), 203–230.

Leth-Petersen, Søren (2010) "Intertemporal consumption and credit constraints: Does total expenditure respond to an exogenous shock to credit?" *American Economic Review*, 100 (3), 1080–1103, 10.1257/aer.100.3.1080.

Luftfartsverket (2000) "Flygplatskapaciteten i Stockholmsregionen: sammandrag av LFVs slutredovisning januari 2000," LFV, Sweden.

McCully, Brett A, Karen M Pence, and Daniel J Vine (2019) "How much are car purchases driven by home equity withdrawal?" *Journal of Money, Credit and Banking*, 51 (5), 1403–1426.

McKay, Alisdair and Johannes F Wieland (2021) "Lumpy durable consumption demand and the limited ammunition of monetary policy," *Econometrica*, 89 (6), 2717–2749.

Mian, Atif, Kamalesh Rao, and Amir Sufi (2013) "Household balance sheets, consumption, and the Economic Slump," *Quarterly Journal of Economics*, 128 (4), 1687–1726.

Mian, Atif and Amir Sufi (2014) "What explains the 2007-2009 drop in employment?" *Econometrica*, 82 (6), 2197–2223.

——— (2018) "Finance and business cycles: The credit-driven household demand channel," *Journal of Economic Perspectives*, 32 (3), 31–58.

Miedema, Henk M.E. and Catharina G.M. Oudshoorn (2001) "Annoyance from transportation noise: Relationships with exposure metrics DNL and DENL and their confidence intervals," *Environmental health perspectives*, 109 (4), 409–416.

Miljödomstolen (2006) "Akt M1414-07, Aktbilaga 1107, bilaga 7," Miljödomstolen, Nacka tingsrätt.

Miljööverdomstolen (2010) "Luftfartsverket and Naturvårdsverket v. Länsstyrelsen i Stockholms län," February, Svea Court of Appeal in Stockholm, case 1441-09.

Murphy, Kevin M. and Robert H. Topel (2002) "Estimation and inference in two-step econometric models," *Journal of Business & Economic Statistics*, 20 (1), 88–97.

Orchard, Jacob, Valerie A Ramey, and Johannes Wieland (2022) "Micro MPCs and macro counterfactuals: The case of the 2008 Rebates," Working Paper.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and III Owens, Raymond (2010) "Housing Externalities," *Journal of Political Economy*, 118 (3), 485–535, 10.1086/653138.

Schiraldi, Pasquale (2011) "Automobile replacement: a dynamic structural approach," *RAND Journal of Economics*, 42 (2), 266–291.

Sodini, Paolo, Stijn Van Nieuwerburgh, Roine Vestman, and Ulf von Lilienfeld-Toal (2023) "Identifying the benefits from homeownership: A Swedish experiment," *American Economic Review*, 113 (12), 3173–3212.

Stockholm (1989) "Bromma flygfält – en utredning om Brommafältets framtid.," Stockholms stad.

Vestman, Roine (2019) "Limited stock market participation among renters and homeowners," *The Review of Financial Studies*, 32 (4), 1494–1535.

Wasi, Nada and Aaron Flaaen (2015) "Record linkage using Stata: Preprocessing, linking, and reviewing utilities," *The Stata Journal*, 15 (3), 672–697.

CHAPTER 2

Should I stay or should I go? The role of housing in understanding limited inter-regional worker mobility

Solo-authored[1]

Workers exhibit lower leaving responses following adverse regional shocks than predicted by spatial models where regional differences in local earnings drive migration. While Glaeser and Gyourko (2005) attribute this to falling housing costs benefiting renters, making the shocked location more attractive, which lowers overall out-migration, declining rents also generate moving heterogeneity (Notowidigdo, 2020). However, many places that experience economic shocks have high homeownership rates. When home prices fall, homeowners do not necessarily benefit from this rent channel but endure housing wealth erosion.

In this paper, I study how homeownership and falling home prices depress the out-migration incentive from a region that experiences an adverse shock. I argue that the housing wealth losses homeowners face diminish other locations' desirability because the mover can purchase less housing and consumption after a move following the shock. This housing wealth effect

[1]This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the author and do not necessarily reflect those of Norges Bank. For helpful comments and discussions, I thank Andy Atkeson, Ariel Burnstein, Chandni Raja, Daniele Caratelli, Gregor Schubert, Joan Monras, Joao Guerreiro, John Mondragon, Jonathan Vogel, Kasper Roszbach, Knut Are Aastveit, Lee Ohanian, Luís Teles Morais, Magnus Gulbrandsen, Marianna Kudlyak, Mauricio Ulate, Mariano Palleja, Pierre-Olivier Weill, Ragnar Juelsrud, Roine Vestman, and seminar participants at the 2024 AREUEA National Conference, CSU Long Beach, the Economics of Migration Junior Seminar, the Federal Reserve Bank of San Francisco (FRB SF), the Riksbank, Stockholm University, UCLA, and Uppsala University; for their hospitality during the writing of this paper, I thank the Department of Economics at Stockholm University, the FRB SF, and Norges Bank. All mistakes are my own.

counteracts the rise in the leaving rate induced by the income shock and varies by housing wealth. Whether the channel is quantitatively important to explain migration following regional income shocks is an empirical question. To assess it requires a scenario featuring 1) changing economic conditions, 2) ensuing home price changes, and 3) data on individual moving decisions, housing tenure, and other economic factors over time. Such settings are rare, but found in Norway.

My analysis examines the impact of the global collapse in oil prices during 2014–2016 on the Norwegian labor market area (LMA) encompassing Stavanger—a city referred to as the oil capital of Norway. Due to a supply glut driven mainly by technological innovation and increased shale oil production in the United States, global oil prices fell over 50% during an 18-month period following June 2014 (see Figure 2.1).[2] In response to the price shock, Norwegian oil producers implemented substantial cutbacks in investments and labor costs, which reduced labor earnings and raised unemployment in oil and the non-tradable sector in Stavanger (Juelsrud and Wold, 2019). Local home prices fell while the rest of the country saw rising prices. However, out-migration was little affected even though the technological origins of the shock caused a permanent change in the region's economic prospects.

This quasi-experimental setting combined with rich administrative data in a country with a high homeownership rate allows me to study housing wealth shocks' influence on migration. I document heterogeneous responses using a continuous difference-in-differences framework. The small, observed rise of 0.37% in the aggregate out-migration of Stavanger following the shock is driven by a 41% increase in the migration of renters and homeowners with little housing wealth and muted by a -26% reduction in the migration of other homeowners, indicating that the housing wealth channel is more important than the rent channel in explaining migration. To isolate the channel, quantify its welfare consequences, and analyze policy in this setting, I estimate a life-cycle model with endogenous location, housing, and saving decisions, location-specific returns to skills, and individual location preferences. I

[2] The decline in price expectations was significant: Brent oil futures maturing in January 2023 traded at \$102 in early January 2014. By the end of 2014, they had plummeted to \$53, and by early 2016, the price reached a trough at \$34.

show that the housing wealth channel is key to understanding the observed differences in migration; when home prices re-equilibrate in response to the shock, the rise in out-migration falls from 29% to 2.6% and the model replicates the behavior that homeowners with more housing wealth reduce their mobility while low-housing wealth owners behave like renters. The model illustrates that the general equilibrium effect amplifies the shock to homeowners in terms of welfare and that renters benefit more from untargeted moving subsidies. The analysis consists of three parts.

First, I document the economic impact of the oil price plunge on Stavanger's labor and housing markets compared to the rest of Norway. I show that, following the shock, the income growth of Stavanger workers both in and outside the oil industry lagged behind that of workers in the rest of the country and that unemployment was elevated throughout the studied period. The section discusses the persistence of the shock and perceptions of it among policymakers and workers. I also discuss the decline in net migration and show that it is driven by a reduction in the in-migration rate; the out-migration probability increased modestly by 0.37% during 2015–2018 and was even 4.5% (0.11 percentage points) below the pre-shock level during 2017–2018, while in-migration was 30% (0.030 percentage points) lower during 2015–2018. Projections for the long-term population size of Stavanger were revised significantly downward.

Second, I empirically document reduced-form facts on the heterogeneity in the change in migration that supports that the housing wealth channel influences migration. I find that people with no or little housing wealth left the area at a higher rate following the shock while those with more housing wealth tended to stay with a higher probability than before. The increase in the out-migration rate of the former group is 41% while the reduction in the latter is −26%. The data allow me to rule out other potential explanations for the divergence by contrasting the impacts among renters and homeowners along with other observables such as age, prior income, net worth, and attachment to the region through family ties. I also run horse-race regressions to simultaneously control for observables that correlate with selecting into a housing tenure. I find that the increase in the departure rate of renters is not explained by other covariates while a rise among the young can be attributed to a significant extent to other observable factors.

I also document changes in other dimensions of migration that are in line with a significant role of the housing wealth channel. While workers who leave Stavanger are more likely to move to locations with higher incomes and home prices following the shock, they are much less likely to become homeowners in the destination, both compared to before the shock and to arrivals from other locations in the same period, and when controlling for the changes in the mover composition. While 32% of movers bought a dwelling before 2014, only 22% do after leaving Stavanger following 2014. The composition of arrivals experiences a shift toward groups that benefit more from cheaper housing: Renters with less labor earnings continue to move to the area at the same rate as before, while homeowners across the board avoid it. Only the age group of 58–66-year-olds do not significantly reduce their arrival rate, and people with family ties in the region also reduce less.

The third part of the analysis is based on a spatial life-cycle model similar to Kennan and Walker (2011) and Giannone et al. (2023), which I use to estimate the welfare impact of the reduction in home prices and elucidate why homeowners respond differently to renters. The model also provides an environment to test the efficacy of moving subsidies, a policy to promote labor mobility. The model incorporates location, housing tenure, housing size, and saving decisions, and it considers location-specific returns to skills and individual location preferences to generate pre- and post-shock heterogeneity. Home prices are determined in equilibrium and thus change with local economic conditions.

The model includes moving costs, preference shocks, adjustment costs of housing, and borrowing constraints that depend on the home value, which yield not strictly concave utility functions. To solve such a model, I combine the nested-value function method and the endogenous-grid method with an upper envelope step, as presented in Druedahl (2021), with a framework of discrete location choices. To my knowledge, this is the first application combining these techniques when solving a McFadden (1974) style model. In addition, I apply a transformation of the savings grid, from expressed in nominal terms to share-of-housing-value, to sidestep the problem of many house-and-location-specific borrowing constraints in the form of a uniform cap on the loan-to-value ratio, a potentially novel innovation.

As I endogenize the cost of housing and worker's housing wealth to local economic conditions, the migration decision, in turn, is indirectly influenced by general equilibrium forces. Following a 6% reduction in overall Stavanger earnings and letting home prices re-equilibrate, renters leave at a 13% higher rate than before. Homeowners do too, but on average to a lesser degree because of heterogeneous wealth effects. Homeowners with low housing wealth respond like renters (14.7% rise in out-migration), while those with high housing wealth reduce, on net, their leaving rate by -8.7%. The average rise across groups is 2.6%. This is consistent with my empirical results. I then decompose the effect of the income shock and home price shock by feeding them into the household problem separately.

Both homeowners and renters who stay at least one more period in Stavanger suffer from the reduction in income. In terms of equivalent variation (EV), homeowners experience a welfare reduction of -1.3% compared to -2.5% for renters. The difference is due to renters, on average, earning less. However, the welfare impact due to the home price shock diverges. I estimate that homeowners who stay experience a welfare loss of -3.7% while renters enjoy an average rise of 2.6%. In the scenario of both an income shock and an immediate home price re-equilibration, the net effect on homeowners and renters who stay is -4.0% and -0.59%, respectively. I.e., due to the housing wealth effect, homeowners experience addition loss, while renters are compensated by a rent reduction.

The value of leaving following the joint shock, measured as EV, also differs across housing tenure. A renter who leaves immediately experiences a decrease in welfare that is negligible compared to the same renter before the shock.[3] The homeowner who sells their house and relocates endures a -3.7% loss of welfare, all due to the housing wealth effect.

Thus, homeowners are in terms of welfare worse off whether they leave or stay, and relative to them, renters are better off in either location. But for migration, it is the differential

[3]The value of leaving exhibits a minute reduction of -0.0029% because there is a non-zero probability that they will return to Stavanger and then experience the reduction in local income, which leads to a reduction in welfare compared to before the shock.

between the present values of leaving and staying that determines migration and due to the housing wealth channel, it has increased the most for renters, which, in part, is why they are more responsive in terms of migration.

Given the role of the housing wealth channel, how can policy alleviate it? An example that has been used historically is the provision of moving subsidies, i.e., to offer financial assistance conditional on a worker moving far enough to accept a new job. In an experiment using the model, workers are offered such subsidies if they leave Stavanger. If the subsidies are conditional only on recipients' leaving (i.e., untargeted), the leaving probability elasticity is four to five times higher for renters. This is because renters are, on average, more financially constrained. The welfare improvement of moving in an environment with moving subsidies is also greater for renters.

This implies that moving subsidies can have unintended spillovers onto homeowners. A policy that encourages renter migration can amplify the drop in housing demand and home values. Thus, independent of how the policy is financed, such subsidies can become welfare transfers from owners to renters. Taking mobility behaviors as fixed, this analysis suggests that location-based policies can be more suitable to address worker welfare. Stimulating business creation, growth in distressed regions, and attracting new workers would counteract home price declines and make use of existing housing.

Related literature. While previous work such as Munch et al. (2006), Battu et al. (2008), and Blanchflower and Oswald (2013) focus on the higher moving costs associated with homeownership and lower inter-regional mobility following idiosyncratic income shocks, this paper focuses on a channel arising from regional economic shocks and general equilibrium effects. I show that the housing wealth channel can influence migration and thus housing tenure matters in the response following economic shocks. It complements the work of Notowidigdo (2020) who shows that low-income workers in the U.S. benefit from declining rents after adverse shocks. Given high U.S. homeownership rates, the wealth effect for homeowners uncovered here is also a relevant aspect and speaks to the work of Autor et al. (2021) and complements the study of migration dynamics in Monras (2018) and Rodríguez-Clare et al. (2020).

This paper also relates to work on how being underwater on a mortgage (Modestino and Dennett, 2013; Valletta, 2013) or how facing credit constraints reduces mobility (Fonseca and Liu, 2023; Giannone et al., 2023). My analysis complements by highlighting a distinct mechanism that not only financial frictions related to housing pose as a hindrance to migration but also housing wealth. The empirical setting is conceptually different because the source of the variation in economic variables is persistent, while financial shocks only cause temporary fluctuations in home prices that do not generate a wealth shock in the theoretical sense.

Foundationally, this work builds on the seminal paper Kennan and Walker (2011) by studying the costs of migration, adding to their work by endogenizing housing wealth to local economic conditions and allowing moving costs to vary. This perspective also builds significantly on Glaeser and Gyourko (2005), who highlight the role of home prices in migration from declining regions. My paper complements such work by providing worker-level evidence on how regional shocks have heterogeneous migration and welfare impacts based on housing wealth. I also add to the location value framework from Bilal and Rossi-Hansberg (2021) by endogenizing location value. Both in the empirical analysis and the model, I demonstrate the importance of the general equilibrium effect to understand the weakened out-migration incentive of homeowners.

The mechanism discussed also has long-term effects on the local demographics. As young workers leave and older and poorer households with relatives in the area enter the depressed region, the labor pool and household demand shift, potentially influencing aggregate outcomes. The housing wealth channel interacts with the location-preference forces in Zabek (2024), which I complement through the study on changes in the in-flow to the depressed location. Endogenizing population responses to economic conditions also adds to the literature on the determination of local housing prices, such as Määttänen and Terviö (2014) and Landvoigt et al. (2015). While existing work often takes population changes as exogenous, to focus on the distribution of local home prices, this model allows for a two-way interaction between the population composition and market clearing. Thus, it provides insight into the joint problem of migration flows and home prices.

Finally, my work is by no means the first to use the impact of the 2014–2016 oil price plunge on Stavanger as an exogenous shock to economic conditions. The first study, to my knowledge, is Juelsrud and Wold (2019), which studies the effect of increased job-loss risk on household savings. Later examples are Fagereng et al. (2022), Lorentzen (2023), and Aastveit et al. (2024) (i.e., Chapter 3). However, this study differs from previous work by focusing on the determination of home prices, regional housing demand, and migration.

Roadmap. The paper proceeds as follows. Section 2.1 describes the data and the sample selection of the empirical analysis and presents summary statistics. It is followed by Section 2.2, where I describe the time period leading up to the 2014 oil price plunge and the economic consequences for Stavanger and its population. In Section 2.3, I introduce a toy model for how to think about the influence of economic factors on migration. Section 2.4 presents the empirical results of the paper. The life-cycle model is described in Section 2.5, where I also present the main results that it produces. Section 2.6 concludes by summarizing the paper and discussing potential future work.

2.1 Data

This section describes the data sources used and presents summary statistics on the sample used in the empirical analysis.

2.1.1 Sample selection and data sources

I combine several registries from Statistics Norway (SSB) to construct a panel of every person above age 24 living in Norway. I observe (anonymized) identifiers for each individual that allow me to match observations across datasets. I refer to these anonymized identifiers as IDs.

First, I select all individuals with a tax record with the Norwegian Tax Administration (*Skatteetaten*) in 2014. I observe all formal income streams such as salaries, business income, capital gains income, government benefits, and unemployment benefits (UB) at an annual frequency. The tax administration also records gross and net wealth and its components, such as the value of the primary residence, other real estate assets, deposits, different financial assets, and debt. From these data, I construct an annual panel for every individual.

Second, I match on data from the National Population Register (*Folkeregister*), which has a dataset containing the year and month of registered moves going back to 1966 and the origin and destination municipalities of the moves. Another dataset provides, for each ID, the associated IDs of parents, siblings, and children who are ever registered in Norway. This allows me to track where a worker's relatives are over time; I define a relative as either the parent or the sibling of the focal worker, ignoring children. I have household identifiers and observe the IDs of registered partners. Thus, using the information above, I can also track where individuals are located in relation to their partner's relatives.

I define homeownership status by checking whether any member of the household has wealth in the form of a primary residence; if so, I define all members of the household as homeowners, and if not, I define them as renters. I focus on primary residence because its location (which is not reported) coincides with the household members' location. Using this definition, I arrive at renter shares of 21.5% for Stavanger and 22.8% for the rest of Norway in 2010–2013, close to the 24% among all households in 2013 reported by Rustad Thorsen (2014). The individual renter rate is not publicly available from Statistics Norway before 2015, but it reports that, in 2015, 19.2% of all individuals in the age range 20–66 were renters. During the sample period, I observe transacted homes, their prices, and their locations.

The main geographical units that I study are labor market areas (LMAs), as defined in Bhuller (2009). There are three levels of regional administrative units in Norway. The lowest are at the municipality level (*kommune*), aggregating to counties (*fylke*), which in turn aggregate to the national level.[4] Bhuller (2009) constructs commuting zones or LMAs based on commuting patterns across municipality borders. This construction allows LMAs to overlap parts of multiple counties, in contrast to alternative methods of defining labor market areas. Under this categorization, Norway comprises 46 LMAs. In the worker panel, I label every move either an intra- or an inter-LMA move. I also note whether the move is to or from an LMA where a relative lives at the start of the year. For individuals who move more than once in one year, I record the first origin and the last destination and categorize the move based on these. Throughout the paper, "Stavanger" refers to the Stavanger LMA, which includes the municipality of Stavanger as well as Bjerkreim, Eigersund, Finnøy, Forsand, Gjesdal, Hjelmeland, Hå, Klepp, Kvitsøy, Lund, Randaberg, Rennesøy, Sandnes, Sokndal, Sola, Strand, and Time. The travel time between the most distant administrative centers in each municipality is approximately 2.5 hours. Stavanger municipality is at the region's center and is home to approximately a third of its population.

2.1.2 Summary statistics

The main analysis studies workers leaving or arriving in Stavanger. The region's population differs in important aspects from that of the rest of Norway, which I attempt to account for in the analysis. However, on many characteristics, Stavanger workers are similar to Norwegian workers overall, as shown in Table 2.1. For example, they are of approximately the same average age and live in households of similar size. Their homeownership rates and the likelihood of living in the same LMA as an immediate relative do not differ from the national average. Where workers in this LMA differ is across labor market observables. Stavanger has a high share of workers employed by the petroleum sector, where workers on average earn more than twice the non–petroleum sector income. However, even non-petroleum workers in Stavanger have above-average earnings, and the share of workers who receive unemployment benefits is lower (3.9% versus 5.4%). Petroleum workers also have higher educational attainment. The share of workers with a graduate degree or PhD is much

[4] More populated municipalities can also be divided into urban districts (*bydeler*).

higher in Stavanger's petroleum industry (26%) than in the rest of the country. Last, people living in Stavanger are less likely to do an inter-LMA move than workers in other LMAs, but they move more within the LMA.

Finally, we turn our attention to the population size. According to Statistics Norway, the age group 25–66 over 2010–2013 numbered between 2,687,785 and 2,785,563. After data cleaning, I arrived at a refined sample of 2,641,729–2,749,366 individuals per year. Notably, the total count of unique individuals in the analysis sample covering 2010–2018 amounts to 3,283,152. This figure is greater than the annual counts because of the inclusion in each year of new cohorts excluded from previous years. Note also that the numbers reported in the final row in Table 2.1 are different: the sums of individuals in the "All" columns do not align with the overall sample size. The difference arises from the exclusion of cohorts entering the analysis after 2013 and the double-counting of individuals who transition between oil and non-oil sectors or move between Stavanger and other LMAs. However, it is reassuring to note that the total count of unique individuals observed in Stavanger during this period falls within the range reported by Statistics Norway, specifically within the bounds of 181,485 to 193,394.

2.2 Economic impact of the 2014 oil price plunge

This section provides more details on the quasi-experiment and documents the impact of the fall in petroleum prices on workers living in Stavanger. I describe the evolution of the oil market leading up to the great correction in 2014–2016, provide several arguments why the shock should be interpreted as a permanent change in the economic conditions of the area, and that the reduction in labor earnings was significant for both oil and nonoil workers. I also discuss several aggregate time series such as forecasts of oil investments and the population size of Stavanger as well as the changes in net migration to the region following the shock.

Table 2.1: Summary statistics

	In Stavanger			In rest of Norway		
	Not oil worker	Oil worker	All	Not oil worker	Oil worker	All
Panel A: Demographics						
Age	44.1	43.1	44.0	45.2	43.5	45.1
Household size	2.88	3.01	2.90	2.72	2.91	2.72
Homeowner (%)	77.0	89.8	78.5	76.1	86.3	76.3
In rel.'s LMA (%)	78.6	75.1	78.2	77.1	82.2	77.2
< HS (%)	33.9	14.6	31.7	34.2	17.2	33.8
HS (%)	31.6	34.7	32.0	30.2	48.6	30.6
UG (%)	25.7	25.0	25.6	26.2	18.9	26.1
GD or PhD (%)	8.77	25.7	10.7	9.40	15.2	9.52
Panel B: Work, income, and wealth measures (NOK)						
Oil workers (%)	0.00	100	13.8	0.00	100	2.61
Skill s_i	1.29	2.33	1.41	1.43	2.37	1.45
Post-tax income	338,000	610,000	369,000	317,000	530,000	321,000
Salaries and wages	359,000	883,000	419,000	329,000	753,000	337,000
Receiving UB (%)	4.26	1.43	3.93	5.46	4.60	5.44
Panel C: Migration probabilities (%)						
Inter-LMA	2.53	2.04	2.48	3.01	3.20	3.02
Inter-muni.	2.70	2.78	2.71	1.88	1.57	1.87
Intra-LMA	9.39	9.42	9.40	8.69	7.73	8.67
Number ind.	156,179	30,907	187,086	2,228,712	93,204	2,321,916

Note: This table presents summary statistics of subpopulations of Stavanger and Norway in 2010–2013. Skill s_i is estimated; see Section 2.5.3. Abbreviations used: rel. is for relative, HS for high-school degree, UG for undergraduate degree, GD for graduate degree, and UB for unemployment benefits. The average exchange rate of NOK to USD was stable over the period and on average 5.84 NOK per USD (source: Norges Bank).

Figure 2.1: Brent oil spot prices and futures for June 2023

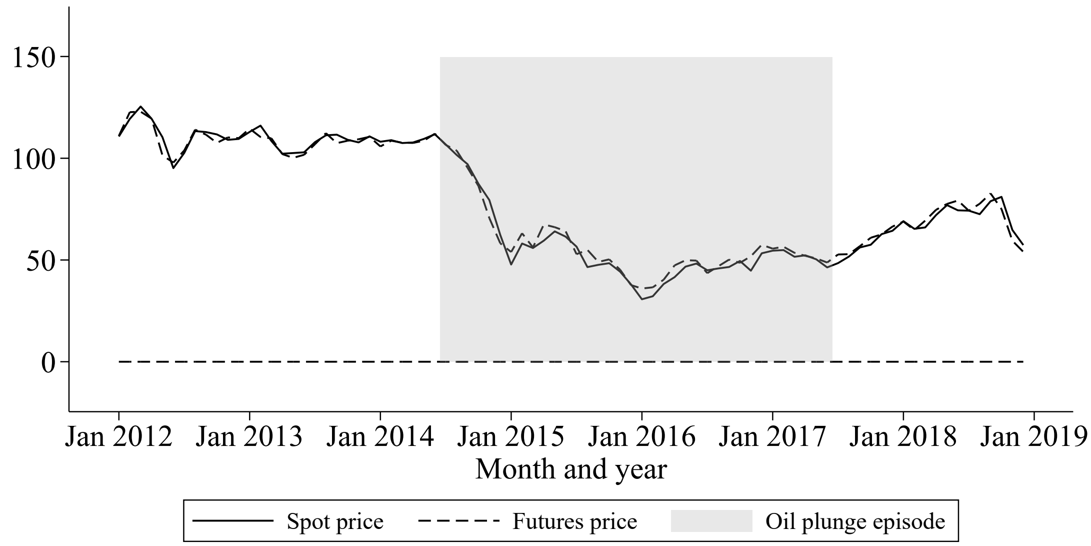

Note: This figure presents the time series of Brent oil prices. The solid line is the monthly average spot price and the dashed line is the monthly average futures price for June 2023 futures.

2.2.1 Supply glut and price plunge

During the first four years of the 2010s, petroleum prices were strong, and the crude oil spot price averaged approximately \$103 (IMF, 2023). This motivated further investments and innovation in unconventional petroleum extraction methods in the U.S. and an increase in production from Canadian oil sands. U.S. shale oil (or, tight oil) had risen since the mid-2000s from 0.35 million barrels per day in 2005 to 0.61 million barrels per day in January 2010 and 3.3 million barrels per day by January 2014, peaking at 8.2 million barrels per day in November 2019 (EIA, 2023). By 2014, the U.S. had become the world's largest producer of crude oil (OECD, 2016). The growth in Canadian production was more modest, with production rising from 2.6 million barrels per day in January 2010 to 3.7 million barrels per day in January 2014 and peaking at 5.0 million barrels per day in December 2019 (CER, 2023).

However, starting in July 2014, the price of oil started to fall; see Figure 2.1. Over the following 18 months, prices fell by over 50%. The literature has not identified a trigger event that caused the market to suddenly reprice oil, but, e.g., Baumeister and Kilian (2016) and

Stocker et al. (2018) both agree that the decline was due mostly to the increased global supply discussed above. The former argue that approximately half of the decline after June 2014 was predictable in June with a VAR forecasting model and that demand shocks were present but of less importance. In addition, a large share of the surprise component of the decline is explained by the more pessimistic outlook for future oil prices. Another contributing factor on the supply side was OPEC's abandonment of price controls in late 2014, which alleviated supply disruptions in the Middle East. On the other hand, both Baumeister and Kilian (2016) and Stocker et al. (2018) document that weaker demand also contributed to the 2014 price shock.

2.2.2 Impact on the Stavanger economy

This section documents the economic conditions in Stavanger following the shock. I contrast Stavanger workers in oil and nonoil sectors to workers in Norway and show for both groups that the reductions in labor earnings were significant. Unemployment rose as well and home prices dropped significantly relative to the rest of the country.

The brunt of the economic impact of the price fall was carried by oil firms, their suppliers, and, consequently, workers in these industries. Investment in new offshore drilling rigs and infrastructure to exploit untapped oil fields fell, and labor expenses were cut, leading to a rise in unemployment and a fall in the earnings of oil workers. Approximately 14% of Stavanger's working-age population was employed directly by the petroleum industry during 2010–2013 and was thus severely impacted by the shock. This is illustrated by the earnings of workers in Stavanger companies in the oil and gas sector in Panel A, Figure 2.2. For every year of a worker's employment history without uptake of unemployment benefits, I sum the earnings from labor and self-employment, call this labor earnings,[5] and estimate

$$\log(\mathbb{E}[LE_{it}|X_{it}]) = \alpha_{bt} + \eta\, X_{it} + \beta^S_{bt} \times \mathbb{1}(\ell_{it} = S), \tag{2.1}$$

[5]The labor earnings measure is winsorized at the 1st and 99th percentiles. This effectively makes the variable nonnegative and handles a couple of very large outliers.

Figure 2.2: Impact on labor outcomes of workers in Stavanger

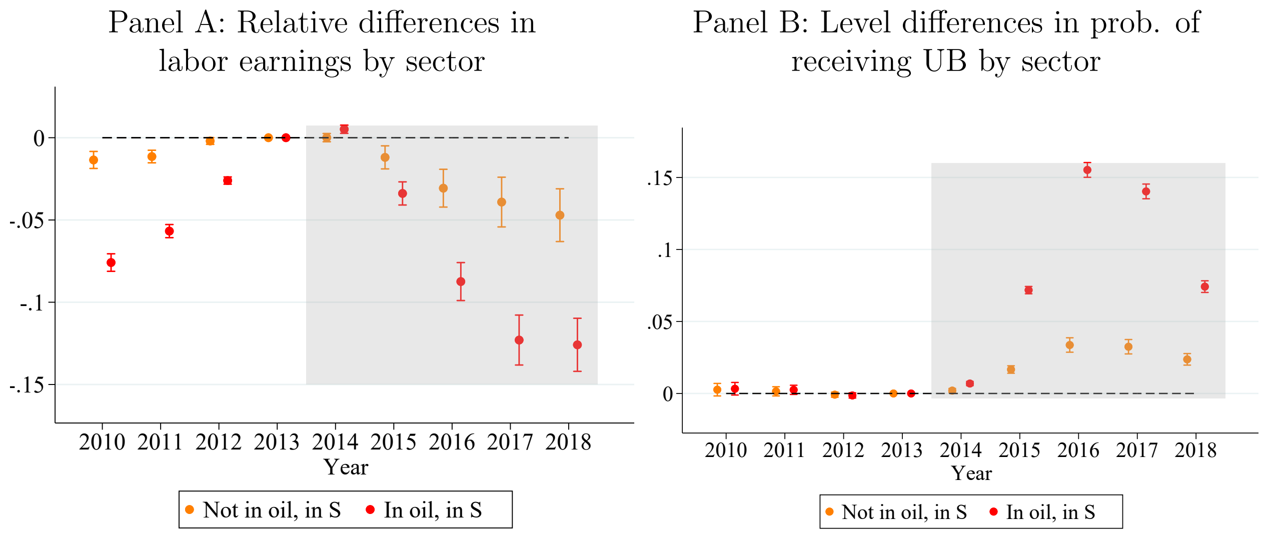

Note: This figure presents the annual differences in labor market outcomes of workers in Stavanger in either oil or nonoil compared to workers in the rest of Norway. Panel A displays the log differences in labor earnings estimated using (2.1) and Panel B the level difference in the probability of unemployment benefits (UB) uptake estimated using (2.2). Labor earnings are the sum of wages, salaries, and income from self-employment. Workers are categorized as receiving UB if they receive any unemployment benefits during the year. The sample is conditioned on no one moving in the post period to avoid reflecting a change in the worker composition of Stavanger. Workers are excluded from Panel A in years of UB uptake.

where β_{bt}^{S} is the per-year earnings difference between either oil or nonoil workers (indicated by b) in and outside Stavanger. To make the pretrends clearer, I set $\beta_{b2013}^{S} = 0$. The model includes worker and year fixed effects and is numerically estimated by the Stata command in Correia et al. (2019). We see that oil workers in Stavanger experienced significant above-trend income growth before 2014, as did the LMA's nonoil workers, albeit not as large. The trend in income growth is reversed in the post-period, where oil workers lose the most relative to the overall trend of the country. The worsened labor market conditions are due not only to the reduction in demand directly from the petroleum sector but also the reduction in workers' overall demand (Juelsrud and Wold, 2019). In my analysis, the sample excludes workers who leave the current labor market in the post-period, so the changes in income growth do not reflect changes in the composition of workers.

Similarly, I estimate the changes in the probability of receiving unemployment benefits (UB), using the model

$$\mathbb{1}(\text{UB}_{it} > 0) = \alpha_{bt} + \eta\, X_{it} + \beta^{S}_{bt} \times \mathbb{1}(\ell_{it} = S) + \varepsilon_{it}, \tag{2.2}$$

which is estimated with ordinary least squares (OLS). This illustrates further that oil workers experienced the most drastic changes in their labor market conditions, with a large jump in the uptake of unemployment benefits; see Panel B. In 2010–2013, unemployment among this group was generally below average (see Table 2.1).

In Figure 2.A.3, I present additional results on the impact on labor earnings and uptake of unemployment benefits. I show that oil workers in other Norwegian regions also experienced elevated unemployment and a similar reduction in labor earnings. Thus, the outside option for Stavanger oil workers has worsened too, and so I drop them from the main analysis. The figure also displays some heterogeneity in the earnings reduction by 2013 labor earnings. Nonoil workers with the lowest earnings in 2013 experienced a greater reduction; however, the rise in unemployment is lower.

The appendix also contains graphs of aggregate time series for the county of Rogaland (see Figure 2.A.4), 77% of whose population resides in the Stavanger LMA. The GDP of Rogaland declined for two years, while the growth in its disposable income, consumption, and employment lagged behind that of all other LMAs in Norway even after the oil price recovery in 2018 and onward.

Home prices in the region boomed during the early 2010s and plateaued during the period prior to the oil price plunge (see Figure 2.3). Relative to the national growth rate of home prices, the loss of housing wealth after the shock among Stavanger homeowners was significant, indicating that the return to living in the region was significantly reduced relative to that of living in other places (Bilal and Rossi-Hansberg, 2021). The population-weighted average of home price indexes in Stavanger fell by 8.5%, while it rose by 23% for the rest of the country's housing markets between the fourth quarter of 2013 and the fourth quarter of 2016.

Figure 2.3: Housing costs

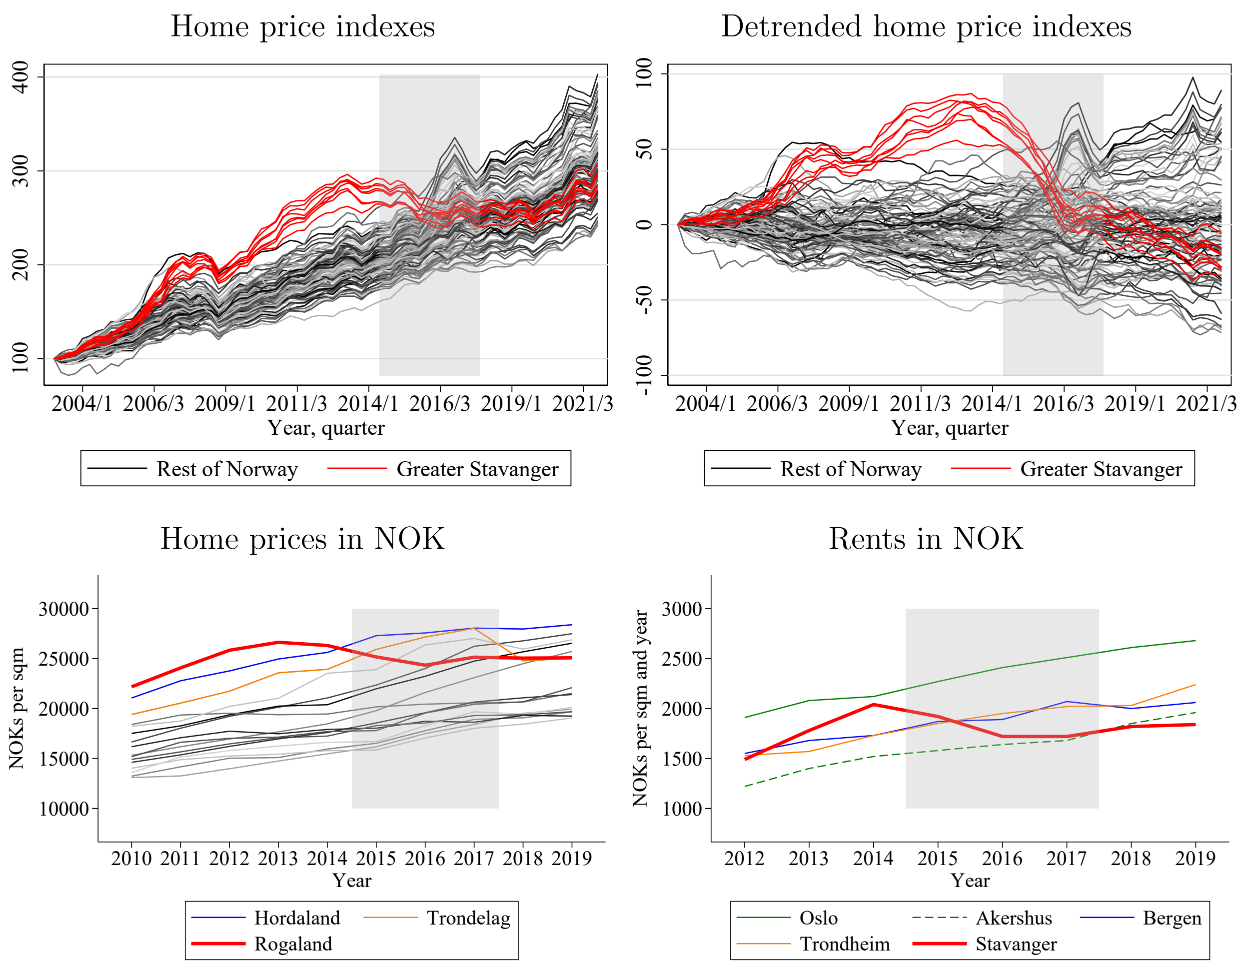

Note: This figure presents time series of the cost and value of housing in Norway. The top-left graph shows home price indexes for different housing markets, from Real Estate Norway (*Eiendom Norge*). I plot the indexes for housing markets in the Stavanger LMA in red and all other Norwegian housing markets in different shades of gray. The oil plunge episode is marked out in gray in all graphs. In the top-right graph, I plot a detrended version of the top-left graph, taking out the annual home price growth rate in Norway. The bottom graphs are in Norwegian krone (NOK) and are constructed using data from Statistics Norway. The bottom-left graph shows the square meter prices across counties, excluding Oslo and Akershus (Figure 2.A.2 includes them). The bottom-right graph shows the square meter rents across select cities and regions.

Figure 2.3 also displays the nominal level of home prices and rents using data from Statistics Norway. The home price time series are available at the county level. In the graph, I exclude the series of Oslo and Akershus because homes are much costlier there and the changes in prices in cheaper areas are dwarfed when plotted using the same vertical axis. If we approximate Stavanger prices by Rogaland County, we see that home prices were above all other counties (excluding Oslo and Akershus) before the plunge episode, and fell below both Hordaland County and Trøndelag County following 2014.[6] I highlight the differences between these counties because Bergen and Trondheim, two other major cities in Norway that are likely destinations for Stavanger workers, lie in them in respective order. Data on regional rents are only available for select cities which are presented in the bottom-right graph in Figure 2.3. They display similar trends: Stavanger rents fall during the plunge episode and do not recover in 2018–2019, while rents continue to rise in other locations.[7]

From these graphs, we conclude that the cost-of-housing differential changed to an economically significant extent over the episode. I quantify it as follows using the data from Statistics Norway. The square meter price in Stavanger (Rogaland) fell by 1970 NOK between 2014 and 2016, or, 7.5%. Meanwhile, prices rose by 4330 NOK in the rest of Norway (weighted by county population). Relative to the 2014 square meter price in Stavanger, that is a rise of 16%. Rents are measured by surveys in October every year, and fell by 320 NOK per square meter and month, or, 16%, while rents in the rest of Norway rose by 198 NOK, or, 9.7% relative to 2014 Stavanger rents. The asymmetry in the rise outside Stavanger and fall inside across housing tenure can be because of the difference in the geographical unit used, but the total change is similar. The cost of a home outside Stavanger rose by approximately 24% and rents by 25%.[8]

[6] The 2018 drop in Trøndelag coincides with the merge of Nord-Trøndelag and Sør-Trøndelag into one county.

[7] The only complete time series are for 3-room apartments, which is the rent I use.

[8] This is essentially adding up the relative changes in Stavanger versus outside.

2.2.3 Post-2014 outlook for oil and Stavanger

For a worker to find it worthwhile to foot the cost of moving, the changes in income differentials need to be persistent; the relatively higher present value of higher future income in another location has to outweigh the short-term costs of moving. While it is difficult to gauge the expectations held by workers during the episode, multiple sources indicate that the shock should be viewed as having fundamentally transformed the prospects of Stavanger for the foreseeable future.

Stocker et al. (2018) and Bjørnland et al. (2021) argue that continuing technological advancements in U.S. shale oil production and the flexibility of the technology make a persistent recovery to pre-2014 price levels unlikely. The new unconventional extraction methods have much shorter lead times, with the time lag from investment decision to first extraction being weeks instead of several years, which makes supply overall more responsive to changes in prices. This feature can also explain why the rise in U.S. production was constantly underestimated for a long time, leading to the plunge itself. Stocker et al. report a 2025 nominal forecast of $65 per barrel, which aligns with the June 2023 Brent oil futures presented in Figure 2.1 and other projections such as those of Norges Bank (2014). This outlook does not rule out the possibility of temporary fluctuations because of geopolitical events or positive demand shocks but indicates that there is a significant downward shift in the expectation of future oil prices.

The potential spread of shale oil technology should make the market even more elastic and lower future profitability (Clerici and Alimonti, 2015), exacerbating the situation from the perspective of Norway, which is not suitable for horizontal extraction methods since oil has been found only offshore.

Forecasts for the Norwegian economy in the wake of the shock were mixed. Projections in a 2015 report from Statistics Norway showed falling investment in the petroleum sector with a projected small recovery that would not return to the 2015 level, which in turn was far below the peak in 2013–2014 (SSB, 2015). Stavanger was not explicitly mentioned in the report but stood to be more severely impacted by this trend. In addition, Norway's fiscal spending would suffer in the long run: in 2013, about 21% of the government's revenue was directly raised from the special oil tax and dividends from government-owned shares in oil companies. In mid-December 2014, the central bank initiated a series of policy rate reductions. Prior to this, the policy rate had been at 1.5%, which was later lowered to 1.25% and eventually reached its lowest point at 0.5% on March 17, 2016. The policy rate remained at this level until September 2018 when it was raised once again. This added further support to other Norwegian exports through a depreciation of the local currency, the krone.

The declines in population projections and home prices in Stavanger indicate that workers likely perceive the region's future prospects to have worsened following the shock. Both news articles such as Hetland and Oppedal (2016) and journal articles such as Jabobsen and Kvittingen (2016) present a grim outlook for the possibilities of finding well-paid employment in Stavanger. They highlight that the decade-long rise in oil extraction costs, attributed to the prolonged period of elevated prices and the limited incentive to reduce expenses, paved the way for potential cost reductions and long-term sustainability—developments that, while potentially positive in aggregate, do not necessarily work to the benefit of Stavanger workers in and outside of oil.

The drop in employment (measured for 15–74 year-olds) was also significant. Having remained at a stable average rate of approximately 74% during 2010–2014, the employment rate fell to an average of 69% during 2015–2019, a drop of -3.1 percentage points relative to the rate in the rest of Norway.[9]

Lastly, the migration to Stavanger fell drastically while the out-migration experienced only a temporary increase. This is depicted in Figure 2.4 as both absolute and relative changes to the out- and in-flow in 2013. In the three first years, out-migration is slightly elevated and peaks in 2015 before falling to a slight decrease in 2018. Averaged across the years excluding 2014, the out-migration only rose by 0.37% (0.9 p.p.). In-migration, on the other hand, responds forcefully and bottoms out in 2015, approximately 50% below the 2013 rate,

[9] I use municipal-level data from Statistics Norway, which compiles employer-reported data on who is employed in the fourth quarter every year (SSB table 06445). To compute the LMA average, I weight observations by the municipality's annual population (SSB table 01222).

Figure 2.4: Changes in in- and out-migration trends for Stavanger

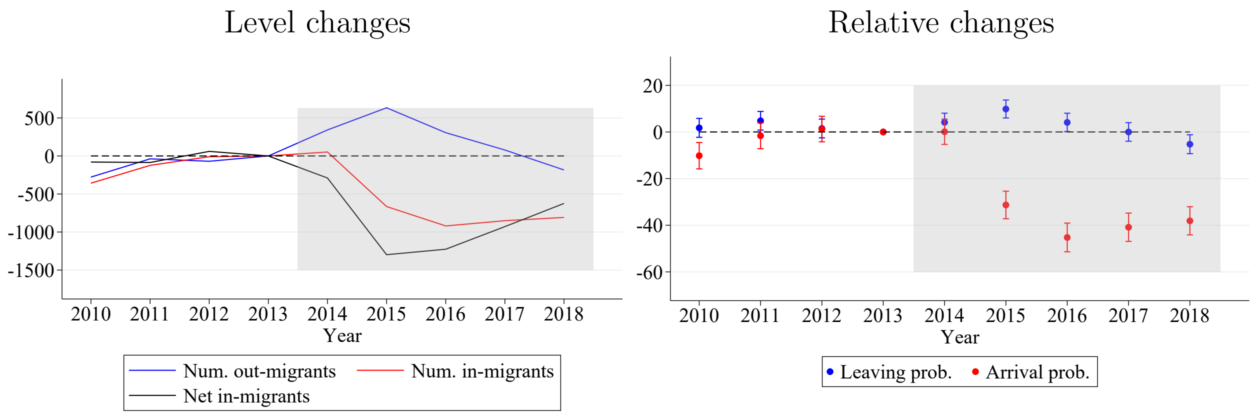

Note: This figure presents the changes in the migration to and from Stavanger relative 2013. The left graph displays the changes in the number of migrations and the net in-migration rate. The right graph displays the relative change. The sample is the same as that in the main empirical analysis. In Figure 2.A.1, I present the per-year migration levels.

and remains depressed throughout the time window considered. The average reduction was 30% (0.030 p.p.) The combined changes led to a great reduction in the net in-migration to the region.

2.3 A simple theoretical framework of migration

To provide an intuition of how shocks to incomes and housing prices affect migration and to guide the empirical analysis, I present here a one-period model inspired by Kennan and Walker (2011). The model is also useful to motivate the use of the logit model in the empirical analysis.

Let $\ell \in \{S, Q\}$ denote the location. A worker chooses where to live. She has a preference γ_S for S, which can either represent an idiosyncratic preference for S or a difference in amenities relative to Q. Location S provides Δw more income than Q and utility is linear. Moving incurs disutility τ and before choosing a location, she is hit by preference shocks ξ_ℓ. Consider the case when the worker starts in location S. Mathematically, the worker solves

$$V = \max \ \{\Delta w + \gamma_S + \xi_S, \ \ -\tau + \xi_Q\}. \tag{2.3}$$

The random variables ξ_ℓ are Gumbel distributed which implies the closed-form expression for the probability of leaving S before the shock is realized: (I provide a discussion on the properties of the preference shocks in Section 2.5)

$$\mathbb{P}(\text{leave } S|\Delta w, \gamma_S, \tau) = \frac{1}{1 + \exp(\Delta w + \gamma_S + \tau)}. \tag{2.4}$$

The expression has several intuitive properties. The greater the income differential Δw between S and Q, the more probable it is that the worker will stay in S. If the worker faces a higher moving cost τ, the lower her leaving probability, and her having a preference for the location has the same effect.

I now introduce housing in the model and consider two cases in parallel: One of a renter, and one of a homeowner. The renter pays rent r_ℓ in ℓ, while the homeowner initially owns a house worth p_S and has to first sell it and then pay p_Q if they move. Depending on the housing price differential $p_S - p_Q$, she either pays or earns from the transaction. The values by housing tenure is

$$V_{Re} = \max \ \{\Delta w + \gamma_S - r_S + \xi_S, \ -\tau - r_Q + \xi_Q\}, \tag{2.5}$$

$$V_{HO} = \max \ \{\Delta w + \gamma_S + \xi_S, \qquad -\tau + p_S - p_Q + \xi_Q\}, \tag{2.6}$$

for the renter and homeowner, respectively. The ex-ante moving probabilities are

$$\mathbb{P}_{Re}(\text{leave } S|\Delta w, \gamma_S, \tau) = 1/(1 + \exp(\Delta w + r_Q - r_S + \gamma_S + \tau)), \tag{2.7}$$

$$\mathbb{P}_{HO}(\text{leave } S|\Delta w, \gamma_S, \tau) = 1/(1 + \exp(\Delta w + p_Q - p_S + \gamma_S + \tau)). \tag{2.8}$$

In this setting, a shock to income influences the probability of leaving similarly, but who responds more is not obvious, especially if the moving cost τ depends on housing tenure (which is frequently argued in the literature). If home prices and rents also respond due to shocks to housing demand, the effect is again ambiguous.

However, if we measure the change in terms of the log odds for leaving[10] and denote the shock to the income differential by ζ_y, the rent differential by ζ_r, and the house price differential by ζ_p, the expression of the change in the log odds is

$$\Delta \log odds_{Re} = -\zeta_y + \zeta_r, \tag{2.9}$$

$$\Delta \log odds_{HO} = -\zeta_y + \zeta_p, \tag{2.10}$$

by housing tenure. Due to the functional form of the logit, the constant terms drop out when we take the difference in the log odds before and after the shock. Thus, for a reduction in the income differential ($\zeta_y < 0$), the log odds change by as much for homeowners and renters, which is an increase in the leaving odds. Reduction in rents and home prices, however, ($\zeta_r, \zeta_p < 0$), reduce, to some degree, the increase in the leaving rate.

Thus, the logit model allows us to abstract from differences in moving costs, as well as preferences for the location

Welfare: The model also allows for a simple welfare analysis which provides a different perspective on the consequences of housing price shocks versus rent shocks. Due again to the Gumbel distribution assumption, the expected present value for renters and homeowners is

$$\mathbb{E}[V_{Re}] = \log\left(e^{\Delta w + \gamma_S - r_S} + e^{-\tau - r_Q}\right) + \gamma, \text{ and} \tag{2.11}$$

$$\mathbb{E}[V_{HO}] = \log\left(e^{\Delta w + \gamma_S} + e^{-\tau + p_S - p_Q}\right) + \gamma, \tag{2.12}$$

respectively. The γ denotes the Euler–Mascheroni constant (approximately 0.5772). In the first equation, a reduction to rent r_S leads to an increase in the present value of staying in S and raises the renter's welfare (the left term on the right-hand side of equation (2.11)). A similar reduction in the home price p_S which appears in the term of the present value of

[10] The log odds for leaving is defined as $\log odds \equiv \log \frac{\mathbb{P}(\text{leave } S)}{1-\mathbb{P}(\text{leave } S)}$.

moving to Q, however, reduces the value of leaving S, and thus reduces the welfare of the homeowner.

In the case of the homeowner, their wealth is tied to the economic prospects of S and a shock to it reduces their welfare by making it less attractive to move. The renter, in contrast, benefits from the reduction in the local rent, as argued in Notowidigdo (2020). This pin-points a tension in models that ignore housing tenure: Renters benefit from lower rents and it motivates them to stay. This acts as an increase in the moving cost, a disutility, and hints that preference shocks are in part reflecting changes in rents. However, the nature of the rent reduction we are considering in this setting is permanent, while preference shocks change from period to period.

In reality, utility is not linear and renters and homeowners differ along other dimensions that influence migration. I address several important aspects of such heterogeneity in the life-cycle model in Section 2.5. However, for the following empirical analysis, this model is rich enough to help us think about the potential role of housing tenure and home price shocks following a regional income shock, and it illustrates that changes in migration can exhibit great heterogeneity.

2.4 Reduced-form results

This section presents the main empirical results of the paper. I present several figures illustrating how migration changes along different worker observables between 2011–2013 and 2015–2018, but renters are consistently leaving at a higher rate. Homeowners with little housing wealth also leave at a higher rate in the post period while homeowners with more reduce their leaving rate. I also present regression results where I account for multiple characteristics that are different across renters and homeowners that could explain why renters move more following the shock. Across potential confounding factors, the effect on the change in renters' leaving rate remains the same.

I also present results on how the choice of destination changes. Those who leave Stavanger move to higher-income, more expensive locations, and are much less likely to become homeowners at their destination, also compared to other arrivers. The composition of people moving to Stavanger changes too. I show that young people and high-income people avoid the region, while poorer, renters, older, and people with family ties in the area either move in at the same or at a higher rate than in the pre-period.

Since the income prospects of oil workers fell uniformly across the country after the shock and thus their outside option did not change, I drop this group from the analysis. However, the results are in general robust to the inclusion of oil workers. I also exclude the year 2014 from the analysis because the price plunge began about halfway through the year.

2.4.1 Changes in the characteristics of Stavanger leavers

As the model in Section 2.3 predicts, the migration response exhibits rich heterogeneity. This section presents the changes in out-migration rates and how they vary across different worker characteristics. I present the changes in terms of levels and the log odds of leaving versus staying. The latter facilitates comparison across groups when the preshock leaving probability varies greatly across them. This is illustrated in Panels A and B of Figure 2.6. In Panel A, the young exhibit overall higher mobility than older workers, and the level change for the young after the shock is large, dwarfing the change for older workers. Meanwhile, in Panel B, the change is expressed in terms of the log odds, which makes the changes in the other groups clearer.

To estimate the change in the probability of moving by different worker characteristics, I employ two model specifications: one is a linear probability model, and the second is a logit model:

$$y_{it} = \alpha + \sum_{b \in \mathcal{B}} \beta^b \, \text{post}_t \times \mathbb{1}(b_{it} = b) + \gamma^b \times \mathbb{1}(b_{it} = b) + \varepsilon_{it}, \tag{2.13}$$

where the outcome is a function of leaving the Stavanger LMA (i.e., inter-LMA migration):

$$y_{it} = \mathbb{1}(\text{leaving } S_{it}), \qquad \text{in the linear probability case, and} \tag{2.14}$$

$$y_{it} = \log \frac{\mathbb{P}(\text{leaving } S_{it})}{1 - \mathbb{P}(\text{leaving } S_{it})}, \qquad \text{in the logit case.} \tag{2.15}$$

On the right-hand side of (2.13), post_t is one following 2014 and zero otherwise and $b \in \mathcal{B}$ indicates the bin to which the worker is assigned. In the linear probability case, β^b is the change in the worker's probability of leaving the Stavanger LMA in terms of percentage points, and in the case of the logit, β^b is the change in the log odds. The changes are within the group across time. All standard errors presented are corrected for clustering within the individual.

The overall finding is that people who rent or have little housing wealth increased their mobility following the shock to Stavanger by 40% while people with greater wealth reduced it (-26%). The changes are illustrated in Figure 2.5, where I show the pre- and post-leaving probabilities of renters to the left and then those of homeowners by housing wealth in five increasing housing-wealth bins (Panel A). The bins are assigned by quintiles computed year by year. Renters are always more likely to leave, and they increase their leaving rate considerably, from approximately 5% to 6% following 2014. The homeowners with the least housing wealth were also more likely to leave before and increased their leaving rate about as much as renters in terms of the log odds (see Panel B). These two groups together are referred to as having "little housing wealth." The change is insignificant for the second housing-wealth bin, and for higher bins, there is a significant reduction in the leaving probability. I refer to the three higher bins as having "greater housing wealth."

The model in Section 2.3 clearly predicts that, if the current utility of living in a location decreases, the leaving probability increases. This is the case for renters and workers with lower housing wealth. However, the fact that homeowners with higher housing wealth show a decreased leaving rate indicates either that the value of living in Stavanger increases or that the disutility of moving, τ, increases by more than the change in the location value.

Figure 2.5: Leaving probabilities by housing tenure and housing wealth

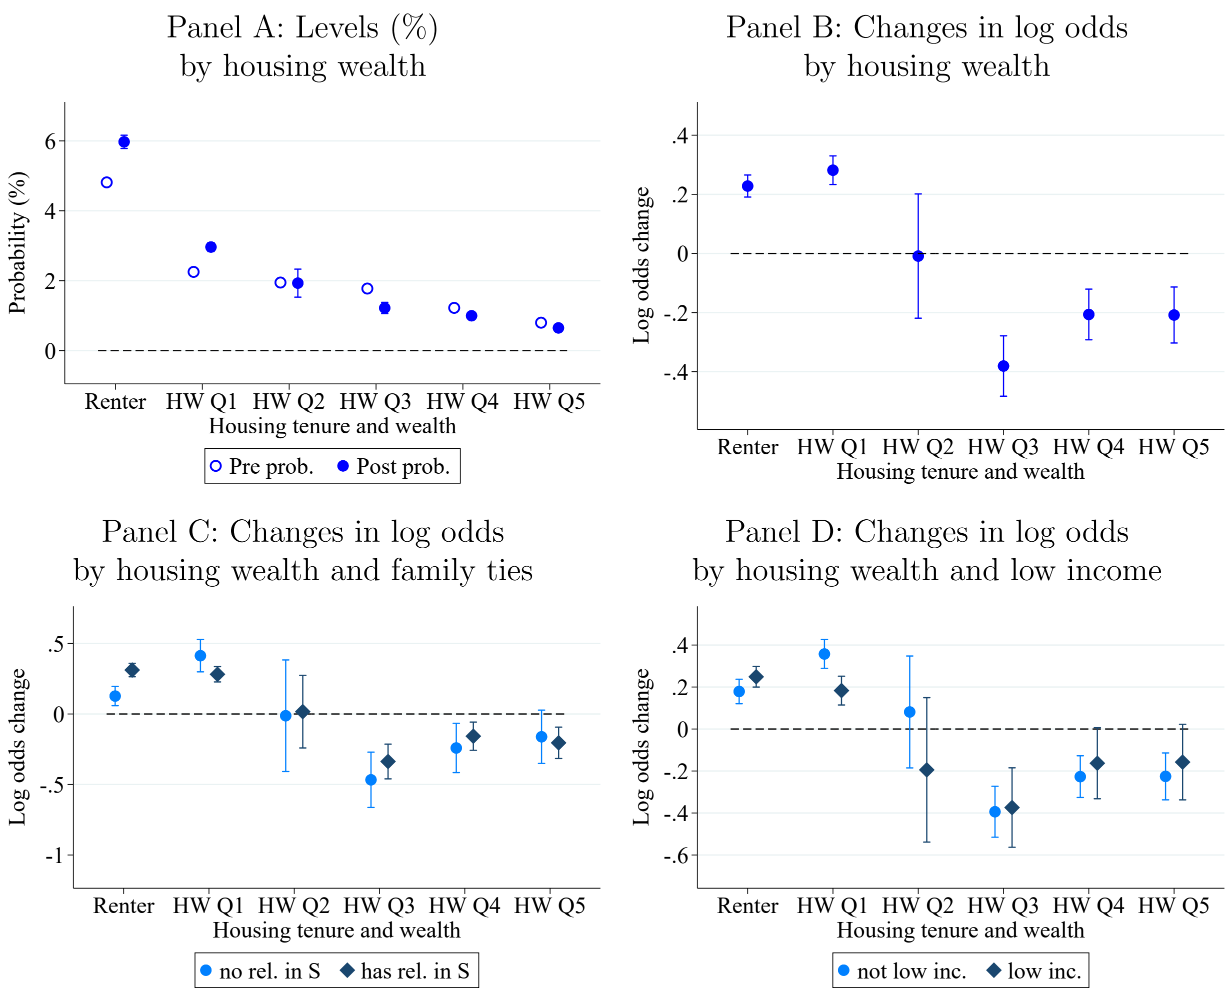

Note: This figure presents changes in the probability or log odds of out-migration from Stavanger following the oil price plunge of 2014. Panel A is produced by the OLS version of (2.13), Panels B–D are produced with logit, and Panels C and D use a third interaction term indicated by the corresponding legend. All error bands are 95% cluster-robust standard errors.

Housing tenure correlates with multiple factors, and given the costs of buying a home, being a homeowner signals a preference for the location or a plan to stay there for a longer time than a renter would. One reason for such a preference is having family in the area, which could make homeowners reluctant to leave. However, in Panel C of Figure 2.9, I split homeowners and renters by the presence of family in the region and find little difference in the change in the moving responses. The comovement across housing wealth bins is striking, and only for renters is there a divergence, where renters with ties are those who showed an increased leaving probability.

As we observed in Section 2.2, the income shock was heterogeneous across income groups, which can lead to different incentives to leave. Income also correlates with housing tenure and could thus explain why people with less housing wealth are more inclined to leave. However, I rule out this conjecture by showing in Panel D that people in the lowest income bin do not change their moving behavior more than others. In Figure 2.A.5, I also show that the differential effect does not have to do with low-income people becoming unemployed, which actually lowers the leaving probability. I show in the appendix that the higher rate of LMA moves among the unemployed is mostly explained by a general tendency of people in this group to move more frequently. When I include individual fixed effects or workers' average moving probability in years when they receive no unemployment benefits, the increase in moving probability explained by being unemployed is greatly reduced. In addition, we saw in Section 2.2 that the incidence of unemployment was lower for the low-income.

At this point, we have ruled out several potential mechanisms that could support the conjecture that workers richer in housing were not impacted by the shock, so the remaining conclusion is that their moving costs must have risen. This could occur through a reduction in the housing wealth of homeowners. When home prices in Stavanger fell relative to those in the rest of the country, homeowners who wished to move would make less from selling their current dwelling, leading to the household's being unable to afford housing of the same quality in other locations that it could have afforded before the shock. This acts as a pull factor on the homeowner to stay; in the framework of Section 2.3, her moving cost rises. The renter, meanwhile, experiences only the greater benefit of leaving through the increased income differential, a push factor. The same holds for homeowners with the least housing wealth, who lost relatively less wealth. However, some potential confounding factors remain to be ruled out.

Age correlates with homeownership, and younger people are generally more mobile. This is clear from Figure 2.6, Panel A, where I bin workers by age quintile. The group that

Figure 2.6: Leaving probabilities by age

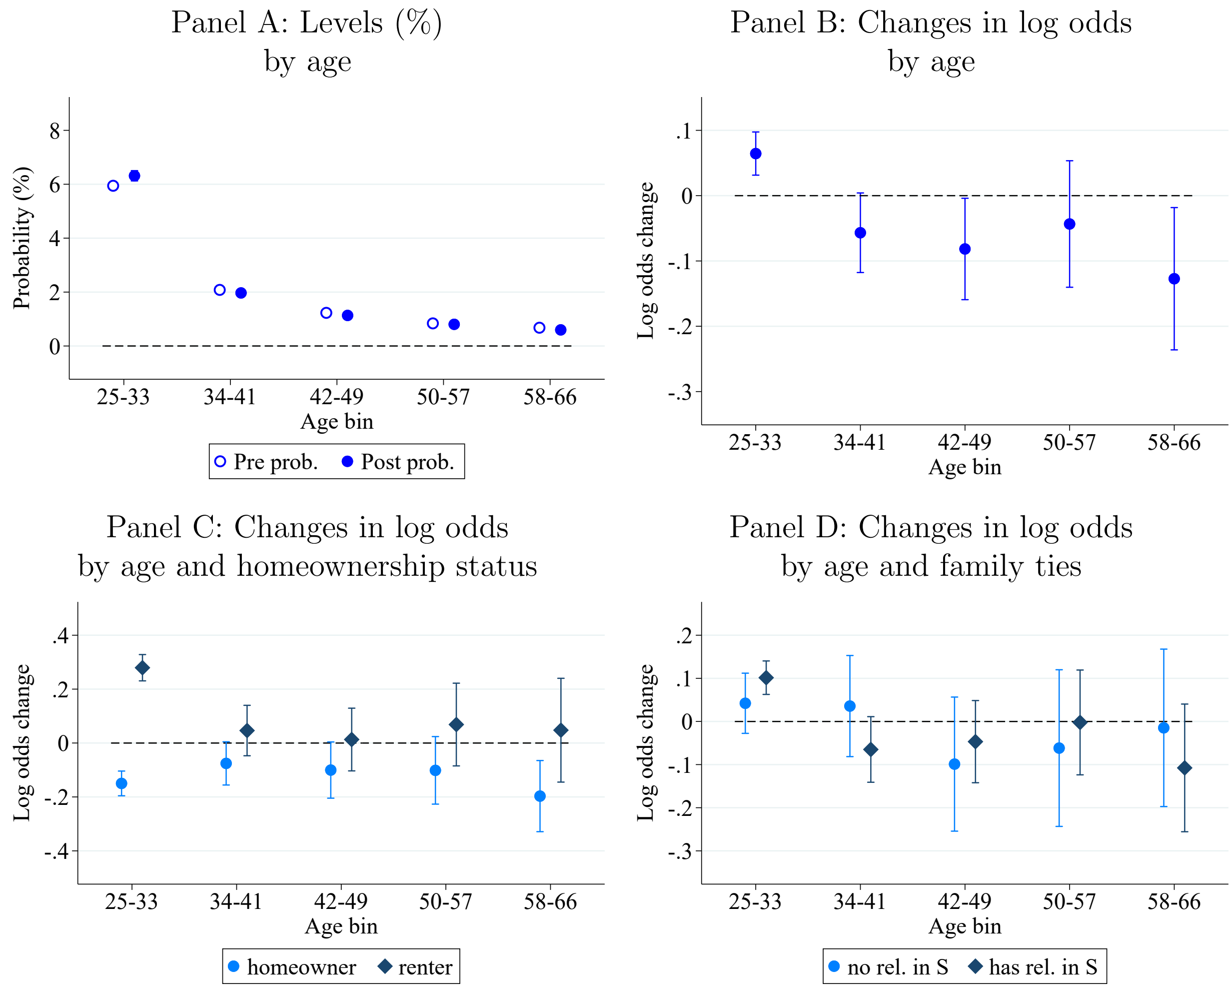

Note: This figure presents changes in the probability or log odds in the out-migration from Stavanger following the oil price plunge of 2014. Panel A is produced by the OLS version of (2.13), Panels B–D are produced with logit, and Panels C and D use a third interaction term indicated by the corresponding legend. All error bands are 95% cluster-robust standard errors.

increases its moving-out rate to a significant extent is the youngest bin of 25–33-year-olds. For older groups, the response is, if anything, in the opposite direction. This is more evident in Panel B, where I present the change in the log odds. If, for some unknown reason, older workers show a reduced leaving probability not because of a loss in housing wealth but for some other reason, this reason should hold for both older renters and older homeowners. This is, however, not the case, as shown in Panel C. Renters across all age groups leave at a higher rate, and homeowners' leaving probabilities are reduced, so the reduction among the homeowners seems not to be because they are on average older but because they own housing.

There is also a weak correlation between age and having immediate relatives in the region. This varies across LMAs, but in Stavanger, people with relatives in the area are slightly younger on average. This also does not explain why younger people leave more, as shown in Panel D.

From a life-cycle perspective, younger people naturally benefit more from leaving a location where expected incomes have fallen. The present value of a move is greater because their lifetime wealth comes from future income. By moving to a better location, they have more years to benefit from the move, while older people, who are closer to retirement, will be less willing to assume the disutility and monetary cost of moving. With this in mind, it is surprising that young homeowners show a reduced leaving probability, which means that the loss in housing wealth, or the increase in the moving cost, outweighs the income differential.

I also conduct a series of horse-race regressions, simultaneously including multiple factors to assess whether different combinations of them can better explain the different behaviors of renters versus homeowners and of the young versus the old. The model specification is the logit version of (2.13), and the results are presented in Table 2.2.

First, I contrast the youngest age group with all the others, controlling for the additional characteristics listed in the caption of Table 2.2, which are not interacted with the post dummy (Column 1). The reduction in the leaving rate is approximately zero for the youngest group. This differs from the result in Figure 2.6, Panel B, because of the inclusion of additional control variables.

Next, in Column 2, I compare homeowners to renters, and again, we see a reduced leaving rate for homeowners but an increased one for renters. This result is consistent with that in Panel B, Figure 2.5, even when I control for other worker characteristics. Then, I add combinations of factors interacted with the post dummy.

In Column 3, I include both age and homeownership simultaneously. The increase in the log odds of the young age group is reduced from 0.094 to 0.072 (albeit insignificantly), while

Table 2.2: Regressions with multiple covariates

Bin	(1)	(2)	(3)	(4)	(5)	(6) Net worth	(7) Labor inc.
post_t	−0.094***	−0.012***	−0.16***	−0.27***	−0.25***		
	(0.020)	(0.018)	(0.022)	(0.031)	(0.031)		
$\text{post}_t \times \text{young}_i$	0.094***		0.072***	0.057**	0.061**	0.085***	0.051*
	(0.026)		(0.027)	(0.027)	(0.027)	(0.027)	(0.027)
$\text{post}_t \times \text{renter}_i$		0.17***	0.16***	0.20***	0.21***	0.18***	0.22***
		(0.026)	(0.026)	(0.028)	(0.028)	(0.031)	(0.029)
$\text{post}_t \times$ rel. in S_i				0.15***	0.14***	0.14***	0.15***
				(0.028)	(0.029)	(0.030)	(0.029)
$\text{post}_t \times$ recently unemployed_i					−0.22***		
					(0.043)		
$\text{post}_t \times \text{bin}_{1,i}$						−0.39***	−0.32***
						(0.042)	(0.040)
$\text{post}_t \times \text{bin}_{2,i}$						−0.23***	−0.26***
						(0.038)	(0.040)
$\text{post}_t \times \text{bin}_{3,i}$						−0.28***	−0.30***
						(0.041)	(0.041)
$\text{post}_t \times \text{bin}_{4,i}$						−0.30***	−0.24***
						(0.044)	(0.042)
$\text{post}_t \times \text{bin}_{5,i}$						−0.047	−0.23***
						(0.051)	(0.042)
Pseudo R-squared	0.13	0.13	0.13	0.13	0.13	0.13	0.13
Num. obs.	1,123,621	1,123,621	1,123,621	1,123,621	1,123,621	1,123,621	1,123,621

Note: This table presents the change in the log odds of a worker leaving Stavanger combining several worker characteristics. In all regressions, I control for the worker being young (25–33 years old), being a renter, having family ties in Stavanger, being recently unemployed, and the previous year's binned labor income and net worth, calculated within age group and year. All errors are cluster-robust at the individual level.

the difference in the effect of post $\times$ renter is minute (0.17 to 0.16). Furthermore, when I add post $\times$ rel. in S, the probability increase for renters is amplified, and the change for the young is mitigated (Column 4).

In the event of layoffs, employers prioritize retaining workers with longer tenure, as stipulated by Norwegian labor practice.[11] Tenure with a company can correlate with housing tenure or age, where renters and younger workers face a higher risk of being laid off, which could motivate them to seek new opportunities in other locations and could explain some of the observed variation. However, the horse-race regressions demonstrate that my controlling for being recently terminated does not significantly alter the effects observed among the young age group or renters (Column 5). However, there is a reduction in the leaving probability among the group of recently terminated. This can be attributed to changes in this group's composition. The oil price shock led to unemployment among workers who had rarely experienced it before. Those who during normal economic times experience unemployment exhibit a generally higher tendency to relocate, even in years when they do not receive unemployment benefits (I discuss this further in conjunction with Table 2.A.1). Thus, when workers who are less mobile and less unemployed enter the group of unemployed, the moving rate naturally decreases.

Another correlated characteristic of the young age group and individuals with varying housing wealth is their net worth. As argued in Bilal and Rossi-Hansberg (2021), workers may divest from their location choices when facing financial constraints and income shocks. However, in Column 6 of Table 2.2, we observe the opposite effect, where the lowest-net-worth bin shows the most significant reduction in leaving probability while individuals in the highest-net-worth bin exhibit an insignificant, less pronounced reduction. The quintiles defining these bins are calculated at annual frequency within each age group. This suggests that the combination of income shocks and other general equilibrium effects is important to consider.

[11] This practice is regulated in the basic agreement (*Hovedavtalen*) resulting from periodic negotiations between Norwegian labor unions and business and industry confederations.

Figure 2.7: Leaving probabilities by labor income

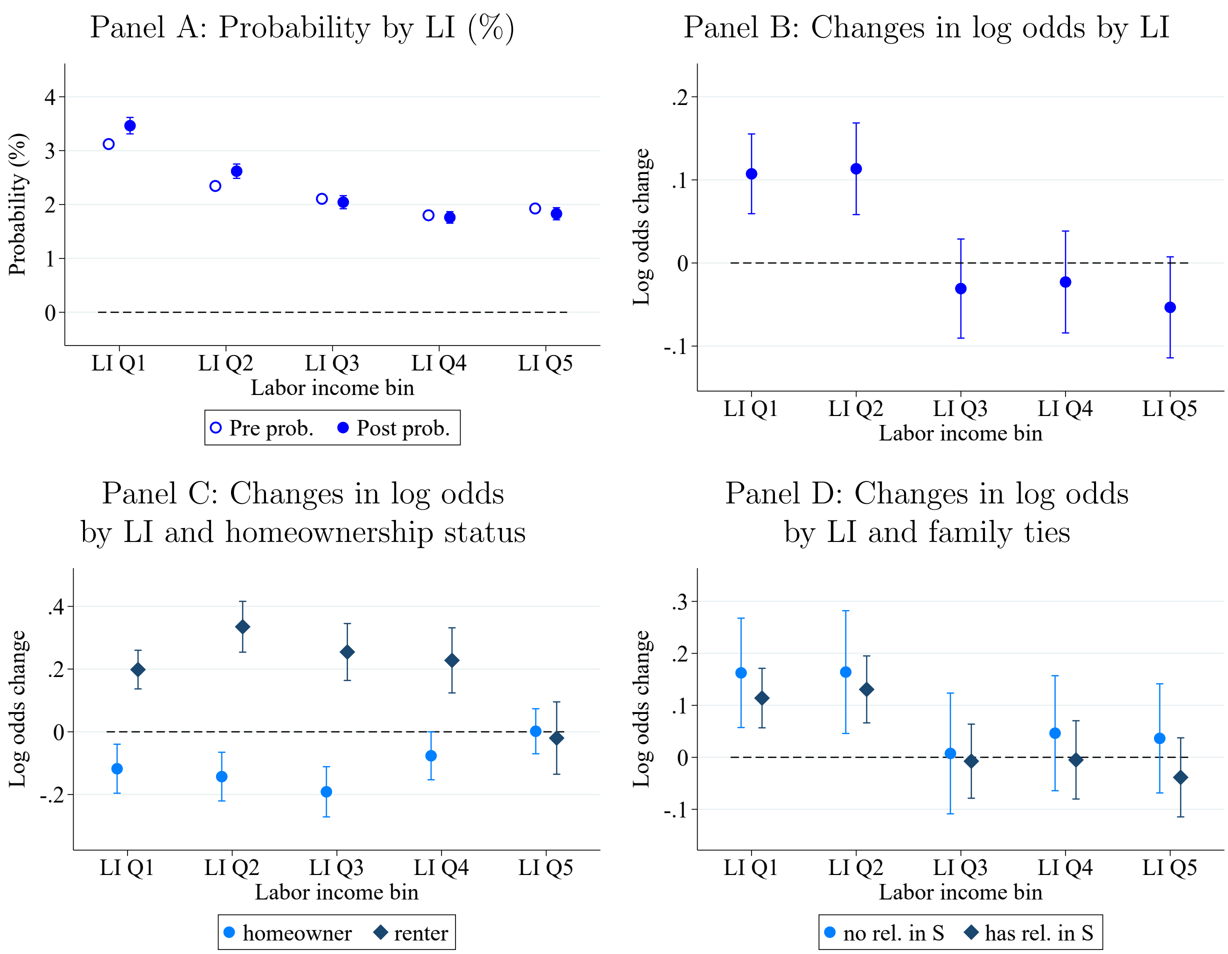

Note: This figure presents changes in the probability or log odds of out-migration from Stavanger following the oil price plunge of 2014. Panel A is produced by the OLS version of (2.13), Panels B–D are produced with logit, and Panels C and D use a third interaction term indicated by the corresponding legend. All error bands are 95% cluster-robust standard errors.

The horse-race analysis also explores effects across labor income bins, determined based on quintiles calculated for each age group and year (see Column 7). The reduction in the leaving odds is relatively consistent across these bins.

For completeness, I present the estimates of the logit model (2.13) in Figure 2.7, which illustrates an increase in the leaving rate among workers in the lower income bins and a reduction among those in the higher income bins (Panels A and B). However, when I split the income bins by homeownership status (Panel C), it again becomes evident that renters are more likely to leave whereas homeowners are more inclined to stay across all income bins except the top one. Combining these findings with the horse-race results, I conclude that the disparate behaviors across income bins are primarily explained by differences in homeownership status. In Panel D, I further stratify by family ties, revealing no discernible difference.

To summarize, renters left Stavanger at a higher rate following the shock to the region. This is not explained by renters' being different in terms of age, previous income, presence of relatives in the region, or incidence of unemployment. In the horse-race regressions, including different covariates has very little effect on the magnitude of the estimated change for renters. In contrast, the young also left at a higher rate than before the shock, but this difference is explained to a large degree by their other characteristics. Last, other confounding factors may exist, but they have to be orthogonal to the factors that I already have tested for if they are to explain the heterogeneity across housing tenure.

2.4.2 Changes in leavers' outcomes

The choice to move involves not only an *whether* to move but also *where*. Another margin along which movers can adjust is whether to buy a home in the destination. I present in this section the results of changes in these decisions together with the labor market outcomes of people who leave their LMA. Following the shock, movers are more likely to move to locations with higher incomes and home prices and exhibit a much lower probability of becoming homeowners in their destination.

To compare destinations, I calculate the destination's mean labor income and the average home transaction price in the labor market region over the period 2011–2013. I find that workers who left Stavanger moved to destinations where labor income and home prices were, on average, 0.0084 and 0.029 log points higher, respectively, than before the shock. The statistics are presented in Columns 1 and 2 in Table 2.3, and the model specification is

$$\log y_{dt} = \beta \, \mathrm{post}_t + \eta \, X_{it} + \alpha, \tag{2.16}$$

Table 2.3: Outcomes of leavers

	Destination mean income	Destination log mean home price	Buys a dwelling		
	(1)	(2)	(3)	(4)	(5)
post_t	0.0084**	0.026**	−0.45***	−0.080***	−0.44***
	(0.0036)	(0.012)	(0.044)	(0.019)	(0.0040)
leaving S_{it}				0.13***	1.18***
				(0.045)	(0.083)
$\text{post}_t \times$ leaving S_{it}				−0.38***	−0.017***
				(0.047)	(0.052)
Sample	S leavers	S leavers	S leavers	All leavers	In S at start of year
Num. obs.	9158	9158	9158	196,099	753,962
(Pseudo) R-squared	0.066	0.071	0.085	0.082	0.067

Note: This table presents changes in destination characteristics and the probability of purchasing a home after a move of different subgroups following the oil price plunge. In Columns 1 and 2, I present the estimates from model (2.16). In Columns 3–5, I present estimates from the logit model (2.17). In all regressions, I control for a fourth-degree polynomial in age, previous year's income bins, homeownership status–by–housing wealth bins, net worth bins, and the number of nonworking family members bins. The errors are cluster-robust at the level of the destination municipality.

where X_{it} contains the control variables age_{it}, age_{it}^2, age_{it}^3, age_{it}^4, last year's post-tax income bin, labor income bin, housing wealth bin, net worth bin, and number of nonworking family members. The documented change in the composition of movers motivates my use of controls, but the results are robust to their exclusion.

For the housing decisions at the destination, I estimate a logit model. This is motivated by my finding above of the important role of housing wealth. The estimates are presented in Columns 3–5, and the model is

$$\log \frac{\mathbb{P}(\text{LMA-move}_{it})}{1-\mathbb{P}(\text{LMA-move}_{it})} = \alpha + \beta\, \text{post}_t + \eta\, X_{it} + \varepsilon_{it}. \tag{2.17}$$

I include the same control variables as in (2.16), again to rule out that it is changes in mover characteristics that explain the results.

I find that workers who leave Stavanger are much less likely to buy a house or apartment at their destination than they were before the shock (Column 3). The reduction is highly economically significant. To rule out that overall pessimism or shocks to local housing supply are changing the home-buying behavior, I contrast people leaving Stavanger with other LMA-leavers in the pre- and post-periods (Column 4). This reduces the effect's magnitude somewhat, from -0.45 to -0.38, and generally, movers buy fewer homes, but the change among Stavanger leavers is much greater than that among other leavers. Both these results hold when I include destination fixed effects to account for the change in where movers go and the destinations' local housing market conditions. If I contrast Stavanger-leavers with Stavanger-stayers (Column 5), I find that Stavanger residents in general have a lower home-buying probability in the post-period.

2.4.3 Changes in the Stavanger arrivals

Overall migration to Stavanger fell dramatically in 2015 and stayed depressed throughout the episode. Like those in out-migration, the changes in in-migration exhibit great heterogeneity. I find that renters and homeowners with little housing wealth change their in-migration rate little, or even increase it if they are renters, while workers of higher incomes reduce their arrival rate. Older people with family ties in the region exhibit a near-zero reduction in the odds of arriving in Stavanger and the arrival population shifts towards people who rely more on welfare transfers.

The levels are overall of smaller magnitudes in all the figures (see the Panel As). This is because the sample now includes everyone who does not live in Stavanger and, while we previously studied moves from Stavanger to anywhere else, we now look at moves from anywhere to Stavanger. This naturally lowers the migration probabilities.

In Figure 2.8, I display the changes in arrival rates by age bin. Following the oil price plunge, it was again the young that responded the most, as measured by level changes, and they were the most frequent arrivals before the oil shock episode (Panel A). In terms of log odds, however, all age groups have a similarly reduced moving-in probability except the oldest (58–

Figure 2.8: Arrival probabilities by age

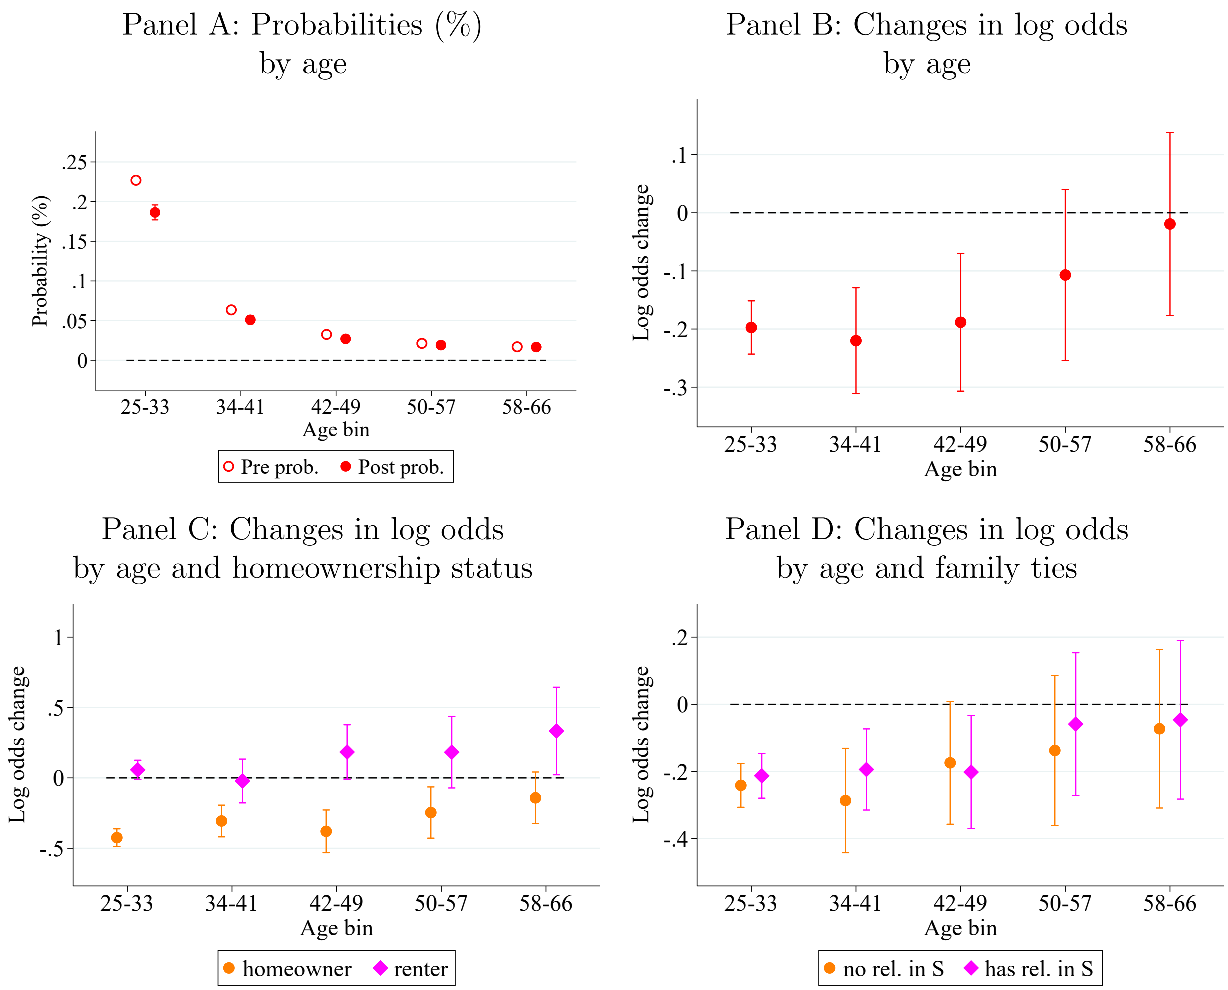

Note: This figure presents changes in the probability or log odds of in-migration to Stavanger following the oil price plunge of 2014. Panel A is produced by the OLS version of (2.13), Panels B–D are produced with logit, and Panels C and D use a third interaction term indicated by the corresponding legend. All error bands are 95% cluster-robust standard errors.

66-year-olds, see Panel B). The reduction is about 26% (30 log points) and significant for the four younger bins, while for the oldest group, it is 11% (12 log points) and insignificant. This is consistent with income differentials being a key motivator of migration (Kennan and Walker, 2011) and with workers close to retirement putting less emphasis on them. In Panel C, I contrast renters and homeowners. The latter group has a similarly reduced in-migration rate across all age groups, while renters show no reduction or even, in the case of the oldest, an increase. In Panel D, I split up age groups by the presence of family ties in Stavanger. The

Figure 2.9: Arrival probabilities by housing wealth

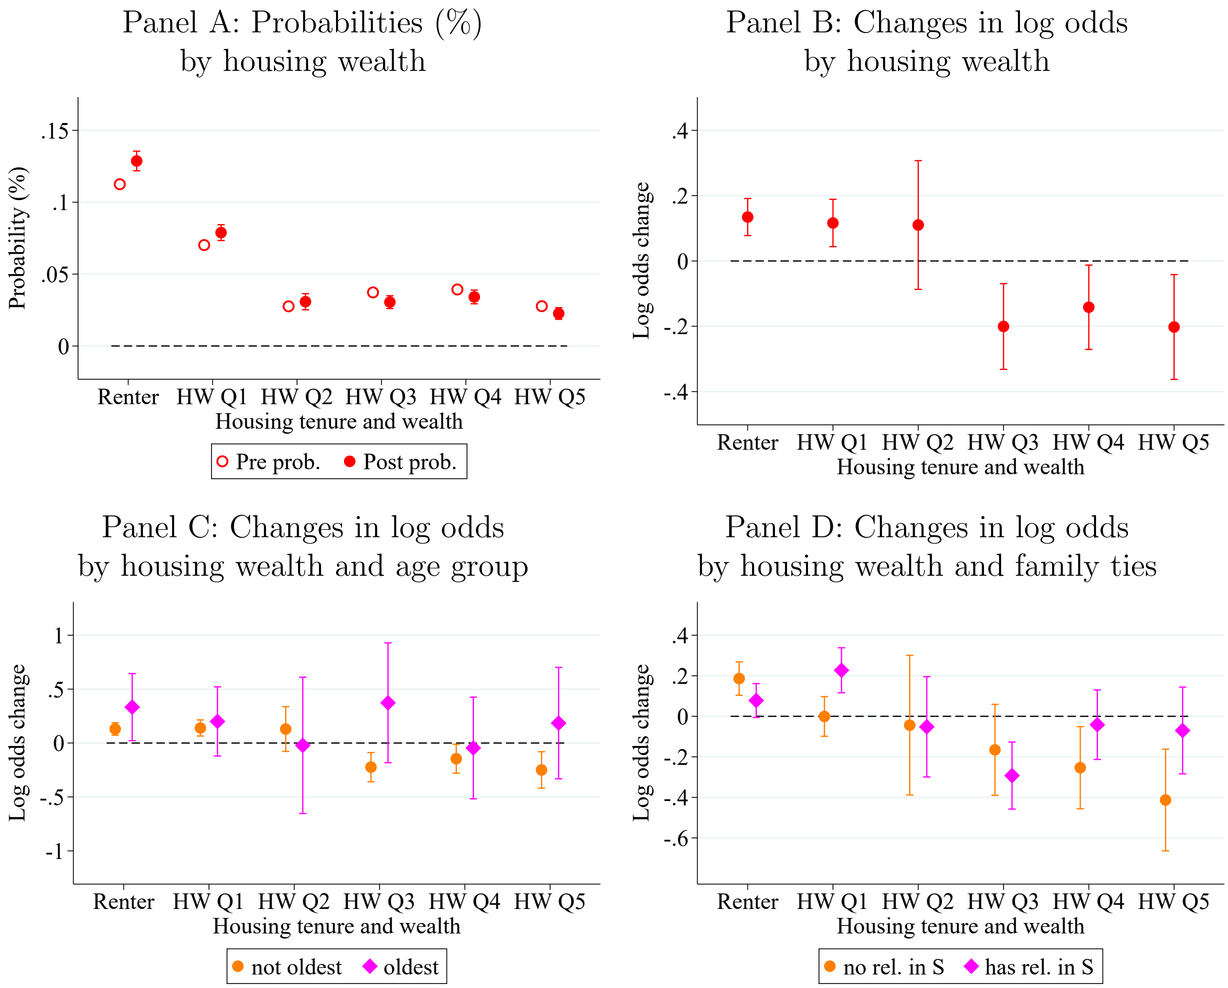

Note: This figure presents changes in the probability or log odds of in-migration to Stavanger following the oil price plunge of 2014. Panel A is produced by the OLS version of (2.13), Panels B–D are produced with logit, and Panels C and D use a third interaction term indicated by the corresponding legend. All error bands are 95% cluster-robust standard errors.

groups with ties consistently show a smaller reduction in the arrival rate, but the differences are insignificant within each age bin.

It seems that how renters and homeowners evaluated Stavanger changed differently in the post-period. To explore this, I first split up homeowners by housing wealth bin and then by labor earnings. In Figure 2.9, results for the former subsamples are displayed. From Panels A and B, we can conclude that the reduction among homeowners is driven by the richest three bins (−57%). The panels also display a small increase of 4.7% among the workers with

Figure 2.10: Arrival probabilities by labor income

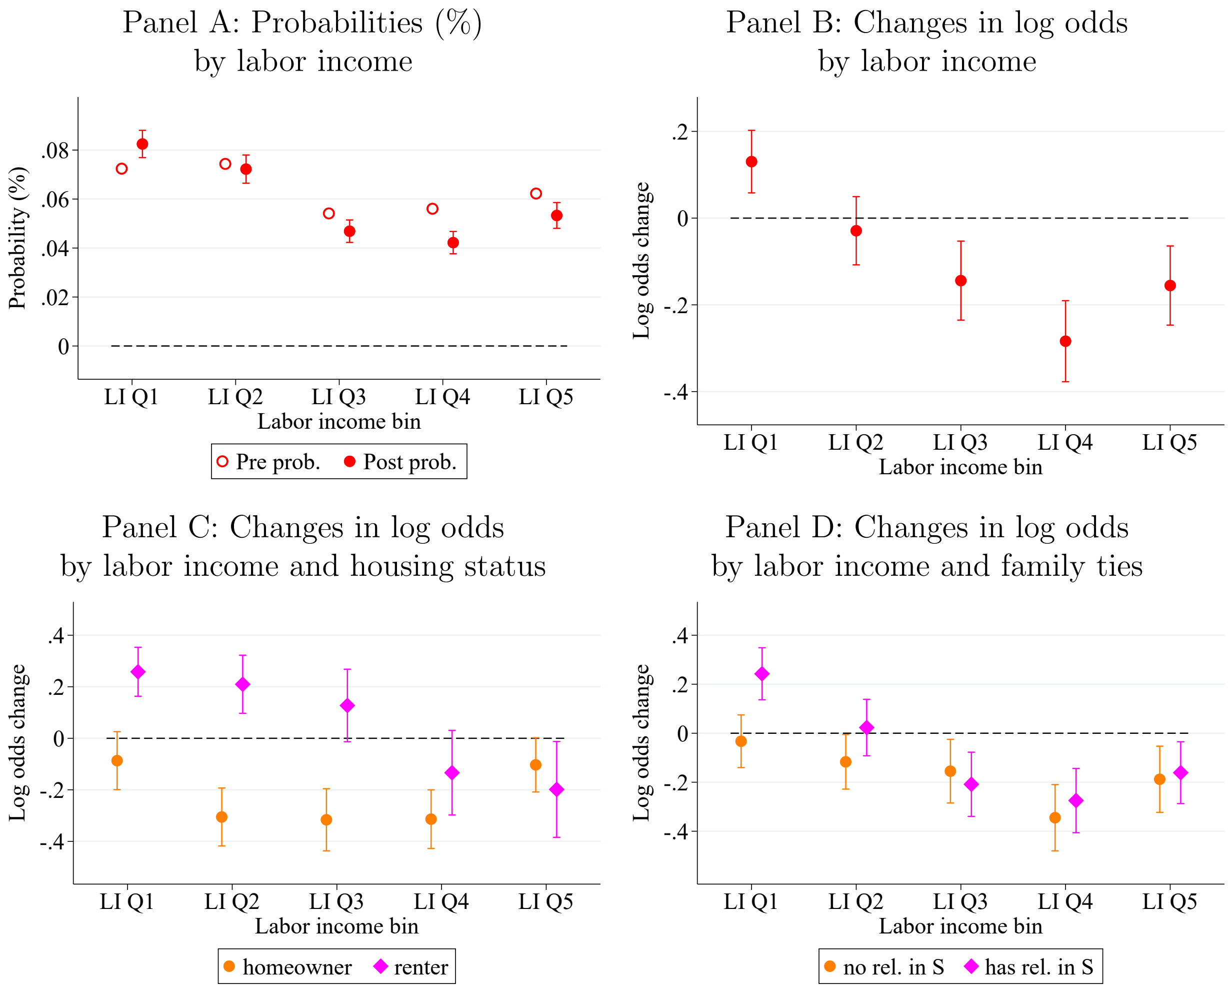

Note: This figure presents changes in the probability or log odds of in-migration to Stavanger following the oil price plunge of 2014. Panel A is produced by the OLS version of (2.13), Panels B–D are produced with logit, and Panels C and D use a third interaction term indicated by the corresponding legend. All error bands are 95% cluster-robust standard errors.

little housing wealth. Panel C in turn illustrates again that people in the oldest bin respond by moving in more across all housing wealth bins, albeit with mixed levels of significance. In Panel D, we observe a starker difference for those richest in housing wealth when we compare workers with and without family ties in Stavanger. Those with relatives in the area and with the most housing wealth display an insignificant decrease in their arrival rate. However, these results do not explain the difference between renters and homeowners.

In Figure 2.10, I present the changes by labor earnings bin. Panels A and B show that the reduction in the moving-in odds is large and significant for all groups except the lowest income bin. This is driven by homeownership status; all renters except in the two highest income bins move in at a rate higher than or similar to their rate before, while all homeowners show a reduced rate (see Panel C). In Panel D, I split the sample by the presence of family ties in Stavanger. People without ties have a reduced arrival rate across labor income bins, but the reduction for people with relatives is smaller. In the two lowest income bins, there is no reduction and even an increase among the poorest. In Table 2.A.2, I test if this is related to the uptake of social welfare in terms of government transfers. I find that Stavanger arrivals in the post-period rely to a higher degree on social welfare, also when controlling for labor earnings, age, and origin LMA. Also when I condition the sample to only study low-income workers the the rise in the share of income from government transfers is significant, both statistically and economically. The rise is 5.5–6.6 percentage points, or, 8.7–26 log points (the ranges are across regression specifications). This group should be less impacted by the regional income shock since social transfers are managed at the national level in Norway.

I conclude that the smaller response of renters and lower-housing-wealth homeowners is because they earn less, are less impacted by the worsened income prospects in Stavanger due to higher reliance on welfare, and are relatively better compensated by the cheaper housing in the region.

2.5 A life-cycle model with location choices

To quantify the importance of the response of home prices in explaining the heterogeneity in migration, perform welfare analysis, and create an environment for policy experiments, I set up a spatial model similar to the models in Kennan and Walker (2011) and Giannone et al. (2023). It has intertemporal decision-making in the form of financial savings and housing choices, and as in Kennan and Walker (2011), workers have location preferences. I also add worker-specific skills and locations have different skill premia so that workers of different skills value locations differently (compare to Bilal and Rossi-Hansberg, 2021). This produces heterogeneous migration patterns. Workers have perfect information except about the impending oil price plunge, which is modeled as an MIT shock to the income process.

2.5.1 Economic environment

The model is an open economy that consists of a set $\mathcal{L}$ of labor market areas, where individual LMAs are denoted by ℓ. Each location has a fixed housing supply H^{ℓ} and some square meter price $p^h(\ell)$ that is determined in equilibrium. A foreign landlord can buy square meters and rent them to workers at an annual rate $\omega^R \times p^h(\ell)$ per square meter. The largest rental is h^R, and the smallest owned house is $h^{HO} > h^R$. All extracted value leaves the country.

The stock of housing is continuously maintained at the same rate as it depreciates by a foreign firm. That is, the quality of the housing stock is also constant. However, the cost of this is $\delta^m \times p^h(\ell)$ per square meter and is paid by the owner of the unit (i.e., renters do not pay for maintenance). All housing requires additional utilities such as water, electricity, and insurance, which cost the resident $\delta^u \times p^h(\ell)$ per square meter (i.e., are paid by both homeowners and renters).

In each location, workers earn a base wage $\mathrm{LP}(\ell)$, and if they have skill s, they earn an additional skill premium $s \times \mathrm{SP}(\ell)$. Over the life-cycle, the wage follows the curve $g(q)$, where q denotes age. For a worker who lives in ℓ, she earns after taxes

$$y(s, q, \ell) = g(q) \exp(\mathrm{LP}(\ell) + s \times \mathrm{SP}(\ell)). \tag{2.18}$$

The income process does not exhibit income risk unless we consider the risk to income from random migration.

(Foreign) banks are willing to lend only to homeowners and there is a cap φ on the loan-to-value ratio. Borrowers pay an interest $r^m > r^s$, where r^s is the return on savings.

It is costly to buy and sell property. A share $\eta^{h,\mathrm{sell}}$ of the value per square meter is lost at a sale, and an additional share η^H per square meter has to be paid when a worker buys a home. Renters do not pay adjustment costs if they stay in the location.

2.5.2 The worker's problem

Workers live from 25 to 66, and their age is denoted by q. Their skill is $s \geq 0$, and they have a location preference $\ell_f \in \mathcal{L}$. Every period, they wake up in a location $\ell_o \in \mathcal{L}$ with housing $h \in [\underline{h}, \overline{h}]$ and savings $a \geq \underline{a}(h, p^h(\ell_o))$. If $h \leq h_R$, the worker is a renter, and otherwise, she is a homeowner. Decisions are made annually.

Every year, workers choose to either stay or move to another location. The decision is denoted by ℓ^d, after which they pick how much to consume (c) and how much housing they want to rent or buy (h). To emphasize that these choices depend on the location choice, I use the superscript d below. What remains of the available cash on hand is carried over to the next period as savings b with return $r(b)$. Every location provides some level of utility $A(\ell)$ for free, which I refer to as amenities. A worker who moves is subject to the disutility τ and has to pay a monetary moving fee $\eta^\ell(\ell_o, \ell^d)$. Anytime a homeowner leaves her current location, she has to sell the current home.

By living in their preferred location, workers obtain additional per-period utility γ_f, and in every period, they are hit with a vector of preference shocks $\bar{\xi}$ for each location. The elements of the vector are denoted ξ^d and are described shortly.

The present value of choosing the optimal destination given the worker's state (s, ℓ_f, q, ℓ_o, a, h, $\bar{\xi}$) is denoted by V. The mathematical formulation of the worker's problem is

$$
\begin{aligned}
V(s, \ell_f, q, \ell_o, a, h, \bar{\xi}) = \max_{\ell^d, c^d, h^d, b^d} \Big\{ & u(c^d, h^d) + \gamma_f \mathbb{1}(\ell^d \in \mathcal{L}(\ell_f)) \\
& + A(\ell^d) - \tau(\ell_o, \ell^d) + \xi^d + \\
& + \beta\, \mathbb{E}_\xi [V(s, \ell_f, q+1, \ell^d, a', h^d)] \Big\}_{\ell^d \in \mathcal{L}},
\end{aligned} \tag{2.19}
$$

$$
y(s, q, \ell^d) = g(q) \exp(\mathrm{LP}(\ell^d) + s \times \mathrm{SP}(\ell^d)), \tag{2.20}
$$

$$
a' = b^d(1 + r(b^d)), \tag{2.21}
$$

where the budget and borrowing constraints depend on the homeownership status and whether the worker is adjusting her housing. In the case of an intra- or interlocation move, I use an auxiliary variable called cost-adjusted cash-on-hand, which also includes the value of housing if the worker is a homeowner:

$$x^A(\ell_o, \ell^d, h) = a + (1 - \eta^{h,\text{sell}})\, h \times p^h(\ell)\, \mathbb{1}(h \geq h^{HO}) - \eta^\ell\, \mathbb{1}(\ell^d \neq \ell_o). \tag{2.22}$$

The different budget constraints are

$$c + (\omega^R + \delta^u)\, p^h(\ell^d)\, h^d + b = x^A(\ell_o, \ell^d, h) + y(s, q, \ell^d), \quad \text{if } h^d \leq h^R, \tag{2.23}$$

$$c + (\delta^u + \delta^m)\, p^h(\ell^d)\, h^d + b = a + y(s, q, \ell^d), \text{ if } h \geq h^{HO} \wedge h^d = h, \tag{2.24}$$

$$c + (1 + \eta^{h,\text{buy}} + \delta^u + \delta^m)\, p^h(\ell^d)\, h^d + b = x^A(\ell_o, \ell^d, h) + y(s, q, \ell^d), \quad \text{if } h^d \neq h \wedge h^d \geq h^{HO}, \tag{2.25}$$

where the first case is that of the worker who ends the period as a renter and the beginning-of-period homeownership status is captured by $x^A(\ell_o, \ell^d, h)$. The second case is that of a homeowner who does not adjust her housing. The last case is that of a worker who decides to buy a home, where again the initial status is reflected by $x^A(\ell_o, \ell^d, h)$.

A worker who chooses to be a renter faces the no-borrowing constraint:

$$a \geq 0, \tag{2.26}$$

while a homeowner can borrow using a mortgage that respects the loan-to-value (LTV) constraint:

$$a \geq -\varphi\, p^h(\ell^d) \times h^d. \tag{2.27}$$

The utility function is

$$u(c, h) = \frac{(c^{1-\alpha}\, (\kappa^{HO}\, h)^\alpha)^{1-\sigma}}{1 - \sigma}, \tag{2.28}$$

where κ^{HO} captures the additional utility of owning one's home. In the last period, there is also additional utility from the remaining cash-on-hand,

$$\Phi(b) = \frac{\phi_0 \times b^{1-\sigma}}{1-\sigma}, \tag{2.29}$$

$$0 < b < x^A(\ell^d, \ell^d, h^d) - c - cost\ of\ housing. \tag{2.30}$$

The constraints make it so that the worker cannot end her working-life indebted. This should be thought of not solely as a bequest but also as an incentive to save for retirement. The choice of functional form is different from that in, e.g., De Nardi (2004), which allows for no bequest, while my functional form rules it out. Alternatively, I could solve for retirement, have a bequest in the case of death, and allow retirees to migrate between locations, but I assume that this is a negligible feature that does not materially affect the results of the model.

All units are in mean annual post-tax income, which in Norway was 322,600 NOK (55,240 USD) in 2010–2013.

The vector of everyday location-preference shocks is a key feature of the model. They are Gumbel distributed with a scale parameter ν that is the same across all locations. The distribution is also known as a type-1 generalized extreme value distribution. The choice of distribution allows for a closed-form expression for the expectation value of the value function and the transition probabilities, given the value function and the fact that the workers pick the utility-maximizing location (see Section 2.C, or McFadden, 1974).

For the purpose of modeling migration, preference shocks combined with moving costs allow the model to show low rates of migration (due to it being costly to move) when there are clear economic benefits to relocating in the form of income and home price differentials. The preference shocks nudge a share of workers to make the move despite the high costs. The size of ν is relatively more important for older than for younger workers in driving moving decisions. The preference shocks can also yield migration decisions that take workers to worse locations, as happens in real data.

The added randomness also makes computing the housing price equilibrium easier. In the case of no preference shocks but significant moving costs, small changes in prices do not always induce a small change in the moving rate between locations; rather, they sometimes trigger large changes if the state space of workers is not fine enough. This in turn creates large swings in housing demand across locations. Through my adding preference shocks, there is always a small flow by every worker type to every location, and adjusting prices changes these probabilities by small amounts, lessening the swings in housing demand, and facilitating finding the equilibrium.

2.5.3 Worker skills and wage premia

The skills used in the model are as of now abstract objects. This section describes how I estimate them from the microdata. This makes it possible to study how worker composition changes in response to different economic shocks. Skills also stand in, in part, for worker–location match quality.

I assume that the income process takes the form

$$\log \text{income}_{it} = \alpha_0 + \sum_{\ell \in \mathcal{L}} \text{LP}_\ell \times \mathbb{1}(\ell_{it} = \ell) + \sum_{\ell \in \mathcal{L}} \text{SP}_\ell \times \mathbb{1}(\ell_{it} = \ell) \times s_i + \eta\, X_{it} + \varepsilon_{it}, \tag{2.31}$$

where ℓ is the location, LP_ℓ is a basic income difference of workers of $s_i = 0$ across locations, and SP_ℓ is a location-specific skill premium that is linear in skill. Equation (2.31) is based on the model in De la Roca and Puga (2017) but without assuming that the benefit of a location is proportional to the population size; this specification allows for more flexible local wage premia. I outline a fixed-point algorithm to solve the nonlinear system in $(s_i, \text{LP}_\ell, \text{SP}_\ell)$ in Section 2.B, where I also describe the data selection used to estimate the model. In the appendix, I also discuss how to separate $(\alpha_0, s_i, \text{SP}_\ell)$ that are only jointly identified.[12] I control for age using a fourth-degree polynomial represented above by X_{it}. I assume that

[12] In brief, there are infinitely many combinations of $(\alpha_0, s_i, \text{LP}_\ell, \text{SP}_\ell)$ that yield the same predicted $\log \text{income}_{it}$. However, they are connected through an affine transformation. Thus, I ensure that the lowest skill s_i is zero, other skills are non-negative, and the variance of skills across individuals is 1 and adjust the SP_ℓs and the intercept α_0 accordingly.

other relevant worker characteristics are constant across time and thus absorbed by individual fixed effects. The parameters are identified by workers being observed in different locations. I estimate the skills of workers who are never observed moving by comparing the worker to others of similar age who at some point move to a different LMA.

The individual skill reflects workers' abilities that are constant across time but that pays off differently across locations. I do not control for industry effects to avoid controlling for high-skill individuals' selection into specific jobs. I ignore the accumulation of experience due to location and age, and I treat the level of education as constant (thus, absorbed by individual skill).

2.5.4 Model estimation

The discrete choices of the worker's problem give rise to not strictly concave value functions, rendering the standard endogenous grid method not applicable. However, by applying an upper envelope step as described in Druedahl (2021), the endogenous grid method can be applied to the problem of the worker who only chooses to save and consume (i.e., who does not move or adjust their stock of housing). I refer to this case as being passive. For the problem of workers who adjust their housing or move, I exploit the nested structure of the problem (again, see Druedahl, 2021) and compute the value of moving or adjusting housing by interpolating the passive worker's problem.

To limit the size of the state space, I merge several LMAs by geographical proximity and similarity of characteristics, from 46 down to 7.

A set of the model coefficients is not estimated with the model but comes from external sources. The values are presented in Table 2.4. The cost of selling an owned dwelling is taken from Kaplan et al. (2020) and is a relative loss of 7%. When a worker buys a house, the government imposes a documentation fee of 2.5%, which I use to proxy the home-buying cost ($\eta^{h,\text{buy}}$). As in Berger et al. (2018), I assume a constant price–rent ratio (ω^R), which I estimate using data from Statistics Norway. I compare per-square meter rents to per-square meter home prices in the period 2010–2013 and arrive at 0.0699, which is close to the 0.06

Table 2.4: Parameter values from external sources

Variable	Name	Value	Source
$\eta^{h,\text{sell}}$	Home-selling cost	7.0%	From Kaplan et al. (2020)
$\eta^{h,\text{buy}}$	Home-buying cost	2.5%	Administrative fee
ω^R	Rent share	6.99%	Estimated outside model using Statistics Norway Tables 05963 and 09895
η^ℓ	Location-adjustment cost	0.221	Estimated using expenditure data
σ	Degree of consumption smoothing	2.0	Standard assumption
r^m	Mortgage interest rate	3.98%	Statistics Norway Table 10748, Dec 2013
r^s	Saving interest rate	1.05%	From Fagereng et al. (2020)
φ	LTV cap	85.0%	Legal requirement, see Aastveit et al. (2022)

Note: This table presents the model parameters that are either estimated without the model or taken from external sources.

in Berger et al. (2018). Using household level expenditure data (see Aastveit et al., 2024), I estimate the cost of moving between locations by regressing the annual expenditure of households on a set of year-fixed effects and a dummy indicating if the household moved in the current, previous, or following year. The interest rate on the mortgage is from Statistics Norway, and the returns on savings are from Table 3 in Fagereng et al. (2020), where I assume that the average return on financial wealth in that paper represents the return of the same portfolio the workers in my setting have available. The cap on LTV φ has been changing over time, but for most of the relevant period, it was 85% (see Aastveit et al., 2022).

The remaining parameters are estimated using the simulated method of moments by matching several moments of the life-cycle profile of Norwegians in 2010–2013. The targets are listed in Table 2.5, and the estimated parameter values are listed in Table 2.6. I simultane-

Table 2.5: Data moments targeted by model estimation

Description	Target value	Source	Simulated value in calibration
Cash-on-hand of 35–45-year-old	−2.18	Microdata	−2.25
Cash-on-hand of 60–66-year-old	0.0874	Microdata	0.0369
Expenditure share of housing expenses	31.2%	SSB CEX	31.5%
Expenditure share of utilities	5.8%	SSB CEX	5.80%
Expenditure share of maintenance costs	5.7%	SSB CEX	5.71%
Share living in preferred location	77.2%	Microdata	68.0%
Share homeowners	74.6%	Microdata	75.3%
Average inter-LMA moving rate	1.81%	Microdata	1.74%
Average inter-LMA moving rate of 57–66-year-old	0.633%	Microdata	0.700%

Note: This table presents the targets of the model estimation, the sources of the targets, and the simulated values of the final estimation. Cash-on-hand is expressed in terms of the sample average post-tax income, which is 322,600 NOK (55,240 USD). Housing expenses are the sum of the cost of home maintenance (paid by homeowners), the cost of utilities (paid by everyone), and interest (paid by borrowers). Moving rates are annual. SSB CEX refers to Statistics Norway's survey of consumer expenditure (Strand, 2014).

ously estimate home prices and update the vector of prices based on the excess or shortage of local housing demand.

The estimation works as follows: For every guess of parameters and home prices, I solve the model and simulate 1000 life paths for every worker in an initial sample of 25-year-olds. The sample is the distribution in 2010–2013. I compute the targeted moments across the simulated sample. Then, I use the full population of 25–66-year-olds in 2010–2013, estimate using the model their migration decisions and housing demand in the period, and compute the excess demand of housing per location. The housing supply is calculated using data from Statistics Norway. For locations with positive excess demand, prices are increased; if excess demand is negative, prices are lowered.

There is a tension between calibrating parameters to the life-cycle profile of migration and matching migration in the cross-section. Forcing an economic model on data on location decisions and other economic variables can make certain observed states highly implausible. The worker's present value in such states is much lower than in other accessible locations, and thus, the worker will have a high leaving probability. An alternative estimation strategy is to use maximum likelihood, however, several parameters relate to expenditure shares which the transition probabilities between states are not informative of. Therefore, I use the simulated method of moments as described above.

2.5.5 A grid-transformation trick

Because of the possibility of binding borrowing constraints, I use a nonuniformly spaced grid for savings, where grid points are more concentrated closer to the constraint. Because home-owners' borrowing constraint depends on the home value, I express savings as a share of the total home value (this can also be used for renters). This requires mapping nominal savings into the grid several times, but the calculation is straightforward, and the transformation of the grid has the benefit of not having to be very dense for all possible negative nominal values of borrowing.

Mathematically, let $\mathfrak{a}^*$ denote the grid point used in the numerical solver that corresponds to savings a. Then,

$$\mathfrak{a}^* \equiv \frac{a}{h \times p^h(\ell)}. \tag{2.32}$$

The savings-per-housing-value $\mathfrak{a}^*$ is then on the grid $\{-\varphi, \mathfrak{a}_2, \ldots, \mathfrak{a}_{N_a}\}$, where $\mathfrak{a}_{N_a}$ is a high enough number to rarely be reached. Since all value functions are solved within a location, interpolating between $\mathfrak{a}^*$s given the location is as accurate as interpolating between the corresponding as.

Unrelated to the grid transformation, the only issue of interpolation to consider is the risk of interpolating between the biggest rental h^R and the smallest owned house h^{HO} when solving the home adjuster's problem when using the nested value function method. The problem is avoided by splitting up the problem into two, one for the worker who chooses to be a renter, and one to be a homeowner.

The spacing between points grows exponentially, and because renters face a no-borrowing constraint, I manually add a point $\mathfrak{a}_{i_0} = 0$ and additional grid points above to the grid to cover the region close to the constraint.

2.5.6 Main model results

I simulate the impact of the oil price shock on moving rates and welfare by reducing the base wage LP in Stavanger by 6% and compute the new vector of home prices that clear all housing markets, taking the current population distribution and the model parameters as given. This leads to a reduction of 14% in Stavanger home prices and a small increase of 0.29% across other locations as the demand for them increases (the rise is the mean across regions, weighted by housing supply). This should be compared to the relative change in the price differential of about 25% documented in Section 2.2. The following sections present the decomposition of the labor market shock and the home price shock, the heterogeneity in migration the model produces, welfare analysis, an analysis of the efficacy of moving subsidies, and the equivalent changes in moving costs the reduction in home prices produces.

Table 2.6: Model parameter estimates

Variable	Name	Value
β	Time discount factor	0.974
ϕ_0	Bequest motive	32.7
α_H	Housing utility parameter	0.477
τ	Disutility of moving	0.575
ν	Preference-shock parameter ($\approx$ standard deviation)	0.143
γ_f	Family-proximity bonus	0.0286
δ^u	Cost of utilities	0.00787
δ^m	Cost of maintenance	0.00924
κ^{HO}	Additional utility of homeownership	1.06
$A(\ell)$	Location bonus (amenities)	See appendix

Note: This table presents the model parameters that are estimated within the model, targeting the moments in Table 2.5. The location-adjustment cost is in units of mean annual post-tax incomes. Over the period 2010–2013, this value was 322,600 NOK (55,240 USD).

2.5.6.1 Changes in migration

Similarly to that in the data, the leaving probability rises only modestly, by 2.6% (compared to 0.37% in 2015–2018 or 5.6% in 2015–2016), and the arriving probability falls by 31% (compared to 30% in 2015–2016). If I hold home prices fixed, the changes are significantly greater, as illustrated in Table 2.7. The change is expressed in terms of the log odds, and in Column 1, the odds for leaving Stavanger increase by 0.50 log points. This corresponds to an increase in the leaving probability of 29%. If we use the preshock leaving-Stavanger probability from Table 2.1 for comparison (2.5%), the leaving probability is 4.2% following a shock to only the base wage, while with the home price adjustment, it is 2.6%. The change in the arrival rate also depends on the re-adjustment of home prices. Without this, the change in the log odds is -1.3, and with it, it is -0.38—a reduction of -72% versus -31%, respectively (see Column 3 and 4). Even with the large home price re-adjustment, the probability of moving to Stavanger is much lower than before.

Table 2.7: Changes in simulated migration, by type of regional shock

	Changes in leaving		Changes in arriving	
	(1) LP shock	(2) LP + HP shock	(3) LP shock	(4) LP + HP shock
Δ log odds	0.50	0.045	-1.3	-0.38

Note: This table presents the model-implied log changes in leaving and arrival odds for either the shock to LP in Stavanger or the joint shock to both LP and home prices in Stavanger (LP + HP shock).

In Table 2.8, I split up the simulated response by homeownership status, and we see that renters' leaving response to the shock is stronger, in line with the fact that they are not weighed down by a loss in housing wealth. When comparing the cases without home price equilibration (Column 1) and with equilibration (Column 2), I observe that the change in the log odds is much smaller in the latter case, as expected (1.1 versus 0.16), but is still of an economically significant magnitude. Homeowners behave similarly across the types of shock, but both responses are smaller than those of renters; in the case of a shock to the wage only, the increase is 0.38, compared to 0.041 if home prices are also reduced. Unlike the situation in the data, the average change is not negative, but the behavior across housing tenures is confirmed.

If, instead of considering an accompanying shock to home prices, I consider a shock to the homeowners' savings of a magnitude that corresponds to the loss in housing wealth, there is also a reduction in the log odds change. However, the change in the moving behavior is not as large as when I let home prices fall. This may be due to how I model the reduction in wealth. In my setup, homeowners with a mortgage that violates the LTV constraint following the reduction in home prices are forgiven the excess debt. If I instead allowed them to be underwater, the housing wealth shock would be considerably more binding.

As in the data, homeowners' arrival rate is reduced by more than renters' in the simulation. The changes in the log odds are large in both the case of the wage shock and the joint

Table 2.8: Changes in simulated migration following an income and home price shock, by type of regional shock and housing tenure

	Renters		Homeowners		
	(1) LP shock	(2) LP + HP shock	(3) LP shock	(4) LP + HP shock	(5) LP + HW shock
Δ log odds leaving	1.2	0.15	0.39	0.031	0.27
Δ log odds arriving	−1.3	−0.40	−1.6	−0.55	

Note: This table presents the model-implied changes in leaving and arrival probabilities expressed in terms of the log odds across homeownership status for either the shock to LP in Stavanger or the joint shock to both LP and home prices in Stavanger (LP + HP shock). Column 5 presents the results of the experiment of shocking homeowners' total wealth by as much as their housing wealth is reduced if home prices fall.

shock, but the change for renters is −0.40, compared to −0.55 for homeowners. To provide further support that it is the change in housing wealth is an important channel, I split up homeowners by housing wealth tertiles. The results are presented in Table 2.9. As in the empirical analysis, the change in the leaving odds falls with housing wealth (see columns 2–4). For the top tertile, the is even a reduction. Also like the empirical results (however, then an insignificant difference), low-housing wealth homeowners exhibit a slightly greater increase, relative to renters (compare columns 1 and 2). The change in the arriving odds is a reduction across all bins and the magnitude is increasing in housing wealth. No group experiences a rise or non-change as in the data (see Figure 2.9), but the sorting is the same.

Note that the earnings shock I use in the simulation affects all workers, but in the empirical analysis, I also find that the composition of arrivals to Stavanger shifted towards people who rely more on government transfers, a source of income that does not depend on regional economic conditions. This can in part explain why all in-migration falls across the dimensions of heterogeneity I study, in contrast to the data.

Table 2.9: Changes in simulated migration, by housing wealth

	Renters	Homeowners		
	(1)	(2) HW 1	(3) HW 2	(4) HW 3
Δ log odds leaving	0.15	0.17	0.040	−0.21
Δ log odds arriving	−0.40	−0.50	−0.60	−0.61

Note: This table presents the model-implied changes in leaving and arrival probabilities expressed in terms of the log odds across housing wealth for a joint shock to both LP and home prices in Stavanger (LP + HP shock). The bins for housing wealth are zero (Renters), below the 1st tertile (HW 1), between the 1st and 2nd tertile (HW 2), and above the 2nd tertile (HW 3).

2.5.6.2 Welfare analysis

I quantify the welfare consequences as the equivalent variation (EV), which is the change in worker income that yields the same mean change in Stavanger workers' present values that the combined income and home price shocks produce. The EV is computed as a change in LP in all locations. The EV cannot apply only to Stavanger because workers can quickly move away from the shock, and the greater the reduction in income, the higher is the migration, and the smaller is the change in Stavanger workers' welfare.[13] Thus, to reflect the lifetime reduction in welfare, ΔLP has to follow workers originating in Stavanger as they move. The results of this exercise are presented in Table 2.10. To center our attention on the significance of local housing wealth and rents and to simplify the interpretation of the analysis, I exclude the minor increases in home prices and rents in other areas due to re-equilibration across housing markets (initially reported as a 0.29% increase). Nevertheless, this adjustment does have a slight impact on the value of relocation, affecting both groups to some extent.

[13] Imagine that the additional disutility due to a reduction in home prices, which hinders migration, is of a great magnitude. Then, the reduction in only Stavanger will have to be very large to create the same change in welfare. However, as ΔLP worsens, workers will leave at a higher rate. Their owned homes do not hold them back because there is no impact on their value in this exercise. As workers move, fewer are left to suffer from the worsened income, and ΔLP will have to become even more negative, which again drives workers away.

Table 2.10: Welfare consequences of regional shocks

	All	Renters			Homeowners		
Welfare change (%)	(1) Both shocks	(2) LP shock	(3) HP shock	(4) Both shocks	(5) LP shock	(6) HP shock	(7) Both shocks
$EV \mid$ stay	−3.0	−2.5	2.6	−0.59	−1.3	−2.3	−4.1
$EV \mid$ move	−2.6	−0.012	0.15	−0.0056	0.0000	−3.7	−3.7

Note: This table presents the model-implied changes in Stavanger worker welfare following different shocks. Welfare is measured in terms of equivalent variation EV, the percentage change in labor earnings from the no-shock case to produce the same group-average welfare as in the shock case, indicated by the column header.

The overall welfare from staying declines by 3.0% (Column 1). The incidence is worse among renters, who are generally poorer, as illustrated by the scenario with the shock only to Stavanger's LP without housing cost adjustment (compare Columns 2 and 5). The welfare change from having only the cost of housing adjusted (Columns 3 and 6) illustrates how renters benefit from lower rents, where they show a positive EV of 2.6%, while homeowners lose 2.3%. The net effect (Columns 4 and 7) shows that homeowners are worse off in general, with an EV of −4.1% versus −0.59% for renters. This seems to indicate that renters should be more willing to stay than homeowners. However, the willingness to move is determined by differentials, which I illustrate by calculating the welfare change in moving.

At the bottom of Table 2.10, I present the EV in the change of the present value of leaving Stavanger. Overall, the present value of moving to another location falls, as shown in Column 1.[14] This is explained by the disaggregation of channels in the following columns.

Leavers experience very small welfare consequences of the direct labor market shock to Stavanger (Columns 2 and 5). The magnitudes depend on the probabilities of returning and then experiencing the negative income shock. The changing cost of housing has clear differential impacts. Renters might return to Stavanger and will then benefit from cheaper

[14] I compute the present value by first calculating each worker's present values of each destination, then I take the expected value across destinations (weighting by the worker's individual moving probabilities), and then I average across workers.

Table 2.11: Corresponding changes in moving costs

	Benchmark	All movers	Renter movers	Homeowner movers
	(1)	(2)	(3)	(4)
τ	0.57	2.01	2.16	1.97
η^{ℓ}	0.22	0.86	0.66	0.90

Note: This table presents the moving costs necessary, when I hold home prices constant, to match the changes in leaving rates following the shock to LP and home prices in Stavanger.

housing by either renting or owning, which increases their welfare from leaving by 0.15%—not as much as it does when they stay. Homeowners who later return would also benefit from cheaper housing, but the initial loss in housing wealth incurred by moving is greater, and the welfare impact is greater than if they stay (−3.7% versus −2.3%, both in Column 6). Note that a share of the value of staying also reflects the value of leaving in the following period. When I combine the shocks, there is a minute welfare loss for renters when they move as well (−0.0056%) and a significant loss for homeowners (−3.7%).

Thus, even if renters are less impacted in welfare terms by the direct and indirect shocks to Stavanger, the change in the welfare differential of moving versus staying increases by more than it does for homeowners because the value of leaving falls by a significant amount for the latter group.

2.5.6.3 Changing moving costs and moving subsidies

I have briefly argued that a reduction in home prices acts as a rise in the cost of moving. I illustrate this by first re-solving the model with a shock to labor income but no accompanying adjustment of home prices; instead, either the disutility of moving τ or the monetary location-adjustment cost η^{ℓ} rises to make the leaving rate match the increase that a labor market shock with the house price shock produces. The numerical values are presented in Table 2.11.

The first column lists the original values of moving costs from the model estimation. Column 2 shows the necessary new levels of the moving disutility τ or the monetary moving cost η^{ℓ} to match the average change in the leaving rate across the whole Stavanger population. The increase in τ is approximately 250%, and for the monetary cost, the increase is approximately 280%, both quite significant.

Second, I similarly target the change in the average moving rate among renters and homeowners separately. This reveals that the disutility of moving τ or monetary moving cost η^{ℓ} has to increase by 360% or 200%, respectively, to make up for the lack of changes in rents. For homeowners, the figures are 230% and 300%, respectively. That is, renters act as if the disutility of moving rises by more than it does for homeowners, but the opposite is true for the monetary moving cost.

The differences across housing tenure illustrate that the changes in the moving costs are not indicative of whether the value of staying or leaving has changed. As previously shown, renters are incentivized to stay by cheaper housing, which raises the staying value, all else equal. Homeowners, in contrast, are worse off measured by welfare and benefit less from a potential move because it is associated with less utility, all else equal. The value of staying also falls but by less than for renters. However, moving costs reduce migration by exclusively lowering the value of leaving. Thus, how to interpret moving costs and shocks to them is an ambiguous exercise because they reflect only the change in the present-value differentials across locations. The exercise also shows the importance of the choice of including disutility versus monetary costs. Renters, who are more likely to be financially constrained, require a smaller increase in the monetary cost to produce a bigger utility loss and disincentivize migration.

Finally, I analyze the role of policy in this environment. Existing literature underscores the substantial costs associated with migration, which suggests that policies aimed at mitigating these costs could have a positive impact on overall welfare. While there are only a limited number of examples, some countries and regions have implemented such policies. For instance, in Germany, certain conditions allow unemployed individuals to receive financial assistance to facilitate relocation for job opportunities (Caliendo et al., 2017).[15] A similar program was in place in Sweden from 1959 to 1987 (Westerlund, 1998). Additionally, the U.S. state of Kentucky and Tulare County, California, offered assistance, primarily focused on welfare recipients and in practice facilitating moves within the region (Briggs and Kuhn, 2008).[16]

To study the effect of moving subsidies, I run an experiment where workers are offered financial support in the form of a one-time payment conditional on leaving Stavanger. Incomes and home prices are shocked as in the main experimental setting and the support is only offered once. For a smaller subsidy of the amount of 5% of an average income (approximately 16,000 NOK or 2800 USD), the leaving rate increases by 15% among renters and by 5.3% among homeowners. This corresponds to an increase in the welfare of leaving by 0.50% and 0.33%, respectively.[17] If the size of the subsidy matches the increase in the monetary moving cost renters experienced in Table 2.11, i.e., 0.44 shares of an average annual income (approximately 150,000 NOK or 24,000 USD), the relative increase in the leaving rate is 250% and 55% for renters and homeowners respectively following the income shock. That corresponds to a welfare raise when leaving of 4.1% and 2.9%, respectively.

The greater response among renters can be understood again by them being more financially constrained. The relative change in the moving rate among renters is three to five times greater than that of homeowners given the same moving subsidy. The simulated policy program does not favor specific types of workers, however, it does raise the welfare of renters who leave more than it does for homeowners. It is cost-effective to offer untargeted financial assistance and let workers decide what is optimal for themselves, but encouraging more out-migration of renters, who are generally younger and less likely to be attached to the area through relatives, can have unintended consequences moving forward. For example, this can

[15]While Caliendo et al. (2017) document positive effects on the job finding rate of program participants, Caliendo et al. (2023) highlight negative effects due to a reduction in the job search rate in the current location.

[16]There are also examples of experimental programs, such as the U.S. Department of Housing and Urban Development's Move to Opportunity experiment targeting people living in high-poverty neighborhoods, that primarily encouraged moves to other neighborhoods within the same region (see, e.g., Chetty et al., 2016).

[17]The increase in welfare refers to the increase in overall income that corresponds to the increased value of leaving and receiving a moving subsidy, i.e., equivalent variation.

further reduce the local housing demand and home prices, making homeowners less likely to leave. The welfare consequences for homeowners can therefore be worse in an environment of moving subsidies.

2.6 Conclusion

Many papers have studied the economics of migration, often through the lens of structural models or census data. This paper adopts an approach utilizing rich Norwegian panel data with annual observations to investigate how changing housing wealth and housing tenure impact the choice to stay or leave a location enduring a persistent adverse labor demand shock. I show how workers in the labor market region of Stavanger were impacted by the large fall in global oil prices in 2014 and follow their movements during the period that followed. The key finding is that renters, who do not have to realize a large loss in housing wealth, are more mobile and leave the region. The loss of housing wealth is a strong enough motive for homeowners to remain in the area, that is, to reduce their leaving rate.

The findings are qualitatively consistent with a life-cycle model with location, housing, and saving choices that highlight that, even if renters on net are partly compensated for the income shock through the accompanying reduction in rents and are better off than homeowners, they leave at a higher rate. This is because the reduction in home prices reduces homeowners' value of moving. The model also shows that moving subsidies are more effective at stimulating the migration of renters because they are more liquidity constrained. If used at a scale such that housing demand is further reduced, this can exacerbate the predicament of homeowners.

One takeaway from this work is the importance of considering how housing tenure is distributed across the economy to understand migration responses. If everyone rents, the reduction in housing prices would be worse, and landlords would bear the consequences. If the incidence of the shock affects groups with a higher rate of homeownership, then the change in home prices will be important for predicting what the moving response will be. Additional results also highlight the importance of the direction of migration for understanding the impact of home prices. From outside the region, lower home prices are attractive and bring in poorer workers, in terms of both income prospects and housing wealth, older workers, and workers who have family ties in the region. The young people and renters leave the area, reducing local housing demand. For them, the decline in home prices is not attractive enough to offset the loss in future labor earnings. Homeowners reduce their moving-out rate because the value of moving has fallen due to the shock to wealth the home price drop implies. This acts as a persistent moving cost shock.

My findings and the setting open up several questions for future research. I have not addressed the firm-side response to the change in local economic conditions, a significant factor influencing workers' labor market opportunities. Instead of focusing on the cost of giving up housing that has lost value, firms have different capital tied to the location in the form of customer bases, physical assets, immobile labor, etc. It is of interest both in itself and for general equilibrium consequences to better understand how able businesses are to relocate in response to local economic shocks.

I have also, as is common in the structural literature, abstracted from the process of job search within and across locations. However, this is at the center of the analysis of Munch et al. (2006) and Battu et al. (2008). The setting here, with complementary data, is well suited to offer further insights into the process. Regarding research on the home price equilibrium, a dimension that I have not fully exploited is the home transaction data available in Norway, which can be used to further study the segmentation of housing markets and the dynamics of the equilibria across the region of Stavanger, expanding on previous work of Määttänen and Terviö (2014) and Landvoigt et al. (2015). Another potential housing-related friction is the potential freeze of the housing market during the home price collapse. Exploring market illiquidity and how it can slow down labor relocation would expand on the findings of Garriga and Hedlund (2020).

2.A Appendix: Additional empirical results

Figure 2.A.1: Stavanger migration in levels

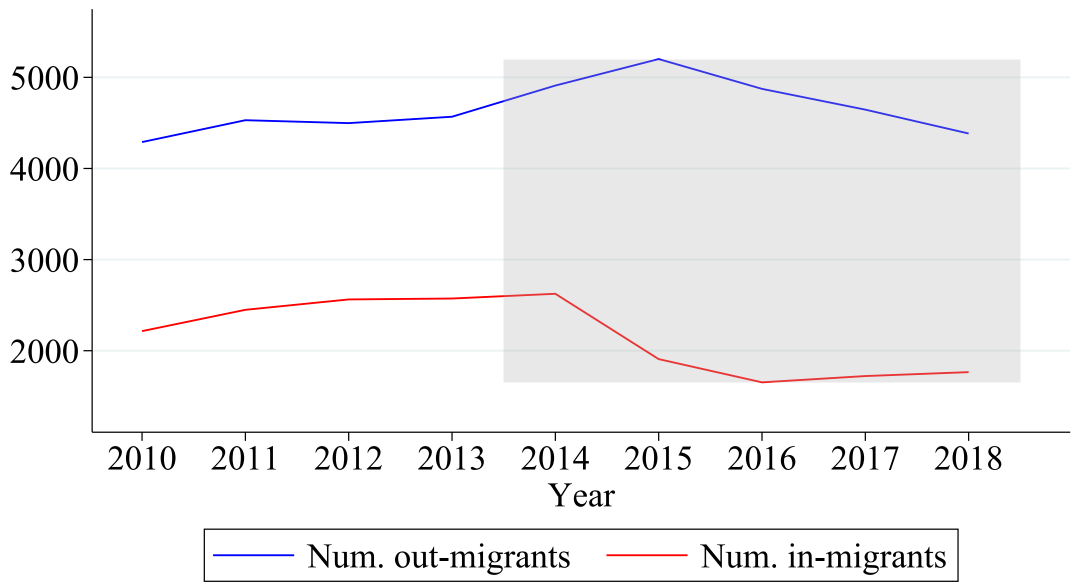

Note: This figure shows the levels of migration in and out of Stavanger in the analysis sample.

Figure 2.A.2: Changes in housing costs over time

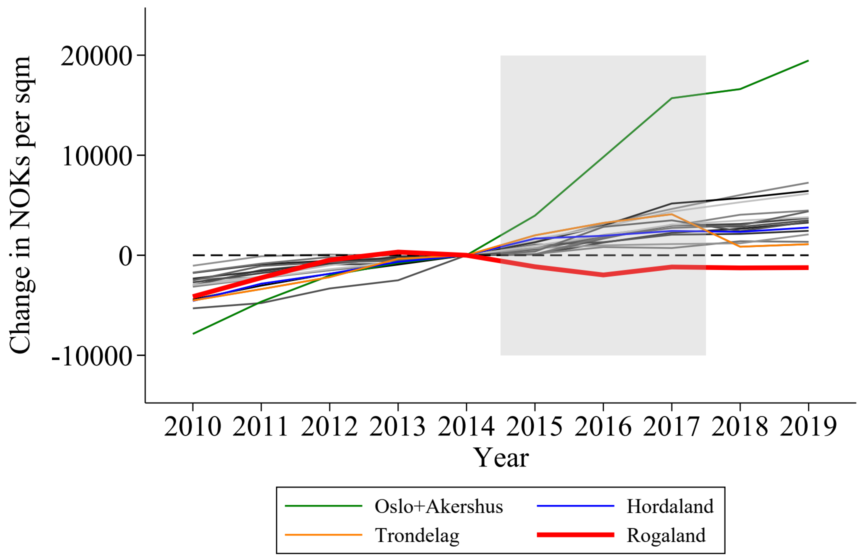

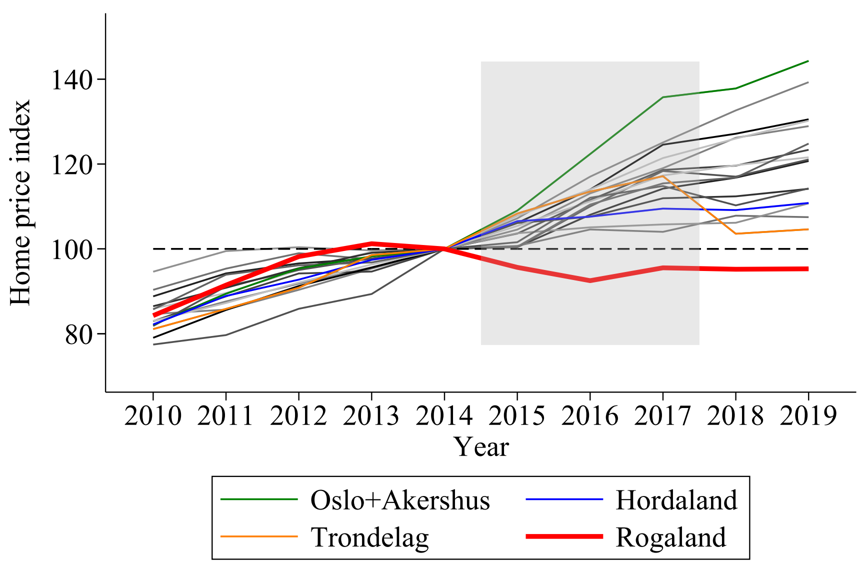

Note: The graph to the left shows the change in home prices in different counties over time, relative to 2014. The graph to the left shows the home price index using 2014 as the benchmark year.

Figure 2.A.3: Additional results on the impact on labor outcomes of workers

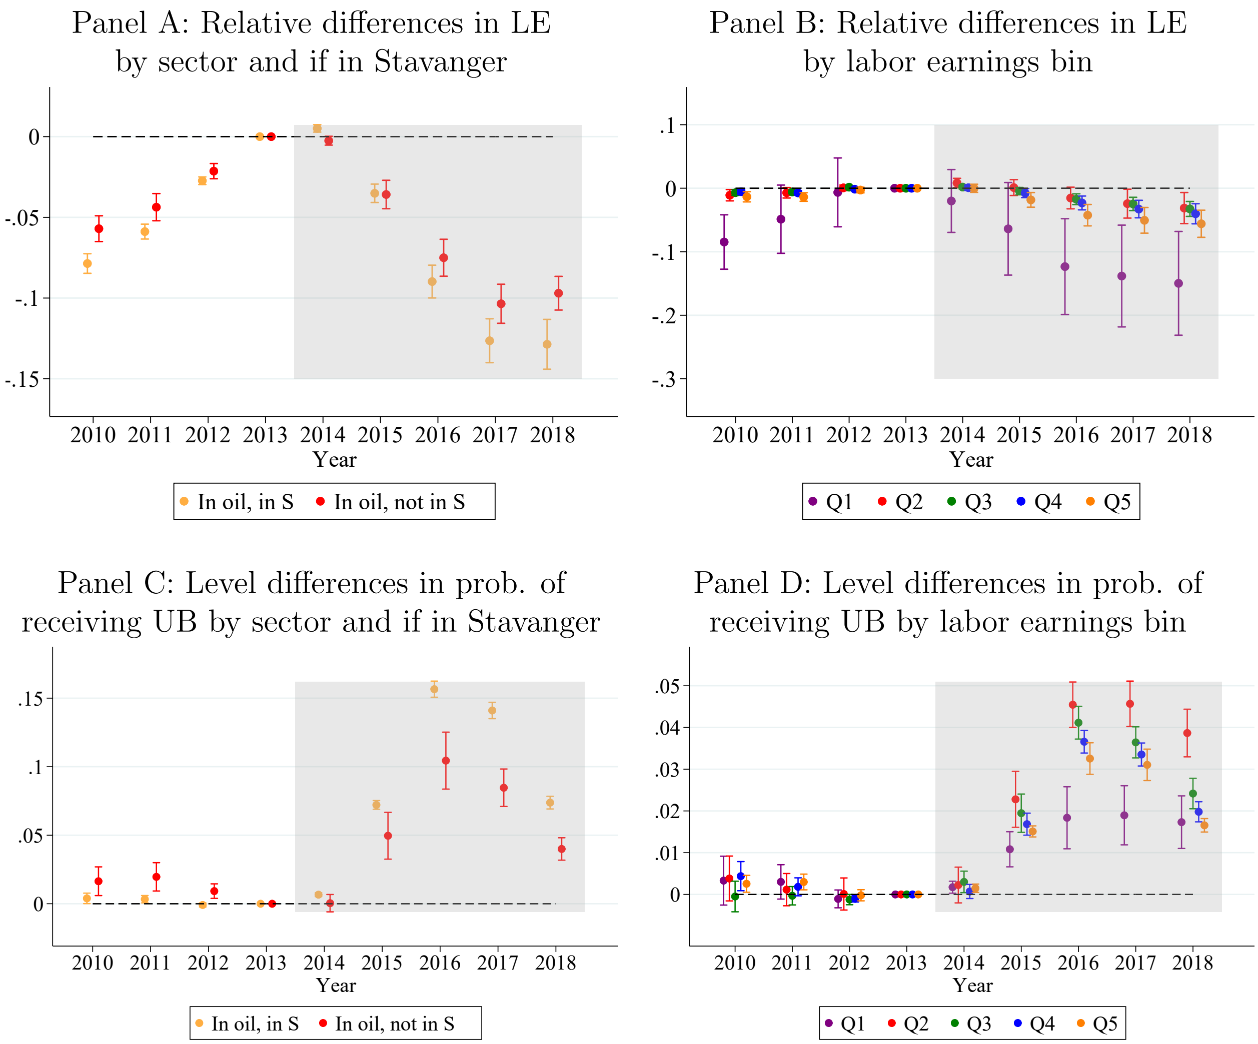

Note: This figure presents the annual differences in labor market outcomes of Stavanger and non-Stavanger workers in oil compared to workers in the rest of Norway. Panel A displays the log differences in labor earnings (LE) estimated using (2.1) and Panel B the level difference in the probability of unemployment benefits (UB) uptake estimated using (2.2). Labor earnings are the sum of wages, salaries, and income from self-employment. For more details, see Figure 2.2.

Figure 2.A.4: Aggregate outcomes for Rogaland county

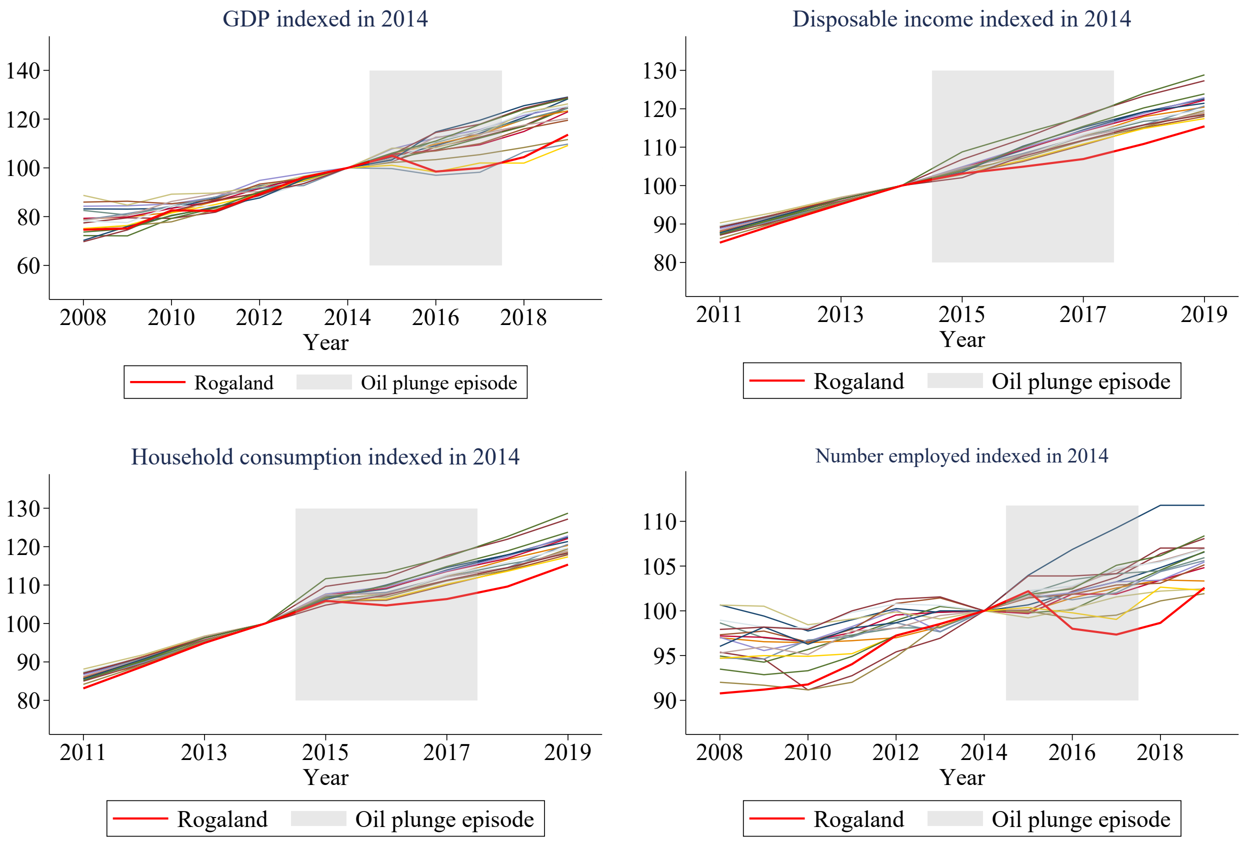

Note: This figure presents aggregate outcomes for the county Rogaland, which the Stavanger LMA is a big share of. Source: SSB.

Figure 2.A.5: Additional results for arrival probabilities

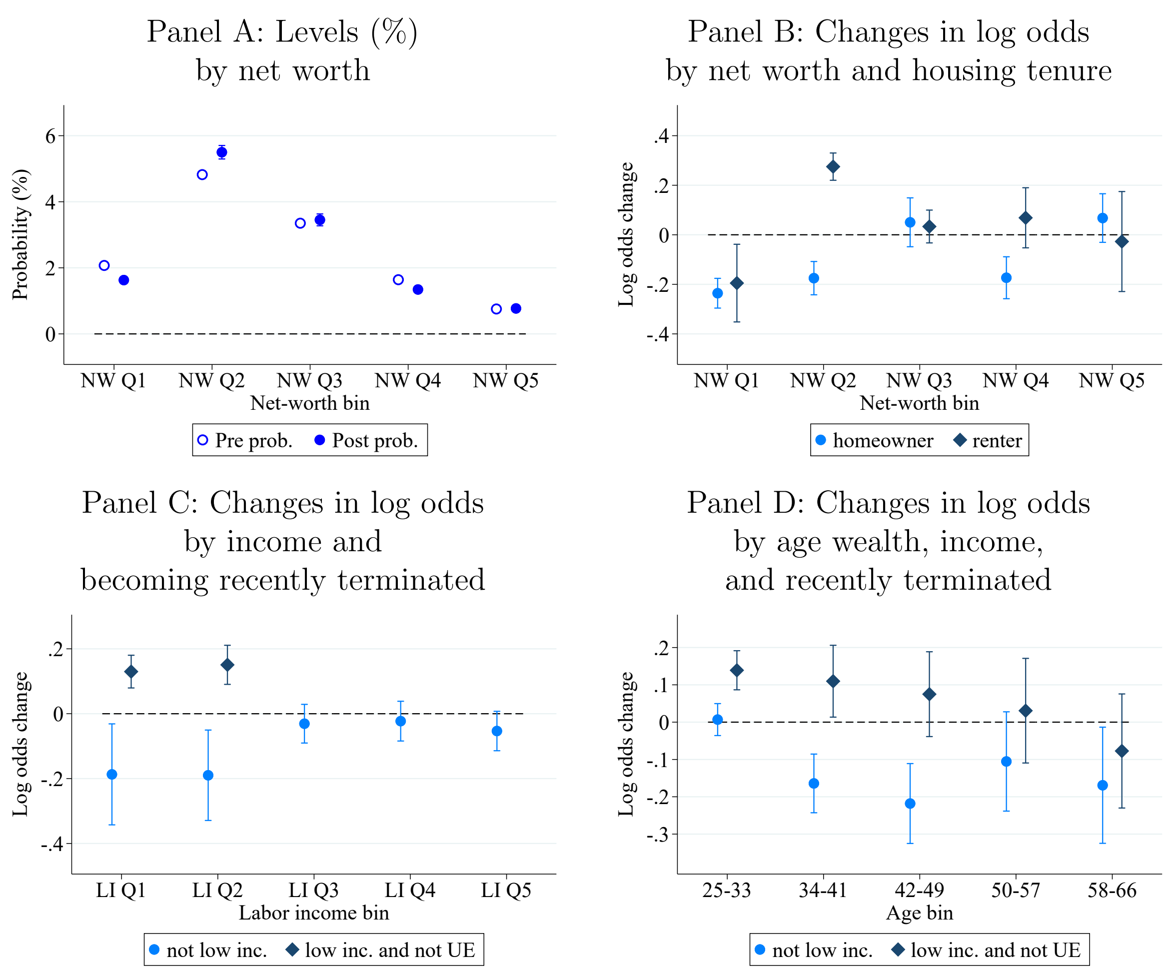

Note: This figure presents changes in the probability or log odds of in-migration to Stavanger following the oil price plunge of 2014. Panel A is produced by the OLS version of (2.13), Panels B–D are produced with logit, and Panels C and D use a third interaction term indicated by the corresponding legend. All error bands are 95% cluster-robust standard errors.

Table 2.A.1: Role of workers' idiosyncratic moving rates in the unemployed's higher mobility

	Intra-LMA move			Inter-LMA move		
	(1)	(2)	(3)	(4)	(5)	(6)
On UB_{it}	0.17***	0.016***	0.044***	0.45***	0.052***	0.18***
	(0.011)	(0.0043)	(0.0050)	(0.030)	(0.0083)	(0.017)
Intra-move rate_i			5.3***			
			(0.20)			
Inter-move rate_i						9.0***
						(0.21)
Worker FEs	No	Yes	No	No	Yes	No
Pseudo R-squared	0.023	0.111	0.071	0.033	0.338	0.168
Num. obs.	30,851,125	30,851,125	30,851,125	30,851,125	30,851,125	30,851,125

Note: This table presents the differences in leaving rates between employed and unemployed workers using a Poisson model. Being unemployed is defined by receiving unemployment benefits (UB). Intra- and inter-move freq. is the individual annual moving rate in years of no uptake of UB. All the estimated models include a fourth-degree polynomial in age and LMA and year fixed effects. The sample covers all Norwegians during 1992–2018 and workers observed less than 3 times are dropped. The standard errors are two-way cluster robust at the LMA and year level.

The higher mobility of workers who experience unemployment: A common finding in the literature on migration and labor market shocks is the elevated mobility of unemployed workers. This is found by regressing a dummy indicating a move on a set of worker characteristics and a dummy indicating whether the worker is unemployed or not.

In Table 2.A.1 I show that this is also the case in Norway, by regressing a dummy indicating either an intra-LMA move or an inter-LMA move on an indicator on UB, indicating whether the worker is receiving unemployment benefits in the year before. I use a lagged variable to not risk picking up the influence of higher unemployment risk following a move. The model specification is a Poisson regression and I include LMA and year fixed effects as well as control for age effects using a fourth-degree polynomial. I correct for two-way clustering in LMA and year. The sample is the full Norwegian population in 1992–2018 except that I remove individuals observed less than three times. The panel is not balanced.

Table 2.A.2: Uptake of social welfare among Stavanger arrivals

	OLS			Poisson		
	(1)	(2)	(3)	(4)	(5)	(6)
post_t	0.066***	0.061***	0.055***	0.26***	0.11***	0.087**
	(0.010)	(0.010)	(0.018)	(0.031)	(0.037)	(0.036)
Additional controls	No	Yes	Yes	No	Yes	Yes
Low inc. sample	No	No	Yes	No	No	Yes
(Pseudo) R-squared	0.007	0.620	0.568	0.004	0.328	0.165
Num. obs.	8083	8083	4133	8083	8083	4133

Note: This table presents the change in the one-year-lagged share of government transfers of post-tax income of Stavanger arrivals following the shock, estimated using OLS and Poisson. Additional controls refer to including origin LMA FEs and two fourth-degree polynomials in age and labor earnings, both standardized and winsorized. Low inc. refers to lagged labor income being below the 40th percentile (i.e., in the two lower quintile bins). The standard errors are two-way cluster robust at the origin LMA and year level.

In both migration cases, there is a strong correlation between unemployment and moving; the values in columns 1 and 4 are log points. However, an overlooked possibility is that the people who experience unemployment differ in their overall migration probability. I test for this by including individual fixed effects. The results are presented in columns 2 and 5. The difference in migration in years of unemployment is greatly reduced, by approximately a factor of nine to ten. The R-squared increases due to the inclusion of individual fixed effects. An alternative approach to make this point is to compute the average moving rate within each individual in years of no unemployment, and include that instead of individual fixed effects in the regression. The results of doing this are presented in columns 3 and 6. The reduction is not as stark as when including individual fixed effects, but the influence of UB uptake in one year is again much reduced.

This exercise shows that the elevated migration rate among the unemployed is to a large extent explained by an overall higher tendency among them to move. However, estimating individual fixed effects or workers' general moving probability requires long panels not often available.

2.B Appendix: Details on estimating worker skills and LMA-specific wage premia

This section describes the estimation process for the individual worker skills and the different wage premia earned in different LMAs.

From the panel used for the reduced-form evidence, I select individuals observed for at least five years. I then create a variable income_{it} that is the sum of labor and business income and government transfers (including unemployment benefits). I then estimate the overall skill of each worker s_i, every labor market area's premium LP_ℓ, and each location's skill premium SP_ℓ as in

$$\log \text{income}_{it} = \alpha_0 + \sum_{\ell \in \mathcal{L}} \text{LP}_\ell \times \mathbb{1}(\ell_{it} = \ell) + \sum_{\ell \in \mathcal{L}} \text{SP}_\ell \times \mathbb{1}(\ell_{it} = \ell) \times s_i + \eta X_{it} + \varepsilon_{it}.$$

2.B.1 Algorithm

The skill premia SP_ℓ and skills s_i have to be jointly estimated, and they enter the income equation as a product. The lion's share of the number of parameters to estimate comes from the individual skills, and maximum likelihood seems infeasible. Instead, I proceed in a fixed-point fashion running OLS estimations (e.g., using `reghdfe`, Correia, 2016). The method is similar to the algorithm used in De la Roca and Puga (2017) to estimate the benefits of learning in larger cities, a procedure that also requires estimating unobserved ability across individuals, which interacts with an unobserved learning effect. A conceptual difference is that my implementation does not assume that the premium is a particular function of city size but is specific to the location. This allows, in theory, the skill premium to be independent of the overall wage bonus from working in a location.

Step 0: To obtain an initial guess for LP_ℓ and s_i, I first estimate

$$y_{it} = \alpha_0 + \eta_i + \eta_\ell + \eta\, X_{it} + \varepsilon_{it}$$

and set $\hat{s}_i^1 = \hat{\eta}_i$ and $\widehat{\mathrm{LP}}_\ell^1 = \hat{\eta}_\ell$ (i.e., I take the estimated individual and labor market area fixed effects). Here, $y_{it} = \log \text{income}$.

Step 1, iteration $j \geq 1$: Estimate the skill premium SP_ℓ for each location using OLS:

$$\bar{y}_{it}^{\text{less LP}} \equiv y_{it} - \widehat{\mathrm{LP}}_\ell^j = \sum_{\ell \in \mathcal{L}} \mathrm{SP}_\ell \times \mathbb{1}(\ell_{it} = \ell) \times \hat{s}_i^j + \eta\, X_{it} + \varepsilon_{it}.$$

I denote the estimates by $\widehat{\mathrm{SP}}_\ell^j$.

Step 2, iteration j: Update the guess of the individual skills by inverting the expression above:

$$\hat{s}_{it}^{j+1} = \frac{y_{it} - \widehat{\mathrm{LP}}_\ell^j - \hat{\eta}\, X_{it}}{\widehat{\mathrm{SP}}_\ell^j}.$$

To obtain the constant individual skill, take the average: $\hat{s}_i^{j+1} = \frac{1}{N_i} \sum \hat{s}_{it}^{j+1}$.

Step 3, iteration j: Update the guess of the labor market area premium by running OLS on

$$\bar{y}_{it}^{\text{less SP}} \equiv y_{it} - \widehat{\mathrm{SP}}_\ell^j \times \hat{s}_{it}^{j+1} = \alpha_0 + \sum_{\ell \in \mathcal{L}} \mathrm{LP}_\ell \times \mathbb{1}(\ell_{it} = \ell) + \eta\, X_{it} + \varepsilon_{it}.$$

Denote the estimate $\widehat{\mathrm{LP}}_\ell^{j+1}$.

Step 4, iteration j. Compute the norm of the relative changes in all the estimated parameters:

$$\text{error}^{j+1} = \left(\sum_i \left(\frac{\hat{s}_i^{j+1} - \hat{s}_i^j}{\hat{s}_i^j} \right)^2 + \sum_{\ell \in \mathcal{L}} \left(\frac{\widehat{\text{LP}}_\ell^{j+1} - \widehat{\text{LP}}_\ell^j}{\widehat{\text{LP}}_\ell^j} \right)^2 + \left(\frac{\widehat{\text{SP}}_\ell^{j+1} - \widehat{\text{SP}}_\ell^j}{\widehat{\text{SP}}_\ell^j} \right)^2 \right)^{1/2}$$

and check if it satisfies the convergence criterion.

If the critical level has not been reached, return to step 1, and increment the iteration counter j by one. We have now estimated LP, SP, and skill s, and use these in the next iteration.

If the criterion is satisfied, use the last estimated LP, SP, and skill s.

Note that all fixed effects are estimated as deviations from the mean, given by the intercept. Otherwise, the estimation suffers from collinearity. By pinning down the mean of the fixed effects to zero, I can identify the fixed effects.

The skill premium, intercept, and individual skills are identified up to an affine transformation. This does not affect the predictions of income differentials across locations or, very importantly, the location value but makes it possible to standardize individual skills and adjust the skill premium and intercept accordingly. The formulas are derived below.

$$\begin{aligned} \hat{\alpha}_0 + \sum_{\ell \in \mathcal{L}} \widehat{\text{SP}}_\ell \times \mathbb{1}(\ell_{it} = \ell) \times \hat{s}_i = \tilde{\alpha}_0 + \sum_{\ell \in \mathcal{L}} \widetilde{\text{SP}}_\ell \times \mathbb{1}(\ell_{it} = \ell) \times \frac{\hat{s}_i - \bar{s}}{\text{sd}(s)} \\ \Rightarrow \quad \widetilde{\text{SP}}_\ell = \widehat{\text{SP}}_\ell \times \text{sd}(s), \quad \forall \ell \in \mathcal{L}, \\ \Rightarrow \quad \tilde{\alpha}_0 = \hat{\alpha}_0 + \sum_{\ell \in \mathcal{L}} \widetilde{\text{SP}}_\ell \times \mathbb{1}(\ell_{it} = \ell) \times \frac{\bar{s}}{\text{sd}(s)}. \end{aligned}$$

2.C Properties of Gumbel distributed random variables

A random variable X that follows a Gumbel distribution Gumbel(μ, β) has PDF, CDF, and expectation value

$$\begin{aligned} f_X(x) &= e^{-\frac{x-\mu}{\beta}+\exp\left(-\frac{x-\mu}{\beta}\right)}, \\ F_X(x) &= e^{-\exp\left(-\frac{x-\mu}{\beta}\right)}, \\ \mathbb{E}[X] &= \mu + \beta\gamma, \end{aligned}$$

where γ is the Euler–Mascheroni constant and is approximately 0.5772.

If g_i is Gumbel$(0, 1)$ and x_i is a sequence of deterministic real numbers, then,

$$\begin{aligned} \mathbb{P}[j = \arg\max_i x_i + \nu \times g_i] &= \frac{e^{x_j/\nu}}{\sum_i e^{x_i/\nu}}, \text{ and} \\ \mathbb{E}[\max_i x_i + \nu \times g_i] &= \nu \left(\log \sum_i e^{x_i/\nu} + \gamma \right). \end{aligned}$$

To avoid floating-point errors (due to taking the exponent of a number of an excessively great magnitude), we can subtract or add an arbitrary $\bar{x}$ to each x_i. The probability expression is unbiased by the transformation, but the expectation is biased and requires a correction term. Thus,

$$\begin{aligned} \mathbb{P}[j = \arg\max_i x_i + \nu \times g_i] &= \frac{e^{(x_j-\bar{x})/\nu}}{\sum_i e^{(x_i-\bar{x})/\nu}}, \text{ and} \\ \mathbb{E}[\max_i x_i + \nu \times g_i] &= \nu \left(\log \sum_i e^{(x_i-\bar{x})/\nu} + \gamma \right) + \bar{x}. \end{aligned}$$

References

Aastveit, Knut Are, Jesper Bojeryd, Magnus Gulbrandsen, Ragnar Juelsrud, and Kasper Roszbach (2024) "What do 12 billion card transactions say about house prices and consumption?" Working Paper, Norges Bank.

Aastveit, Knut Are, R Juelsrud, and Ella Getz Wold (2022) "The leverage-liquidity trade-off of mortgage regulation," Working Paper 6/22, Norges Bank.

Autor, David, David Dorn, and Gordon H Hanson (2021) "On the persistence of the China shock,"Technical report, National Bureau of Economic Research.

Battu, Harminder, Ada Ma, and Euan Phimister (2008) "Housing tenure, job mobility and unemployment in the UK," *The Economic Journal*, 118 (527), 311–328.

Baumeister, Christiane and Lutz Kilian (2016) "Understanding the decline in the price of oil since June 2014," *Journal of the Association of Environmental and resource economists*, 3 (1), 131–158.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra (2018) "House prices and consumer spending," *The Review of Economic Studies*, 85 (3), 1502–1542.

Bhuller, M (2009) "Classification of Norwegian labor market regions," *SSB Notater*, 24.

Bilal, Adrien and Esteban Rossi-Hansberg (2021) "Location as an Asset," *Econometrica*, 89 (5), 2459–2495.

Bjørnland, Hilde C, Frode Martin Nordvik, and Maximilian Rohrer (2021) "Supply flexibility in the shale patch: Evidence from North Dakota," *Journal of Applied Econometrics*, 36 (3), 273–292.

Blanchflower, David G and Andrew J Oswald (2013) "Does high home-ownership impair the labor market?"Technical report, National Bureau of Economic Research.

Briggs, Brianna and Peter Kuhn (2008) "Paying for the relocation of welfare recipients: Evidence from the Kentucky Relocation Assistance Program," Discussion Paper, University of Kentucky Center for Poverty Research.

Caliendo, Marco, Steffen Künn, and Robert Mahlstedt (2017) "The return to labor market mobility: An evaluation of relocation assistance for the unemployed," *Journal of Public Economics*, 148, 136–151.

——— (2023) "The intended and unintended effects of promoting labor market mobility," *Review of Economics and Statistics*, 1–52.

CER (2023) "2023 estimated production of Canadian crude oil and equivalent," report, Canada Energy Regulator.

Chetty, Raj, Nathaniel Hendren, and Lawrence F Katz (2016) "The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment," *American Economic Review*, 106 (4), 855–902.

Clerici, A and G Alimonti (2015) "Oil shale, shale oil, shale gas and non-conventional hydrocarbons," in *Epj web of conferences*, 98, 03001, EDP Sciences.

Correia, Sergio (2016) "Linear models with high-dimensional fixed effects: An efficient and feasible estimator," Working Paper, http://scorreia.com/research/hdfe.pdf, Board of Governors of the Federal Reserve System.

Correia, Sergio, Paulo Guimaraes, and Thomas Zylkin (2019) "PPMLHDFE: Stata module for Poisson pseudo-likelihood regression with multiple levels of fixed effects," Working Paper, Board of Governors of the Federal Reserve System.

De Nardi, Mariacristina (2004) "Wealth inequality and intergenerational links," *The Review of Economic Studies*, 71 (3), 743–768.

Druedahl, Jeppe (2021) "A guide on solving non-convex consumption-saving models," *Computational Economics*, 58 (3), 747–775.

EIA (2023) "Tight oil production estimates by play," table, U.S. Energy Information Administration, accessed on 10/31/2023.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri (2020) "Heterogeneity and persistence in returns to wealth," *Econometrica*, 88 (1), 115–170.

Fagereng, Andreas, Magnus AH Gulbrandsen, and Gisle J Natvik (2022) "Household responses to export prices – Evidence from an oil exporting country," Working Paper, Available online.

Fonseca, Julia and Lu Liu (2023) "Mortgage lock-in, mobility, and labor reallocation," Working Paper, Mobility, and Labor Reallocation (March 24, 2023).

Garriga, Carlos and Aaron Hedlund (2020) "Mortgage debt, consumption, and illiquid housing markets in the Great Recession," *American Economic Review*, 110 (6), 1603–1634.

Giannone, Elisa, Qi Li, Nuno Paixao, and Xinle Pang (2023) "Unpacking moving," Working Paper.

Glaeser, Edward L and Joseph Gyourko (2005) "Urban decline and durable housing," *Journal of political economy*, 113 (2), 345–375.

Hetland, Kjersti and Mathias Oppedal (2016) "Forstå oljekrisa på tre minutt," *NRK*, https://www.nrk.no/rogaland/forsta-oljekrisa-pa-tre-minutt-1.13112473.

International Monetary Fund (2023) "Primary commodity price system," table.

Jabobsen, Siw Ellen and Ida Kvittingen (2016) "Burde sett krisen komme," *Forskning.no, 2016-05-24*, https://forskning.no/okonomi-naeringsliv-olje-og-gass/burde-sett-krisen-komme/418314, Accessed 2023-08-11.

Juelsrud, Ragnar E and Ella Getz Wold (2019) "The saving and employment effects of higher job loss risk," Working Paper 17/2019, Norges Bank.

Kaplan, Greg, Kurt Mitman, and Giovanni L Violante (2020) "The housing boom and bust: Model meets evidence," *Journal of Political Economy*, 128 (9), 3285–3345.

Kennan, John and James R Walker (2011) "The effect of expected income on individual migration decisions," *Econometrica*, 79 (1), 211–251.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider (2015) "The housing market(s) of San Diego," *American Economic Review*, 105 (4), 1371–1407.

Lorentzen, Linnea (2023) "Domino effects: Understanding sectoral reallocation and its wage implications," Working Paper, University of Oslo.

Määttänen, Niku and Marko Terviö (2014) "Income distribution and housing prices: An assignment model approach," *Journal of Economic Theory*, 151, 381–410.

McFadden, Daniel (1974) "Conditional logit analysis of qualitative choice behavior," *Frontiers in Econometrics*.

Modestino, Alicia Sasser and Julia Dennett (2013) "Are American homeowners locked into their houses? The impact of housing market conditions on state-to-state migration," *Regional Science and Urban Economics*, 43 (2), 322–337.

Monras, Joan (2018) "Economic shocks and internal migration," Working Paper, Federal Reserve Bank of San Francisco.

Munch, Jakob Roland, Michael Rosholm, and Michael Svarer (2006) "Are homeowners really more unemployed?" *The Economic Journal*, 116 (514), 991–1013.

Norges Bank (2014) "Monetary policy report with financial stability assessment 2014/4," policy report.

Notowidigdo, Matthew J (2020) "The incidence of local labor demand shocks," *Journal of Labor Economics*, 38 (3), 687–725.

OECD (2016) "OECD Factbook 2015–2016: Economic, environmental and social statistics," Report, https://doi.org/10.1787/factbook-2015-en.

De la Roca, Jorge and Diego Puga (2017) "Learning by working in big cities," *The Review of Economic Studies*, 84 (1), 106–142.

Rodríguez-Clare, Andrés, Mauricio Ulate, and Jose P Vasquez (2020) "Trade with nominal rigidities: Understanding the unemployment and welfare effects of the China shock," Working Paper 27905, National Bureau of Economic Research.

Rustad Thorsen, Lotte (2014) "Økt boliggjeld, men også økte inntekter."

SSB (2015) "Economic Survey 2/2015," survey, Statistisk sentralbyrå.

Stocker, Marc, John Baffes, Y Modeste Some, Dana Vorisek, and Collette Mari Wheeler (2018) "The 2014–16 oil price collapse in retrospect: Sources and implications," *World Bank Policy Research Working Paper* (8419).

Strand, Pål (2014) "Forbruksundersøkelsen 2012," table 10235, Statistisk sentralbyrå.

Valletta, Robert G (2013) "House lock and structural unemployment," *Labour Economics*, 25, 86–97.

Westerlund, Olle (1998) "Internal migration in Sweden: The effects of mobility grants and regional labour market conditions," *Labour*, 12 (2), 363–388.

Zabek, Mike (2024) "Local ties in spatial equilibrium," *American Economic Journal: Macroeconomics*, 16 (2), 287–317.

CHAPTER 3

What do 12 billion card transactions say about house prices and consumption?

with Knut Are Aastveit, Magnus Gulbrandsen, Ragnar Juelsrud, and Kasper Roszbach[1]

How does household consumption respond to home price movements? The question concerns both the overall magnitude of total spending as well as the distribution across different consumer items. Its answers are central to understanding the role of the housing market as a driver or amplifier of business cycle fluctuations and important for how wealth and collateral affect household consumption decisions.

The predominant theoretical perspective, rooted in the permanent income hypothesis, suggests that home price effects on consumption should be minimal (e.g., Campbell and Cocco, 2007; Sinai and Souleles, 2005). However, an expanding body of empirical research, including studies such as Mian et al. (2013) and Aladangady (2017), contradicts this view, indicating substantial impacts. In response to recent empirical findings, Berger et al. (2018) argue that models of consumption with incomplete markets can predict substantial consumption responses. They propose a simple rule-of-thumb formula for the marginal propensity to consume (MPC) out of housing wealth: the marginal propensity to consume out of temporary income multiplied by the value of housing. Since MPCs out of temporal income vary greatly across household characteristics, the formula implies that the effects of home prices on consumption vary considerably among households (Fagereng et al., 2021; Kaplan and Violante, 2014).

[1]This chapter should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. We would like to thank Yann Cerasi for his excellent research assistance.

In this paper, we utilize a quasi-experimental setting in Norway with highly detailed data on household-level spending and characteristics to answer the question of how much spending changes with housing wealth, what determines the response, and what types of spending are affected.

Our approach exploits the regional variations in home prices that resulted from the 2014–2015 oil price shock and differences in the regional exposure to the oil sector. We focus on government workers and contrast between those who live in the Greater Stavanger region—where the Norwegian petroleum industry is concentrated—and those living in the rest of Norway. The earnings of government workers are stable and centrally negotiated at the national levels, providing us with a treatment and control group with essentially the same expectations on unemployment and income risk but different exposures to the changes in local home prices. The setting is discussed in more detail in Chapter 2, Juelsrud and Wold (2019), and Lorentzen (2023).

First, we document a per-year marginal propensity to consume out of housing wealth of approximately 2 percent over a twenty-four-month period, a number somewhat lower than in most prior studies.

Second, our detailed microdata lets us unravel how home values influence spending. We discover that the marginal propensity to consume out of housing wealth increases with leverage and decreases with liquidity and age. As home values decline, borrowing constraints tighten for households near the collateral limit, hindering credit-constrained households from borrowing against their homes to boost consumption. When controlling for age, wealth, and liquidity, we observe that households with a loan-to-value ratio close to a legal cap exhibit a marginal propensity to consume out of housing wealth of roughly 4 percent, double that of the average household.

Third, we show that the adjustment of spending varies across consumption categories. General equilibrium effects hinge on which spending categories households alter. Our findings indicate that durable goods such as cars and furnishings exhibit the strongest negative response to falling home prices, while semi-durables like clothing are less affected. Essential goods such as food and beverages show an opposite response, suggesting a shift away from luxury consumption.

Our study integrates information from three distinct sources to offer a comprehensive understanding of household behavior. First, we utilize Norwegian tax data, providing detailed third-party-reported information on household balance sheets, including their income, wealth, leverage, liquidity, debt, and labor market information. Second, for assessing consumption, we use electronic payment data from NETS Branch Norway, available at the weekly frequency and covering 26 different consumption categories through debit card transactions. Notably, debit cards constitute around 80% of all card transactions in Norway. Our consumption data does not only include debit card transactions but also incorporates bank wire transfers (invoice payments) from households to firms, such as monthly credit card bill payments. All data is aggregated at the household level. The third source of data we use is regional home price data from *Eiendom Norge*.

The main identifying assumption underlying our analysis is that the oil shock only affects government workers' consumption through its impact on local home prices. However, there are potential challenges to this assumption. First is one regarding parallel trends: do government workers have similar consumption patterns across different regions? If not, that indicates structural differences between regions. We address this by examining consumption trends before the oil price collapse and find that consumption among government workers in the Greater Stavanger behaves similarly to those in other regions before the collapse.

A second potential issue is the exclusion restriction. There might be other factors affecting government workers differently due to the oil price collapse. For example, income expectations could worsen for government workers in the oil region compared to those elsewhere. We address this by analyzing two alternative samples: homeowning retirees and renters. Homeowning retirees are not influenced by changes in local labor markets but may be affected by shocks in home prices, while renters are influenced by local labor market conditions but not by changes in home prices. Our results offer reassurance as they consistently demonstrate disparities similar to the main results when comparing home-owning retirees or renting government workers in the oil region to their counterparts elsewhere.

Our findings carry significant policy implications. Although there is a prevalent policy concern that fluctuations in home prices may lead to substantial contractions in consumer spending, our results indicate a limited overall impact of home prices on consumption, with minimal negative spillover effects. Additionally, our findings suggest that reducing household indebtedness is unlikely to result in a substantial effect on the aggregate MPC. This is because the heterogeneous response across household characteristics is relatively small. Specifically, only households near the LTV cap or those with very low liquidity exhibit a more aggressive response. In our setting, this constitutes a relatively small group of households, and their impact on the overall response is likely to be limited. Thus, the aggregate response is a function of the distribution of indebtedness in the economy (Eggertsson and Krugman, 2012).

The remainder of this chapter is organized as follows. In Section 3.1, we describe the data and variables. Section 3.2 outlines our research design, the empirical methodology, and the identification challenges. Section 3.3 presents the causal effects of home price changes on household consumption and discusses possible mechanisms behind our results. Section 3.4 concludes.

3.1 Data and sample construction

This section describes our data sources, sample construction, variable definitions, and summary statistics.

3.1.1 Data sources

We base our analysis on three datasets. The first dataset captures consumption at the individual-week level. It is provided by NETS, an international provider of payment and information services in Scandinavia, the Baltic states, and Switzerland. It contains information on debit card payments through BankAxept for the universe of Norwegian residents. Over our sample period, BankAxept accounted for 71% of the volume of total card transactions and is an open system any bank business in Norway can join and to which nearly all firms are connected. In addition, the data also contains invoice payments and direct remittances. The data is at the individual $\times$ week $\times$ zip-code level for 26 different consumption categories based on the Classification of Individual Consumption According to Purpose (COICOP).

The second dataset contains background data on all Norwegian individuals at an annual frequency and is provided by Statistics Norway. In practice, the data is built from tax returns containing information on income and balance sheet variables for all Norwegians. Since Norway levies both income and wealth taxes, income and wealth data from the tax authorities provides a complete and comprehensive overview of the financial position of Norwegian households. It contains granular information about income sources, including information about wage income, business income, capital income, and welfare transfers, in addition to a breakdown of household balance sheets including bank deposits, primary and secondary housing, and debt. It also includes *household* IDs, which we use for aggregation as described in Section 3.1.1.

In addition to the two large databases on consumption and background information, we use regional information about home prices from *Eiendom Norge* ("Real Estate Norway"). These data provide house indices for a large number of Norwegian regions. The home price indices are used—in combination with information about the value of an individual's house—to compute housing wealth changes at the individual level. We describe this procedure in more detail in Section 3.1.3. The evolution of the home price indices is shown in Figure 3.1, where the red lines indicate the home price indices relevant for the Greater Stavanger region.

Figure 3.1: The evolution of home price indices across Norway

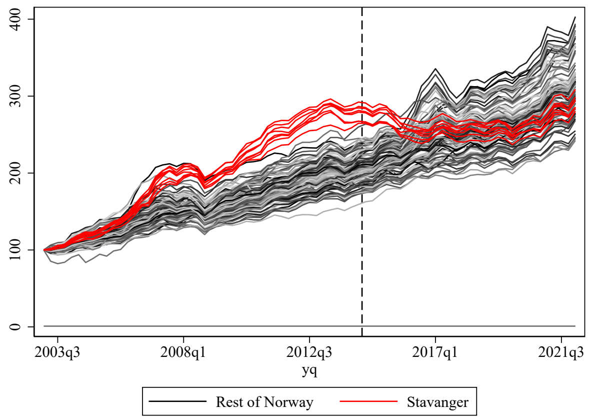

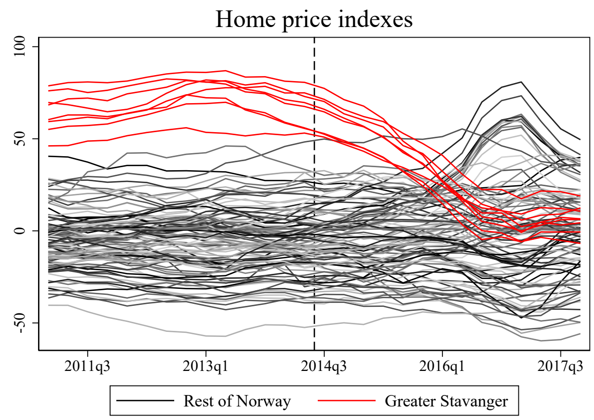

Note: The left graph shows home price indices for different housing markets across Norway, with Greater Stavanger markets in red and the rest in gray. The right graph shows the same indices, detrended by the average year and quarter index. Source: *Eiendom Norge*.

3.1.2 Sample construction

For our analysis, we focus on households in which the majority of the adult household members work in the government sector in 2013 and 2014.[2] Our initial sample consists of 281,000 government worker households, and we aggregate consumption by household, coarser consumption categories, and year-quarter. To observe a housing wealth effect, we condition inclusion on whether households are homeowners at the beginning of 2012. We also require that households have self-employed income of less than 5% of total income to avoid including government workers who also depend on non-government earnings. Our final sample consists of about 206,000 households.

For the cross-sectional analysis, we compare outcomes in 2012 with post-oil plunge outcomes. We then have 1,844,835 period-individual pairs for which we observe all card transactions.

[2]We define a worker as in government if their "institutional sector" (*institusjonell sektorgruppering*) is 1110 or 1120 (non-financial enterprises owned by central government), 1510 or 1520 (non-financial enterprises owned by local government), 3100 or 3900 (public financial corporations), 6100 (central government), or 6500 (local government) when using the version from 2012. The codes we use for older observations use the 1987 version and are 110 (central government) 510 (county municipalities), 550 (municipalities), 150, or 190 (public financial corporations), 610, 630, 635, 660, or 680 (public, non-financial enterprises). We exclude workers in industries (Standard Industrial Classification 2007) that start with 06 (extraction of crude petroleum and natural gas) and 09.1 (support activities for petroleum and natural gas extraction).

3.1.3 Variable definitions and summary statistics

In this subsection, we define the key variables used in our analysis and present summary statistics. All monetary values are nominal values in Norwegian krone (NOK).[3]

3.1.3.1 Demographic information

For each household, we compute the average age of all tax-filing members and the number of individuals within the household who work in the government sector as defined above. We also observe the municipality, number of children, and several other demographic characteristics.

3.1.3.2 Consumption

In the main analysis, we focus on four different notions of consumption: total consumption (the sum of all consumption categories less public sector payments[4]) and three major subcategories of consumption, i.e. food and beverages, furnishing, and vehicles. All are nominal amounts and we cannot observe individual items purchased. These consumption categories are salient examples of non-durable, semi-durable, and durable consumption.

Compared to consumer surveys conducted by Statistics Norway, our measure of total expenditures are within 1.8–15% of those reported. Furnishing is similar (1.8–10%). However, our measure of food and beverages is about 33–46% greater and vehicles 27–36% less (see Table 3.A.1). The latter is likely due to survey questions include used cars which our measure does not capture, if the car is transacted between private individuals.

[3] The USD/NOK exchange rate is 10.05 as of 31 January 2023.

[4] To ensure anonymity of the subjects, consumption categories have been aggregated before they are delivered to us. Payments to the public sector are aggregated with other smaller expenditures into "Consumption, miscellaneous" which we do not use in our analysis.

3.1.3.3 Household income and balance sheet variables

We measure the home value of a household by the implied market value of the household's primary house at a given point in time, based on the estimated market value reported by the tax authorities in 2012 and the evolution of the local home price indices. The household's loan-to-value (LTV) is calculated by dividing total household debt by the sum of the estimated market values of primary and secondary housing. We measure debt-to-income (DTI) by dividing total household debt by after-tax income. Finally, we measure liquid-wealth-to-income (LTI) by dividing bank deposits by after-tax income.

3.1.3.4 Summary statistics

Table 3.1 presents descriptive statistics of our sample. Average total consumption in 2012Q4 constituted approximately 99,000 NOK, but the distribution is skewed with a median of about 86,000 NOK. Out of the three subcategories of consumption, food, and beverages are the most important with an average consumption of approximately 23,000 NOK, while average furnishing and vehicle consumption is approximately 6,000 and 4,000 NOK.

The average primary home value is 2.7 million NOK, with a fairly large standard deviation of 1.3 million NOK. The LTV distribution is very skewed, with a mean of 0.90, a median of 0.48, and a standard deviation of 168.

The average household age is 47, and households on average have 1.42 members working in the government sector. Approximately 5% of our sample resides in the Greater Stavanger region.

3.2 Research design and identification challenges

This section presents our research design. We discuss the empirical setting and underlying assumptions, and provide a section on robustness tests to address potential challenges to the identification.

Table 3.1: Summary statistics

	Mean	Median	Std. dev.	Num. obs.
Panel A: Consumption (in 2012Q4) in NOK				
Food and beverages	22,970	20,834	14,805	205,853
Furnishing	6234	3429	10,338	205,853
Vehicles	3857	0.00	24,883	205,853
Total consumption	99,444	86,601	66,853	205,853
Panel B: Household income and balance sheet variables				
Home value (primary)	2,704,311	2,454,016	1,308,990	205,853
Loan to value (LTV)	0.90	0.48	167.69	205,853
Debt to income (DTI)	2.48	2.20	29.08	205,852
Liquid wealth to income (LTI)	6.96	3.12	13.96	205,852
Panel C: Demographics				
Average age in the household	46.56	45.00	11.10	205,853
No. individuals in gov. sector in household	1.42	1.00	0.63	205,853
Greater Stavanger dummy	0.05	0.00	0.22	205,853

Note: This table presents summary statistics for the data used in the main analysis. Food and beverages, furnishing, and vehicles refer to the NOK amount spent on those goods. Total consumption is the NOK amount spent on total consumption, including further categories in addition to food and beverages, furnishing, and vehicles. The loan to value (LTV) is defined as the debt of a household divided by the value of the household's primary housing. Debt to income is total debt to total income for that household. Liquid wealth to income is defined as bank deposits multiplied by 12 and divided by total income.

3.2.1 Experimental setting

In the ideal experimental setting, the econometrician would observe one-off, random, unpredictable variations in housing wealth at the household level to identify its dynamic effect on consumption. Our approximation to this ideal scenario is to exploit the variation in housing wealth induced by the oil price collapse of 2014. As documented and explained in Chapter 2, the depression in the petroleum sector due to the 2014 price collapse led to big drops in local home prices. In Figure 3.1 we plot the evolution of home price indices in Greater Stavanger—the most exposed region to the oil price collapse—and other areas of the country. We also plot the distribution of the changes in home prices at the household level in Figure

Figure 3.2: Distribution of change in home prices

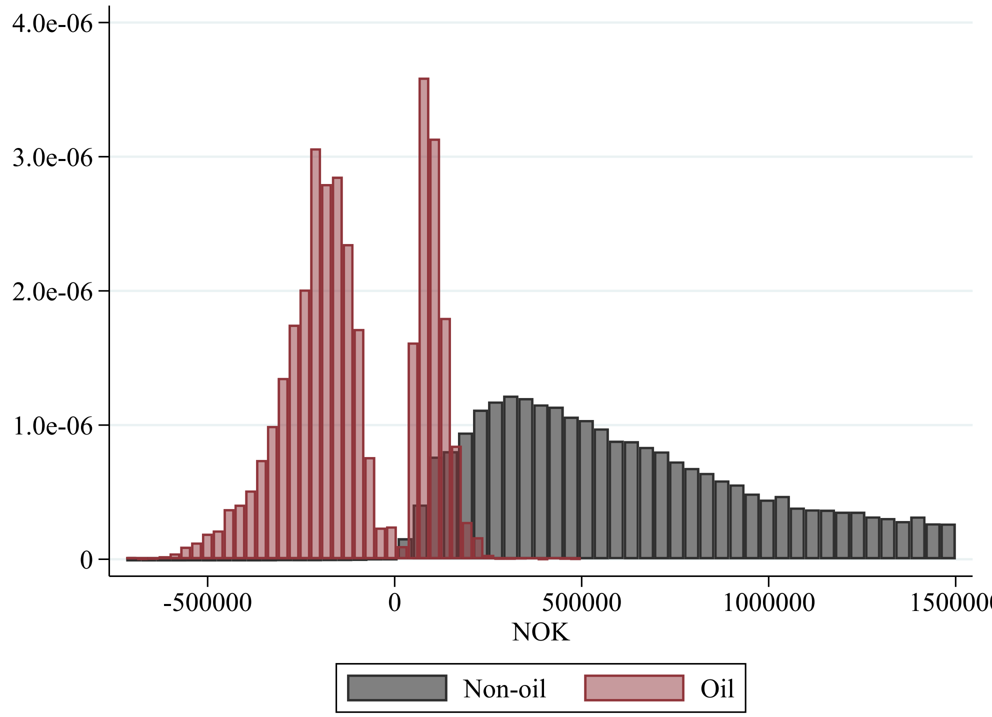

Note: This figure shows the distribution of the home price decline at the household level. The red bars depict home price growth for households in the oil region. The grey bars show home price growth for households living outside the oil region.

3.2. This variation in home prices that materializes in 2015 and 2016 is what we exploit, as further described in Section 3.3.1. For identification, we restrict attention to households where the majority (> 50%) works in the government sector.

We focus on government sector workers because the oil price shock—if it affects them at all—should have a uniform impact on their current and prospective earnings across labor markets. In Norway, government jobs are highly secure against unemployment risk, and wages are determined through national-level collective bargaining between labor and employer organizations. Although a significant portion of national fiscal spending is funded by a special tax on petroleum activities, the distribution of these funds is not influenced by a location's exposure to the petroleum industry. Consequently, any changes in the current and prospective earnings of government workers should be uniform across the country. However, a government worker's housing wealth depends on local housing demand, which varies with the heterogeneous exposure to the petroleum industry.

3.2.2 First test of assumptions

The first test of our assumptions is to regress government workers' home values, total spending, and labor income in Greater Stavanger versus the rest of Norway, following 2014. We estimate

$$y_{i,t} = \beta \, \text{post}_t \times \text{oil region}_i + \eta \, X_{i,t} + \varepsilon_{i,t}, \tag{3.1}$$

where $y_{i,t}$ is the logarithm of either household i's home value, total consumption, or labor income in year t. The vector of controls $X_{i,t}$ contains household mean age and size, and fixed effects for households, housing markets, and years. The results are presented in Table 3.2. We observe a significant drop in relative home values of -0.18 log points. Total consumption also drops, albeit, the change is not significant, or, the distribution of the change is too wide. But, as of now, the important point is that labor income is not changing differently across regions. The estimated coefficient in column (3) is -0.0060, with a standard error of 0.016. This soundly rejects that government workers experience different changes in earnings when contrasting between the Greater Stavanger region and the remainder of Norway.

3.2.3 Threats to identification

The key identifying assumption is that the oil price shock only affects government worker expenditures through its effect on home prices and not through other location-specific channels. In this section, we outline two potential violations of this assumption and how we address them.

3.2.3.1 Structural differences between treatment and control

A natural concern is whether government workers are similar in terms of consumption patterns across regions. For instance, to the extent that there are different regional business

Table 3.2: Difference-in-difference model results

	(1) log home value	(2) log tot. consump.	(3) log labor inc.
$\text{post}_t \times \text{oil region}_i$	-0.184^{***}	-0.109^{*}	-0.00604
	(0.0398)	(0.0547)	(0.0156)
Controls	Yes	Yes	Yes
Num. obs.	1,472,945	1,472,945	1,472,945
R-squared	0.916	0.543	0.518
Household FEs	Yes	Yes	Yes
Region FEs	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes

Note: This table presents the estimates of (3.1). Errors are three-way cluster-robust at the level of household ID, housing market, and year. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

cycles, government workers residing in areas that are relatively booming could have higher consumption growth than government workers residing in other areas. This could be due, for example, to earnings being correlated with the local business cycle, or to peer effect (government workers in Stavanger consuming more because their non-government-worker neighbors are consuming more).

To investigate this concern, we start our empirical analysis by estimating a flexible differences-in-differences model where we compare consumption in Greater Stavanger versus the rest of the country both before and after the oil price collapse. To the extent that consumption evolves similarly prior to the oil price collapse, this would suggest that government workers in different areas have relatively similar consumption patterns prior to the oil price collapse. In the data, there is significant within-household seasonality which makes the analysis very noisy. We remove this by standardizing the spending of each type within household and quarter. Mathematically,

$$\frac{c_{i,t}^k - \text{mean}_{i,q(t)}(c_{i,t'}^k)}{\text{std}_{i,q(t)}(c_{i,t'}^k)} = \sum_{\forall t' \neq 2015\text{Q1}} \beta_{t'}^k \times \mathbb{1}(t = t') + \eta\, X_{i,t} + \varepsilon_{i,t}^k \tag{3.2}$$

Figure 3.3: Variation in consumption across time, location, and consumption category

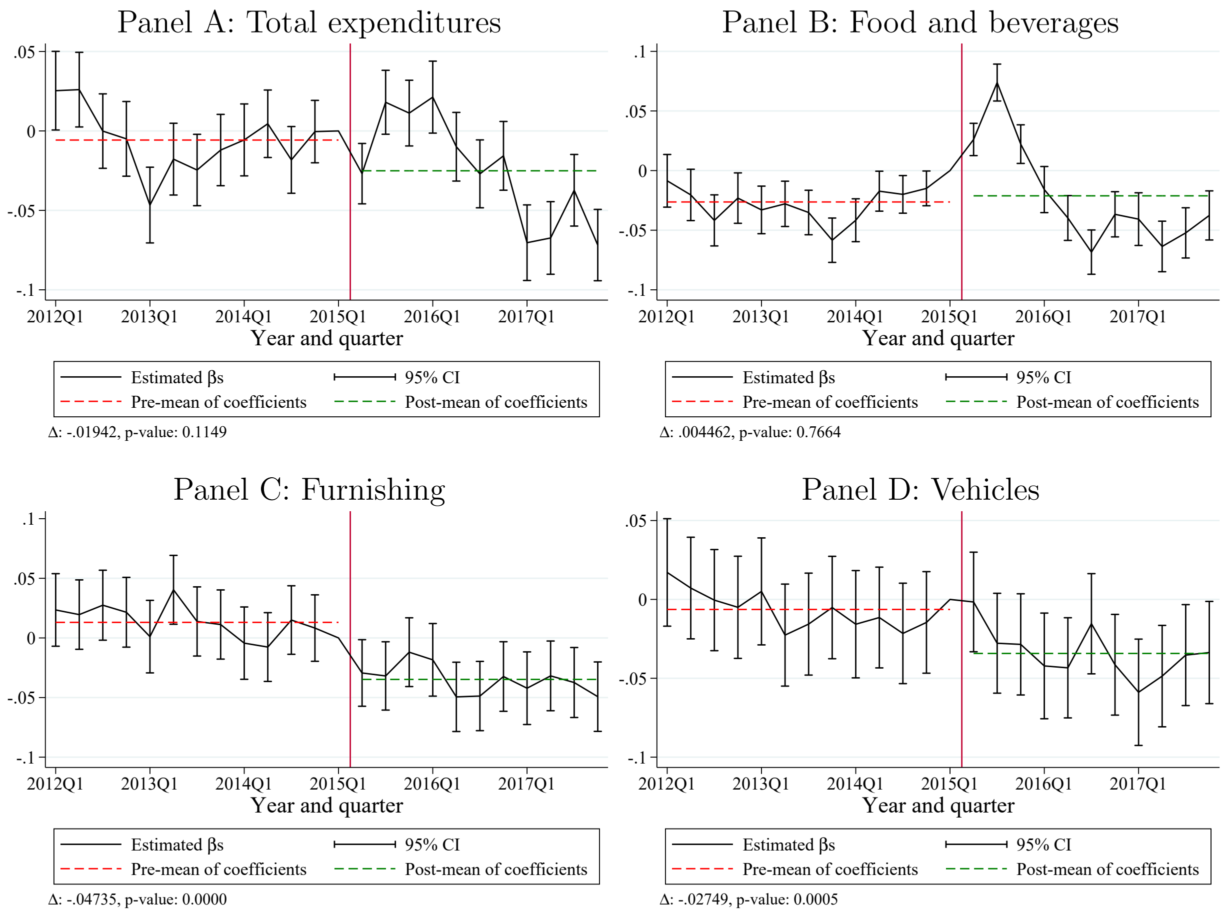

Note: This figure presents the dynamic response of expenditures for different categories of consumer goods estimated using (3.2). The difference Δ and the following p-value under each legend refer to the difference between the red and green dashed lines, indicating the means of the coefficients in the pre- and post-periods, respectively. This is estimated using a difference-in-difference model similar to (3.2) but with the year-quarter dummies replaced by one post dummy.

where $q(t)$ is the quarter of year-quarter t and k indicates the type of good. Using this specification, we can see if the consumption of good k, while accounting for within-household seasonality, varies differently over time across Greater Stavanger and the rest of Norway. The estimates are plotted in Figure 3.3. The vector $X_{i,t}$ contains labor and capital income, gross wealth, liquid wealth, and debt, each standardized and raised to powers from one to four, and household and year-quarter fixed effects. Errors are two-way cluster-robust at the level of household ID and year-quarter.

The differences between the government workers in Greater Stavanger and the rest of Norway do not exhibit any persistent variation in responses to potential unobserved shocks in the pre-period. However, there are significant downward shifts in spending on furnishing and vehicles (Panels C and D) following 2015.

3.2.3.2 Confounding shocks

A second potential violation of the identifying assumption is that, while there are no differences between government workers prior to the shock, there are confounding shocks that differentially affect government workers after the oil price collapse. An example of one such shock could be that income expectations deteriorate for government workers in the oil region relative to government workers otherwise.

To address this issue, we adopt two approaches. First, we condition a model of earnings on a wide set of predetermined observables to ensure that the comparison is based on government workers that are similar based on observables. Second, we do an additional analysis on an alternative sample, namely homeowning retirees, that are unlikely to be affected by the evolution of the local labor markets meanwhile still being exposed to movements in local home prices.

3.3 Results

In this section, we present our main empirical analysis where we estimate and discuss the marginal propensities for expenditures of selected consumption goods. We also perform robustness tests discussed earlier and heterogeneity analysis to separate between the potential channels at work.

3.3.1 IV model

To estimate MPXs out of housing wealth, we want to estimate

$$\Delta c^k_{i,t+\Delta t} = \beta^k \, \Delta hw_{i,t} + \eta \, X_{i,t} + \varepsilon^k_{i,t},$$

where β^k is the marginal propensity for expenditures of good k in period $t+\Delta t$ of a change in housing wealth in period t. Usually, this model is not possible to estimate because of endogeneity between $c^k_{i,t+\Delta t}$ and $hw_{i,t}$ (often attributed to the omitted variables income and wealth expectations). We avoid this by estimating an exogenous change in $hw_{i,t}$ for government workers in different housing markets in Greater Stavanger using a two-stage method akin to the two-sample IV method.

For each household, we observe the value of their primary home at the beginning of 2012 which we denote by $hw_{i,2012}$. Using quarterly home price indices from *Eiendom Norge*, we compute an estimated value of the home in each quarter over 2012–2017. Based on this baseline housing wealth, we then start in the third quarter of 2014, take a 4-quarter window, compute each household's mean primary home value in that window, and denote it by $hw_{i,2012+\Delta t}$. We compute the log-change within household i since 2012 and regress this difference on a dummy if the household lives in an oil-region home market ℓ (see equation (3.3) below). By formula (3.4), we recover the relative change in the local home price index but this approach allows us to also control for household variables that could be correlated with the local change in home prices and drive the change in consumption.

$$\log hw_{i,2012+\Delta t} - \log hw_{i,2012} = \sum_{\ell \in L_S} \delta_\ell \times \mathbb{1}(i \text{ lives in } \ell) + \eta_1 \, X_{i,2012} + \alpha_1 + \varepsilon_{1\,i,\Delta t} \tag{3.3}$$

$$\widehat{\Delta hw}_{i,2012+\Delta t} = (\exp(\hat{\delta}_\ell \times \mathbb{1}(i \text{ lives in } \ell)) - 1) \times hw_{i,2012}. \tag{3.4}$$

Here, L_S denotes all the housing markets in Greater Stavanger as defined by *Eiendom Norge*. The vector of controls $X_{i,2012}$ contains labor income, capital income, gross wealth, liquid wealth, and the household's total debt in 2012, each raised to powers from one to four. All control variables are standardized. The intercept α_1 is the average change in home prices across the country over the considered period excluding the Greater Stavanger markets and is not explained by changes in the characteristics of local government workers. The predicted change in housing wealth estimated in (3.4) is thus the change in housing wealth due to the locations' exposure to the oil price plunge. For households living outside the Greater Stavanger region, the *exogenous* change in housing wealth is zero.

In the second stage, we estimate for each consumption good category k a linear model.

$$c^k_{i,2012+\Delta t} - c^k_{i,2012} = \beta^k_{\Delta t}\, \widehat{\Delta hw}_{i,2012+\Delta t} + \eta^k_2\, X_{i,2012} + \alpha^k_2 + \varepsilon^k_{2\,i,\Delta t}. \tag{3.5}$$

The controls in $X_{i,2012}$ are the same as in the first stage. The interpretation of $\beta^k_{\Delta t}$ is then: for every NOK the household-specific housing wealth rises by between 2012 and $2012 + \Delta t$, expenditures of good category k between the periods 2012 and $2012 + \Delta t$ rises by $\beta^k_{\Delta t}$. The outcome variables are multiplied by four to make the coefficient in terms of per-year instead of per-quarter.

We compute the test statistics by bootstrapping and account for within-housing market clustered errors.[5] The whole procedure keeps the benchmark year fixed (i.e., 2012) while moving the end-year forward by one quarter at a time. We start with 2014Q3–2015Q2 and then 2014Q4–2015Q3, etc.

3.3.1.1 Results on the average consumption responses

In Figure 3.4 we show the MPXs of total expenditures and three categories of consumer goods: Food and beverages, Furnishing, and Automobiles. The interpretation is at an annual frequency.

What we do not observe is homeowners' expectations about *future* home prices. In our main specification (3.5), we implicitly assumed that the change in spending is due to the *current* reduction in home prices in Stavanger relative to the rest of Norway. However,

[5] We cannot use standard correction techniques because of the non-linear transformation in the first stage.

Figure 3.4: The dynamic marginal propensity for spending out of housing wealth shock

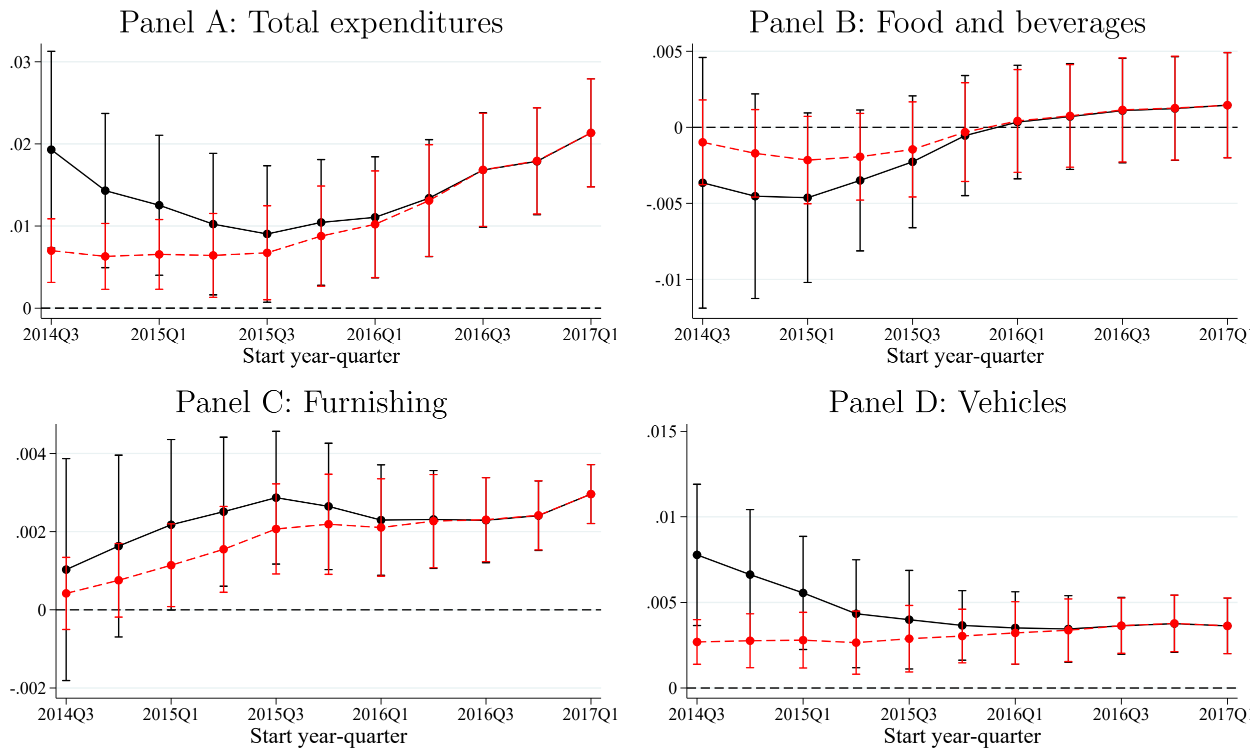

Note: This figure presents the dynamic response of total consumption expenditures for homeowning government workers. Each period refers to a 4-quarter episode that starts in the time on the horizontal axis. The black solid lines display the MPXs estimated using changing home prices, while the red dashed lines display the corresponding estimates when the household-specific housing wealth change is computed over the longest range. The confidence intervals are 90% and are cluster-robust at the housing market level.

these workers are probably internalizing that prices will diverge even further. To study the potential consequences of this, we provide a second set of estimates where the change in housing wealth is held fixed at the longest window we consider (the change over 2012–2017). These are presented in red in Figure 3.4 while the original set of estimates is displayed in black.

Panel A shows that total expenditures fall with decreasing home prices and the magnitude of this change is within the span but in the lower range of previous estimates. We are closest to Disney et al. (2010) and in Chapter 2 (Vestman et al., 2023). The earlier responses are about half the size but more accurately estimated when we use the longer-horizon housing wealth change compared to the current-housing-wealth changes (red versus black). The difference disappears as we look further ahead which is natural since the changes in housing wealth converge.

Food and beverages (Panel B) show no significant change and are if anything of the opposite sign during the first years. As we see that total expenditures are responding, this could indicate a very weak substitution effect away from luxury goods (e.g., restaurant visits) to essential goods.

Panel D shows that a one-NOK fall in home prices causes a 0.08 NOK reduction in car purchases in the beginning and a small but statistically significant drop in expenditures on furnishing (Panel C).

3.3.1.2 Robustness analysis

We perform two robustness checks. Our key assumption is that Stavanger government workers' income prospects do not change relative to other government workers in Norway, but the former group experiences a big unexpected change in housing wealth. If income prospects do change due to the oil price shock, this would bias our results. Another group that is more isolated to income risk is retired homeowners. We define people as retired based on their labor earnings and social benefit payments linked to retirement. We rerun our main specification (3.5) on a sample of retired homeowners and contrast similarly between those living in Greater Stavanger and the rest of Norway. The results are presented in Figure 3.5.

The responses of retirees are also significant and in magnitudes greater than for the main sample of government workers. However, the estimates are not precise enough to tell if the results are significantly different.

Our second robustness test was to test if renters who work in government respond similarly or not. By the assumptions we have made, government workers who rent should be isolated from both income and housing wealth risk. However, we cannot compute an MPX for them since there is no housing wealth change. Thus, we re-estimate the model 3.1 for the total expenditures of renters and present the results in Figure 3.6.

Figure 3.5: The dynamic marginal propensity for spending out of housing wealth shock for retired homeowners

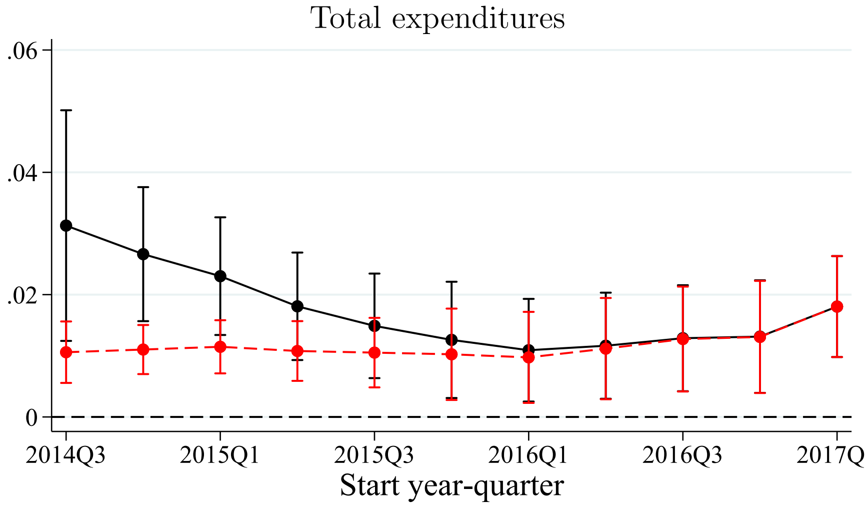

Note: This figure presents the dynamic response of total consumption expenditures for homeowning retirees. Each period refers to a 4-quarter episode that starts in the time on the horizontal axis. The black solid lines display the MPXs estimated using changing home prices, while the red dashed lines display the corresponding estimates when the household-specific housing wealth change is computed over the longest range. The confidence intervals are 90% and are cluster-robust at the housing market level.

The figure displays an insignificant downward shift in total expenditures. The p-value is big, about 0.45, but the sample is also much smaller. However, the results are in line with our assumptions being valid, that renters who do not experience any changes in earning prospects are also not affected by home prices.

Figure 3.6: Variation in consumption across time for renters

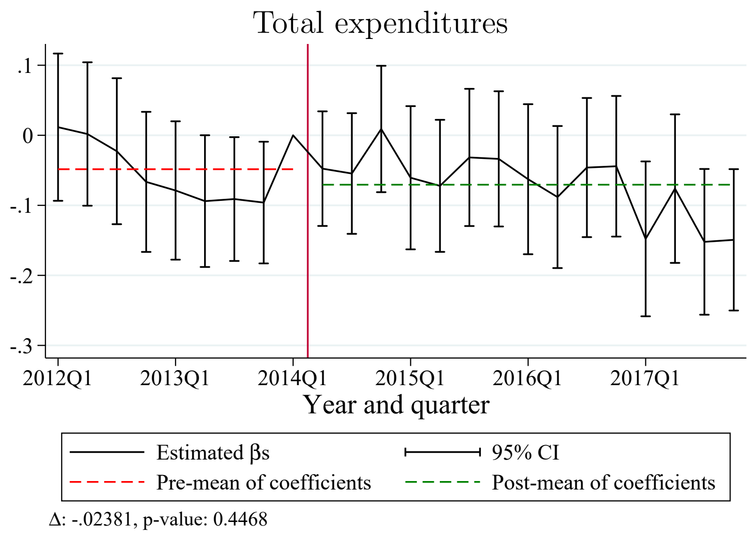

Note: This figure presents the dynamic response of expenditures for total expenditures for renters. Each period refers to a 4-quarter episode that starts in the time on the horizontal axis. The confidence intervals are 90% and are cluster-robust at the housing market level. The Δ and the following p-value under each legend refer to the difference between the red and green dashed lines, indicating the means of the coefficients in the pre- and post-periods, respectively.

3.3.2 Why do home prices affect consumption?

Our goal in this section is to investigate which dimensions of heterogeneity that affect the MPC out of the housing wealth shock. Based on theoretical predictions and previous empirical studies we focus on credit- and liquidity constraints, as well as age, as important dimensions of heterogeneity (see, e.g., Aladangady, 2017; Browning et al., 2013; Campbell and Cocco, 2007; Disney et al., 2010; Jappelli and Pistaferri, 2010).

The pure wealth effect (i.e., how the lifetime wealth of homeowner changes with home prices) has been argued to likely be small (Buiter, 2010). In theory, the housing wealth can only contribute to consumption through a sale that produces a positive net value. This value is then discounted to the present, leading to consumption smoothing in the case of no credit constraints. For older homeowners, consumption is then smoothed out over fewer periods, leading to a greater MPC if it is mainly due to the pure wealth effect.

The sign of the wealth effect depends on the *expected* future net sale (Fagereng et al., 2024). However, a home sale is often in conjunction with a home purchase, which reduces the net value of the transaction, and negative in the case of upsizing. Only at downsizing events or moves to cheaper housing markets, can value be extracted. On the other hand, for renters and young households who expect to buy more housing in the future, changes in home prices affect their perceived lifetime wealth.

The collateral effect works through moving the borrowing constraint of homeowners who use their homes as collateral for mortgages. In the case of credit constrained homeowners, small moves in home prices can generate great shifts in the nominal cap on credit. These households, as well as liquidity constrained households, exhibit greater MPCs out of transitory income shocks as well as home price shocks (Berger et al., 2018; Lustig and Van Nieuwerburgh, 2005; Parker et al., 2013). High indebtedness correlates with being younger, and thus, if the collateral channel is stronger than the pure wealth effect, we expect younger homeowners to respond stronger than older.

Our procedure for disentangling these channels is to estimate differences in MPCs by household types. We construct dummy variables that split our sample population into several groups and include these dummy variables as an interaction term with our main explanatory variable. That is, in the second stage, we estimate for each consumption good category k,

$$\begin{aligned} c^k_{i,2014} - c^k_{i,2012} &= \sum_{g \in \mathcal{G}} \beta^k_{g,\Delta t} \, \widehat{\Delta hw}_{i,2012+\Delta t} \times z_{g(i),2013} \\ &\qquad + \eta^k_2 \, X_{i,2012} + \alpha^k_2 + \varepsilon^k_{2\,i,\Delta t}. \end{aligned} \tag{3.6}$$

Compared with (3.5) the difference is thus the inclusion of $z_{g(i),2013}$s.[6] In our preferred specifications, $z_{g(i),2013}$, are time-invariant dummy variables splitting the households into

[6]Note that, since $z_{g(i),2013}$ is time-invariant, it is absorbed by the household fixed effects and should therefore be excluded from the regression.

equally-sized groups based on the value of the heterogeneity variable in 2013, that is, before the housing shock.

We construct four groups of dummy variables, and run regression models for each of them. The groups consist of two measures of credit constraints—debt-to-income (DTI) and loan-to-value (LTV)—in addition to age (of oldest household member), and liquidity, measured as the sum of bank deposits divided by after-tax income (per month), (i.e. liquidity-to-income (LTI)). Households are then assigned to one group.[7] For LTV, we split up the sample in intervals of 10%, i.e., the group with the lowest LTV has a value in the range of 0–0.1, the second-lowest group is in the range of 0.1–0.2, and so forth. DTI is split up by 0.05, age by 6 years, and LTI by units of 2.5.

The estimates of (3.6) are presented in Figure 3.7 and display rich heterogeneity in the response. First of all, as debt-to-income rises, so does the MPX. For homeowners with the least DTI, the effect is zero, and it then rises almost linearly with DTI. For the most indebted group, the total MPX is 0.04, twice that of the baseline. Also for LTV we see an increasing effect. The effect peaks at an LTV of 85%, coinciding with the legal cap on mortgages. After this point, the effect is if anything diminishing. Liquidity-to-income displays the opposite relationship, inline with the importance of credit: as LTI increases, the effect falls. For very high values of LTI, the effect seems to rise but it is inexactly estimated. The heterogeneity in age is inaccurately estimated. The middle of the sample shows significant effect (35–60) and older zero or negative.

All our results are in line with an important role for housing as collateral to finance consumption. Homeowners who are likely credit constrained are responding the most while older homeowners do not respond—i.e., the pure wealth effect is small.

We provide further evidence that credit plays an important role in understanding the heterogeneity. First, we construct a variable *new* $debt_{i,t}$ which is one if total debt has increased by more than 10% and 20,000 NOK (approximately 1000 USD) since previous year and zero

[7]Different measurements and splits of the samples produce overall the same qualitative conclusions. Recall that, in Norway, since 2010 borrower-based requirements have limited households' maximum DTI and LTV. In 2013 these requirements were DTI less than five times gross income, and LTV to less than 80 percent.

Figure 3.7: Heterogeneity in the MPXs

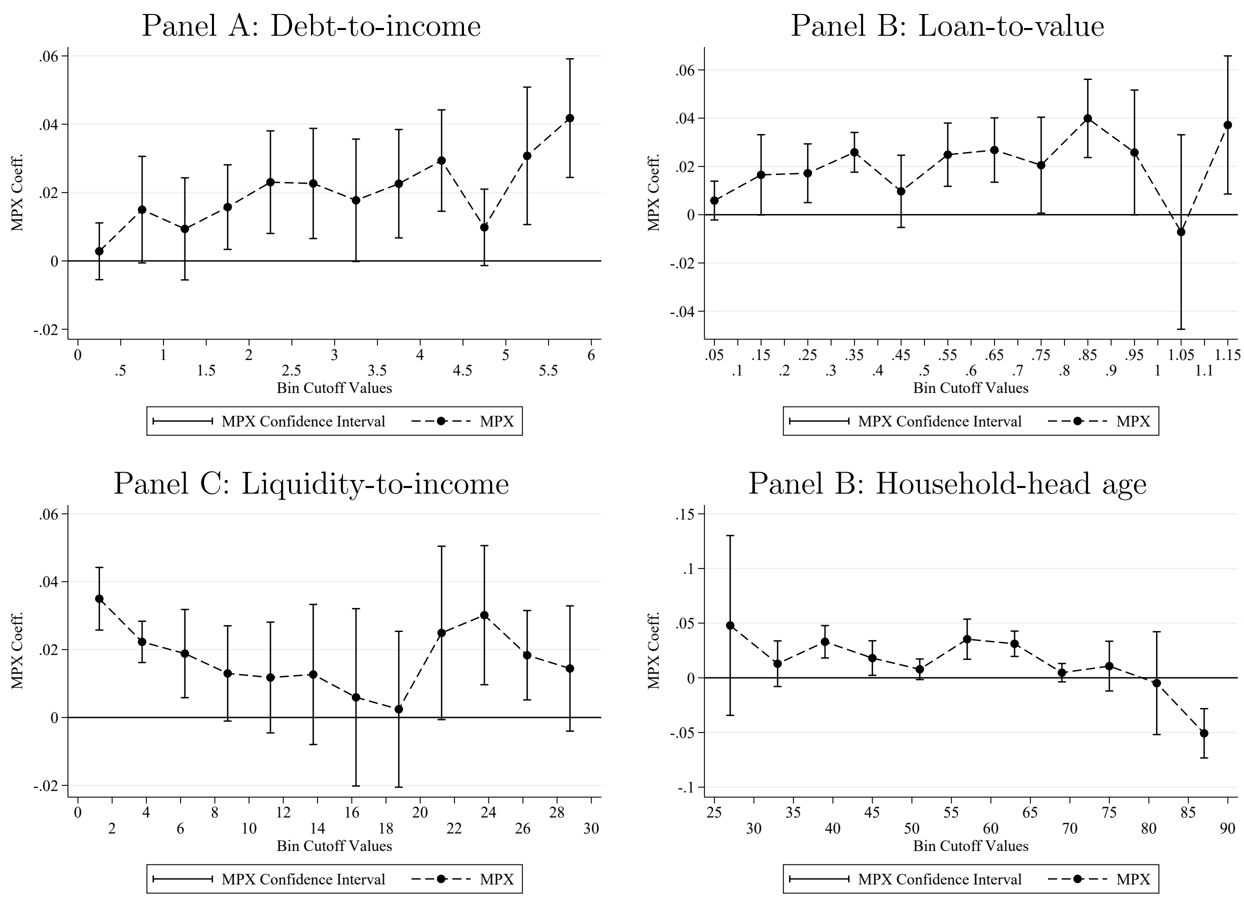

Note: This figure shows the heterogeneity of the total MPX for different household characteristics estimated using (3.6). Errors are cluster-robust at the level of the housing market and the confidence intervals are of 95%.

Figure 3.8: Relative propensity to borrow

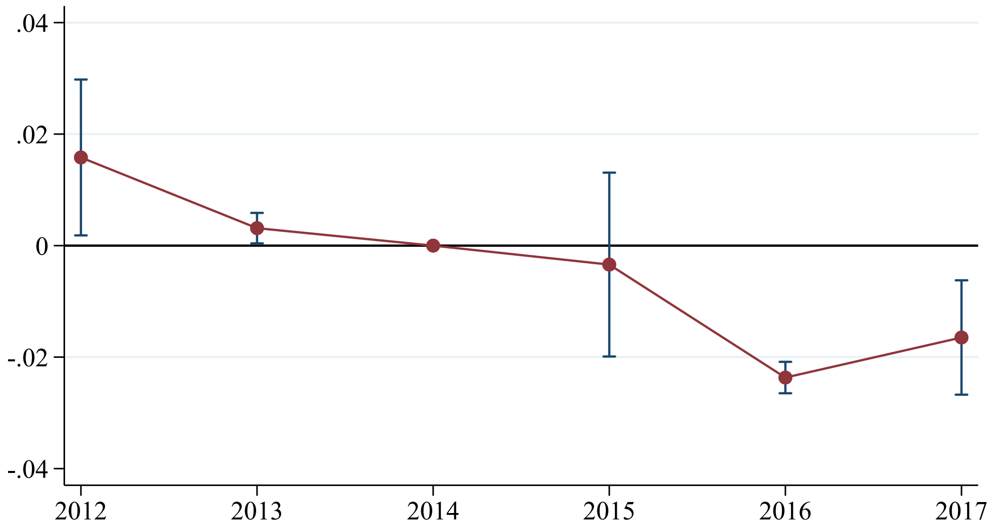

Note: This figure shows the relative propensity to take up a new loan for affected vs. non-affected households estimated using (3.1).

otherwise (similar to Bhutta and Keys, 2016). Second, we estimate a model like (3.1) but replacing the post_t dummy with year dummies (excluding 2014). The estimates are presented in Figure 3.8 and exhibit a significant drop in the increase in debt following 2014, albeit, no significant effect in 2015.

3.4 Conclusion

The question of how homeowners' expenditures are affected by shocks to housing wealth is a question of great importance for policy. In a world were spending responds strongly to swings in home prices, policy makers might want to respond to it, otherwise, they can focus their attention to other economic shocks hitting the economy. In this paper, we empirically show that the fall in home prices in the region of Greater Stavanger in Norway reduced the spending on furniture and vehicle purchases among government workers. Food and beverages are unaffected. We focus on government workers because of the institutional setting that their wages are bargained at the national level by unions and employers, and their job security is very high. We show that this holds in the data and perform robustness tests using renters and retired homeowners to support the validity of our identifying assumptions.

We also perform heterogeneity analysis and find support that the collateral channel dominates the pure wealth channel. Homeowners with higher debt-to-income, loan-to-value, and lower liquidity-to-income and of lower age respond stronger per NOK reduction in home prices. This explains also why more durable spending is more affected. These are often credit financed and possible to postpone purchases of in case of tighter credit conditions.

Our findings indicate that changes in home prices by themselves do not affect spending as much as previously thought. Also, the goods that are mostly affected have small impact on the local economy due to often being imported. The consequences of this correlation is a topic for future research.

3.A Appendix: Comparison of measured expenditures and survey data

Table 3.A.1: Comparison to consumer surveys

	Total consumption	Vehicles	Furnishing	Food & beverages
Microdata	397,776	15,428	24,936	91,880
All Norway[a]	435,507	43,226	24,495	63,146
Rural (*spredtbygd*)[b]	404,862	57,742	24,224	65,185
Urban (not Olso, Bergen, and Trondheim)[c]	445,068	42,282	25,962	64,175
Agder & Rogaland[d]	440,794	44,203	26,950	64,417
Two-person households[e]	466,821	50,358	27,694	69,030

Note: This table compares the expenditures in the microdata versus survey data by Statistics Norway. The wave used here is 2012 and split up by different characteristics. The expenditure categories do not map exactly to our definitions. Here, "vehicles" refer to consumption group 071 ("vehicle purchases") which also includes used vehicles, "furnishing" refers to consumption group 05 ("furniture, household items, and maintenance"), and "food and beverages" is the sum of groups 01 and 02 ("groceries and alcohol free beverages" and "Alcohol beverages and tobacco").

[a] Table 10235, *Forbruksundersøkelsen*, 2012
[b] Table 10238, *Forbruksundersøkelsen*, 2012
[c] Table 10238, *Forbruksundersøkelsen*, 2012
[d] Table 10237, *Forbruksundersøkelsen*, 2012
[e] Table 10240, *Forbruksundersøkelsen*, 2012

References

Aladangady, Aditya (2017) “Housing wealth and consumption: Evidence from geographically linked microdata,” *American Economic Review*, 107 (11), 3415–3446.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra (2018) “House prices and consumer spending,” *The Review of Economic Studies*, 85 (3), 1502–1542.

Bhutta, Neil and Benjamin J Keys (2016) “Interest rates and equity extraction during the housing boom,” *American Economic Review*, 106 (7), 1742–1774.

Browning, Martin, Mette Gørtz, and Søren Leth-Petersen (2013) “Housing wealth and consumption: A micro panel study,” *The Economic Journal*, 123, 401–428.

Buiter, Willem H (2010) “Housing wealth isn’t wealth,” *Economics: The Open-Access, Open-Assessment E-Journal*, 4.

Campbell, John Y. and Joao F. Cocco (2007) “How do house prices affect consumption? Evidence from micro data,” *Journal of Monetary Economics*, 54, 591–621.

Disney, Richard, John Gathergood, and Andrew Henley (2010) “House price shocks, negative equity, and household consumption in the United Kingdom,” *Journal of the European Economic Association*, 8 (6), 1179–1207.

Eggertsson, Gauti B and Paul Krugman (2012) “Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach,” *The Quarterly Journal of Economics*, 127 (3), 1469–1513.

Fagereng, Andreas, Matthieu Gomez, Emilien Gouin-Bonenfant, Martin Holm, Benjamin Moll, and Gisle Natvik (2024) “Asset-price redistribution,” Working paper, Available online.

Fagereng, Andreas, Martin B Holm, and Gisle J Natvik (2021) “MPC heterogeneity and household balance sheets,” *American Economic Journal: Macroeconomics*, 13 (4), 1–54.

Jappelli, Tullio and Luigi Pistaferri (2010) “The consumption response to income changes,” *Annu. Rev. Econ.*, 2 (1), 479–506.

Juelsrud, Ragnar E and Ella Getz Wold (2019) “The saving and employment effects of higher job loss risk,” Working Paper 17/2019, Norges Bank.

Kaplan, Greg and Giovanni L Violante (2014) “A model of the consumption response to fiscal stimulus payments,” *Econometrica*, 82 (4), 1199–1239.

Lorentzen, Linnea (2023) “Domino effects: Understanding sectoral reallocation and its wage implications,” Working Paper, University of Oslo.

Lustig, Hanno N and Stijn G Van Nieuwerburgh (2005) "Housing collateral, consumption insurance, and risk premia: An empirical perspective," *The Journal of Finance*, 60 (3), 1167–1219.

Mian, Atif, Kamalesh Rao, and Amir Sufi (2013) "Household balance sheets, consumption, and the Economic Slump," *Quarterly Journal of Economics*, 128 (4), 1687–1726.

Parker, Jonathan A, Nicholas S Souleles, David S Johnson, and Robert McClelland (2013) "Consumer spending and the economic stimulus payments of 2008," *American Economic Review*, 103 (6), 2530–2553.

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

Vestman, Roine, Jesper Bojeryd, Björn Tyrefors, and Dany Kessel (2023) *The housing wealth effect: Quasi-experimental evidence*: Centre for Economic Policy Research.